aws elemental

CONFIGURATION GUIDE

AWS ELEMENTAL CONDUCTOR LIVE 3

AWS Elemental
1320 SW Broadway
Portland, Oregon, 97201

+1 503 222 3212
www.elemental.com

This guides applies to AWS Elemental Conductor Live 3, version 3.7.

Contents

About This Manual

This guide is intended for engineers who are performing the initial configuration on one or more AWS Elemental nodes that are in an AWS Elemental Conductor Live 3 cluster. In other words, this guide describes how to set up Conductor Live 3 nodes and worker nodes (AWS Elemental Live nodes and optional Statmux nodes) in a "Conductor cluster."

Phase 2 of Installation

This guide provides detailed information on phase 2 of installation:

- Configure all the nodes into a cluster so that they can be controlled by AWS Elemental Conductor Live 3.
- Enable user authentication so that users must log in to use any product.
- Add users if user authentication is enabled.
- Configure the time zone, DNS server, NTP servers, and firewall.
- Configure other Ethernet interfaces, as required.
- Configure routers and other input devices.

Pre-requisite Knowledge

It is assumed that you know how to:

- Connect to the AWS Elemental Conductor Live 3 web interface using your web browser.
- Log into a remote terminal session, in order to work via the command line interface.

Note:	To receive assistance with your AWS Elemental appliances and software products, see the forums and other helpful tools on AWS Elemental User Community.

1 PROCEDURES

1.1. Initial Configuration

This section describes the steps to perform to set up an initial cluster of Conductor nodes and worker nodes.

- It is assumed that the Conductor nodes are “brand new.”
- Worker nodes may be new or may have been previously deployed in stand-alone mode.

1.1.1.1. Getting Ready

Networks

We strongly recommend that, if your deployment involves several Conductor Live 3 clusters, you set up each cluster in its own network.

CPU-only versus GPU-enabled

Some worker nodes in the cluster can be running GPU-enabled versions or some can be running CPU-only versions of the AWS Elemental software: Conductor Live 3 can manage a “mixed architecture” cluster.

Redundant Conductor Configuration

You can configure the cluster to operate with two Conductor nodes in a redundant configuration. One Conductor node is the primary and is active and managing the cluster. Another Conductor node is the backup; it is also active but it is not managing the cluster; instead, it is continually keeping its database synchronized with the primary node.

1. Find out if you should set up for Conductor redundancy.
2. If you are in a redundant configuration, find out which of the Conductor hardware units is the primary node and which is the backup. Make a note of the IP addresses of the two nodes.

Note:	In this guide, we assume you have a redundant configuration. If not, treat references to “primary Conductor node” as references to “the Conductor node” and ignore instructions for the backup Conductor node.

Redundant Worker Nodes

You can configure the worker nodes to operate in a redundant configuration. Worker nodes are organized into groups, with some nodes active and doing encoding or muxing, and other nodes in a “passive reserve” state.

1. Find out if you should set up for AWS Elemental Live redundancy and/or Statmux redundancy.
2. Identify the hardware units that will initially be active and those that will be passive reserve. Make a note of the IP addresses.

User Authentication

You can configure the entire cluster so that users and administrators must log on to view or work with the web interface of a node or work with the REST API of a node.

User authentication is an all or nothing proposition: either all users must log on to all nodes for access by all methods (web interface or API) or no users must.

1.2. SSL Configuration

SSL (Secure Socket Layer) enables the secure version of HTTP (HTTPS) and encrypts communications between the client and server. When configuring SSL on your cluster, note that:

- Every node in the cluster must match the Conductor's SSL setting.

  If the Conductor has SSL enabled but a worker does not (or vice versa), the worker will experience errors and/or service impairment.
- For existing clusters being reconfigured with SSL on or off, all nodes must be de-clustered, reconfigured, and then added back to the cluster.

  To avoid service interruption, we recommend that SSL configuration on existing nodes be completed in a maintenance window.
- Once SSL is enabled, you must continue to use the --https command with subsequent node reconfigurations.

  If you omit the command, SSL is disabled.
- If you are enabling user authentication but do not use the --https option in the command, the installer will only proceed after you have acknowledged the warning that encryption will not be used when transmitting usernames and passwords.

1.2.1.1. Summary of Initial Configuration

Action	Page
Configure Network Features on each Worker Node	
☐ Access the web interface on the worker node	9
☐ Verify that licenses are installed on the worker node	10
☐ Configure the time zone for the worker node	14
☐ Configure network devices (Ethernet devices) for the worker node	14
☐ Set up additional DNS servers and NTP servers for the worker node, if required	19
☐ Configure the firewall for the worker node	20
☐ Provide the worker node with access to remote servers: mount points	31
☐ Configure the worker node for RTMP inputs	38
☐ Configure SNMP for the worker node	42
Set up Network Features on Primary Conductor Node	
☐ Access the web interface on the primary Conductor node	9
☐ Verify that licenses are installed on the Conductor node	10
☐ Configure the time zone for the Conductor node	14
☐ Configure network devices (Ethernet devices) for the Conductor node	14

- ☐ Set up additional DNS servers and NTP servers for the Conductor node, if required 19
- ☐ Configure the firewall for the Conductor node 20
- ☐ Provide the Conductor node with access to remote servers: mount points 31
- ☐ Configure SNMP for the Conductor node 42

Set up Network Features on Backup Conductor Node

- ☐ Access the web interface on the backup Conductor node 9
- ☐ Verify that licenses are installed on the Conductor node 10
- ☐ Provide the Conductor node with access to remote servers: mount points 31

Configure the Cluster

- ☐ Access the web interface on the primary Conductor node 9
- ☐ Add worker nodes to the cluster 21
- ☐ Set up for Conductor redundancy 23
- ☐ Set up for worker redundancy 24
- ☐ Configure for SDI direct inputs on worker nodes 32
- ☐ Configure for SDI router inputs on worker nodes 32
- ☐ Configure for database backup 39
- ☐ Enable user authentication as needed 52
- ☐ Create users for the cluster (as required) 56
- ☐ Address role policies (if PAM authentication is enabled) 62
- ☐ Enable Conductor redundancy. This must be the last step! 65

1.2.1.2. Summary of Changing SSL Settings

Enable or Disable SSL

- ☐ Access the web interface on the primary Conductor node 9
- ☐ Remove all nodes from the cluster 22
- ☐ Run configuration to enable/disable SSL on all nodes 9
- ☐ Add nodes back to the cluster 21

2 TASKS

1.4. Configuring SSL

All nodes in the cluster must have the same SSL setting (on or off).

Perform this procedure on	Conductor nodes and Worker nodes

You can optionally enable SSL on the node in order to secure traffic over the communications layer. Traffic that goes over this layer includes traffic that uses the HTTP protocol. By default, the AWS Elemental product is configured with SSL disabled.

If you enable SSL, then all traffic over the communications layer must use a secure protocol. If you disable SSL, then all traffic must use the unsecured version of the protocol. Traffic that uses the wrong version of the protocol will fail.

SSL configuration is accomplished through the Linux command line interface.

1.4.1.1. Warning

Once SSL is enabled, then every time you enter one of the following commands (to change some other aspect of the configuration), then you must always include the --https option:

- The run command (for example, sudo sh ./elemental_production_conductor_file_2.8.n.nnnnn.run).
- The configure command (sudo ./configure).

If you have enabled SSL and omit the –https option, you will disable SSL.

1.4.1.2. Enabling SSL

SSL configuration must be the same throughout the cluster.

1. Remove all Conductor and worker nodes from the cluster, as described on page 22.
2. Enable SSL on all nodes:
   - At your workstation, start a remote terminal session to the primary Conductor node.
   - At the Linux prompt, log in with username “elemental” and the default password (if not changed by admin). Appendix A. Default Password Information on page 68 for more information.
   - Change to the directory where the configuration script is located:

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

   - Run the configuration script as follows:

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

     Where:

     --https enables SSL

Note:	If you run this command when SSL is in fact already enabled, nothing will change in either the SSL configuration or any other aspects of the configuration.

- The configuration prompts appear. At each prompt, accept the suggestion in order not to change other aspects of the configuration.
- Perform these steps through a remote terminal session with the secondary Conductor, as well as at each worker node.

3. Rebuild the cluster, as described on page 21.

1.4.1.3. Disabling SSL

Remove all nodes from the cluster, then run the configure script (above) *without* the --https option. SSL will be disabled. When SSL is disabled on all nodes, return them to the cluster.

If you run this command when SSL is in fact already disabled, nothing will change in either the SSL configuration or any other aspects of the configuration.

1.5. Accessing the Web Interface

Perform this procedure on	Conductor nodes and Worker nodes

At your workstation, open a web browser and enter the IP address of the desired node. For example:

```
10.4.136.90
```

1.6. Verifying Licenses

Make sure that licenses have been installed on the nodes. The licenses to look for depends on the type of deployment.

2.1.1 Node-locked Licenses

Perform this procedure on	Both Conductor nodes and worker nodes

In a node-locked license deployment, each node has its own license.

To check the license on each Conductor and worker node:

1. From the web interface for the node, choose Settings >Licenses. Make sure the screen looks like the relevant screen below.
2. If a license is missing and you installed on physical hardware units, follow the procedure in these guides:
   - Installing AWS Elemental Conductor Live 3 with Node-locked License – Quick Guide
   - Installing AWS Elemental Live with Node-locked License – Quick Guide
   - Installing AWS Elemental Statmux with Node-locked License – Quick Guide

   Or if you installed on VMs:
   - Node-locked License Deployments on a VM – Install Guide

Note:	Conductor redundancy – High Availability (HA) – must be disabled in order to install licenses. If performing a first-time configuration, install licenses before you enable HA (page 65). If you are installing licenses after initial configuration, you must first disable HA (page 66).

Conductor Nodes

Make sure that each Conductor node has:

- conductor.lic

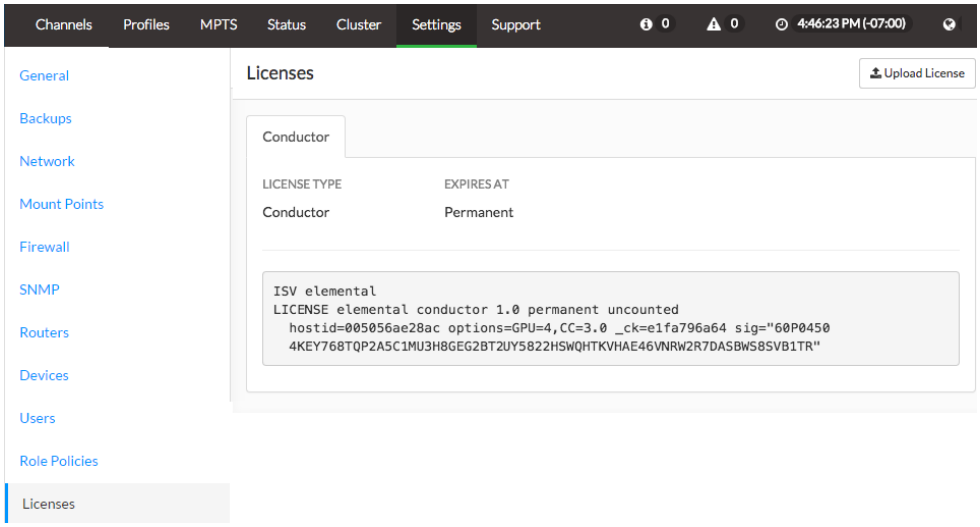

Worker Nodes

Make sure that each worker node has:

- eme.lic

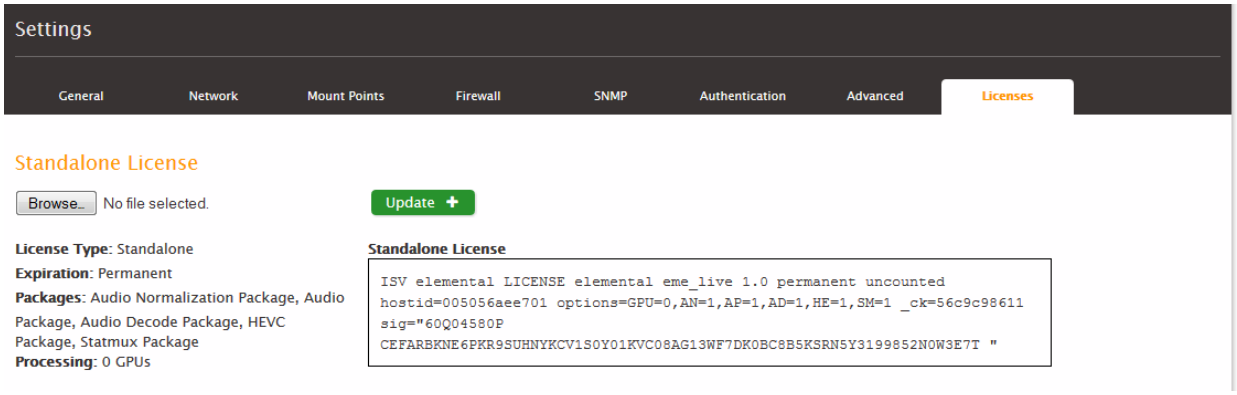

2.1.2 Pooled Licenses

Perform this procedure on	Both Conductor nodes

In a pooled license deployment, the Conductor node acts as a “license server” for the worker nodes.

To check licenses on each Conductor node:

1. From the web interface for the primary Conductor node, choose Settings>Licenses. Make sure the screen looks like the following.
2. If a license is missing, follow the procedure in:
   - Pooled License Deployments on a VM - Install Guide

Conductor Nodes

Make sure that each Conductor node has:

- conductor.lic
- pooled.lic

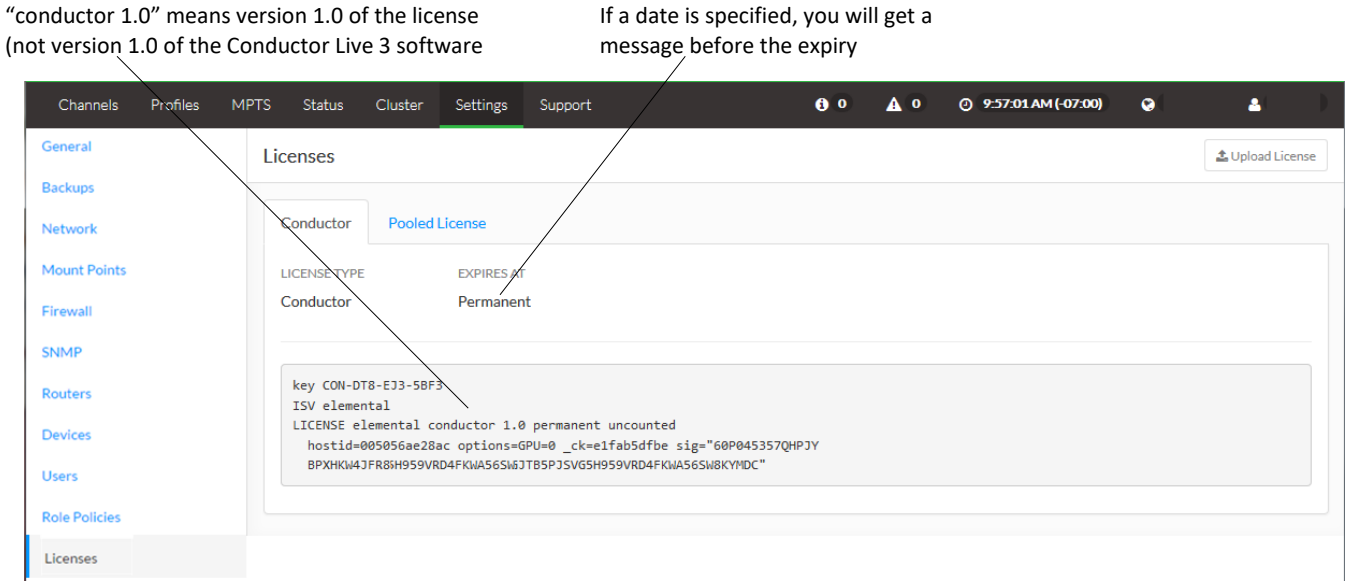

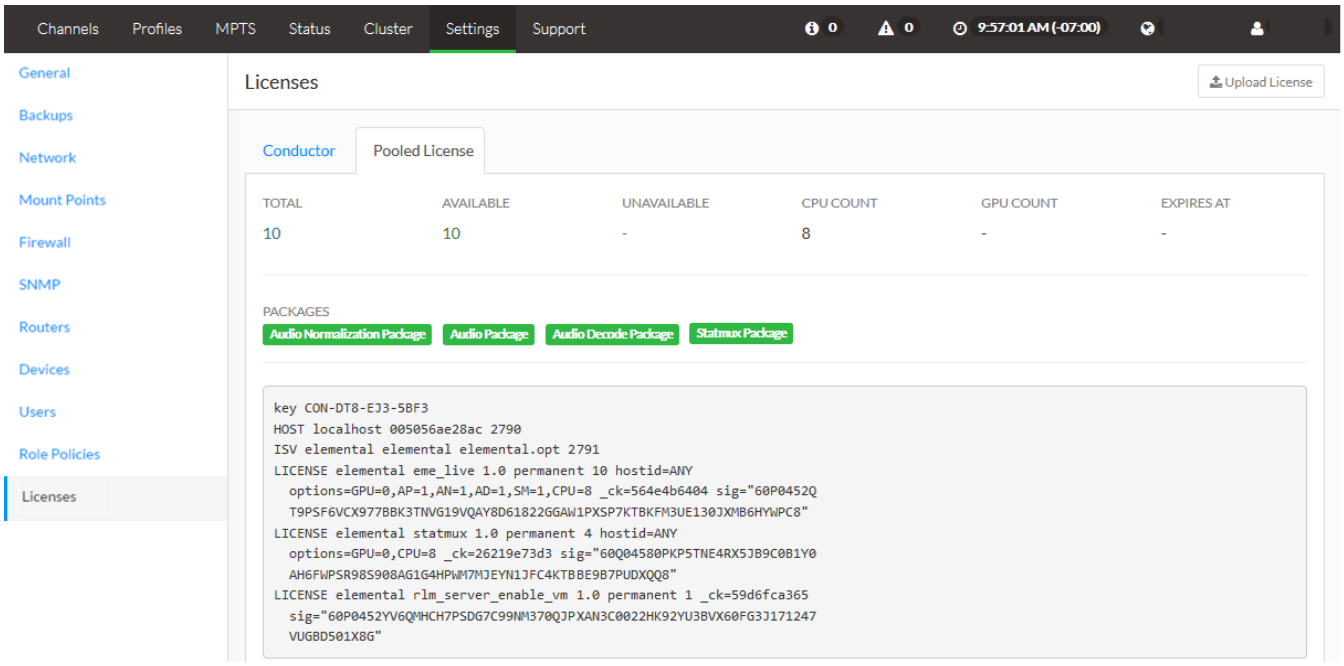

Worker Nodes

The worker nodes do not need to have a license installed; instead, when the node is added to the cluster (page 21), the primary Conductor node will provide it with an eme.lic.

1.7. Configuring the Timezone

Perform this procedure on	The primary Conductor node and worker nodes

Follow this procedure if you did not set the time zone when installing (via the -t prompt) or if you want to change the time zone.

1. From the web interface for the node, go to the Settings>General screen and set the time zone from the dropdown menu.
2. Click Update.

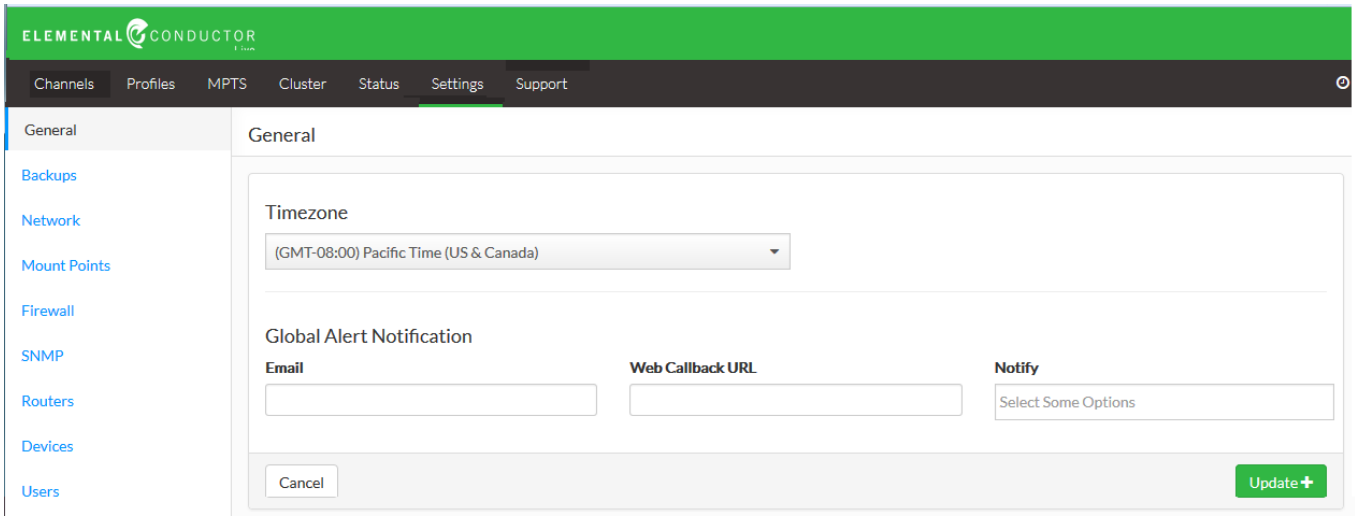

The Conductor web interface will show all activity with a timestamp for the specified time zone.

This setting does not affect:

- Time reporting on the individual AWS Elemental Live nodes: you should have set the time zone for each node when you installed AWS Elemental Live.
- Activity via SSH or via the REST API.

1.8. Configuring Network Devices

Perform this procedure on	The primary Conductor node and worker nodes

When you installed Conductor Live 3, eth0 would have been automatically set up. You may also have chosen to set up eth1 via installation. If you did not do so, you can set it up on the Settings>Network screen.

In addition, you can optionally set up bonding.

The possible network configuration possibilities are therefore:

- Set up eth0 and eth1 with no bonding
- Set up eth0 and eth1 with bonding0

2.1.3 Ethernet Devices with No Bonding

Note:	Conductor redundancy – High Availability (HA) – must be disabled in order to configure network devices. If performing a first-time configuration, configure devices before you enable HA (page 65). If you are configuring network devices after initial configuration, you must first disable HA (page 66).

The eth0 device was automatically set up during installation. If you also set up eth1 at that time, no further configuration is required. If you did not set up eth1, do so now.

The following describes the procedure for the Conductor node. The procedure and screens for a worker node are nearly identical.

1. From the web interface for the node, go to the Settings>Network screen.
2. In the Network Devices section, click Add Device and choose Ethernet. The Add New Network Device dialog appears.
3. Complete the fields as follows:

Field	Description
Device Name	Select “eth1.”
Description	Optional.
Master	Should specify “No devices with port bond settings available.” This wording indicates that you have not created a bond-type device, so bonding is not available.
Address Mode	Select DHCP or Static or None. If you choose Static, then IP Address, Netmask, and Gateway fields appear.
Management	Typically unchecked because eth0 is usually set up as the management interface.
IP Address, Netmask, Gateway	The fields appear only if you set Address Mode to Static.
Static Routes	Optional.

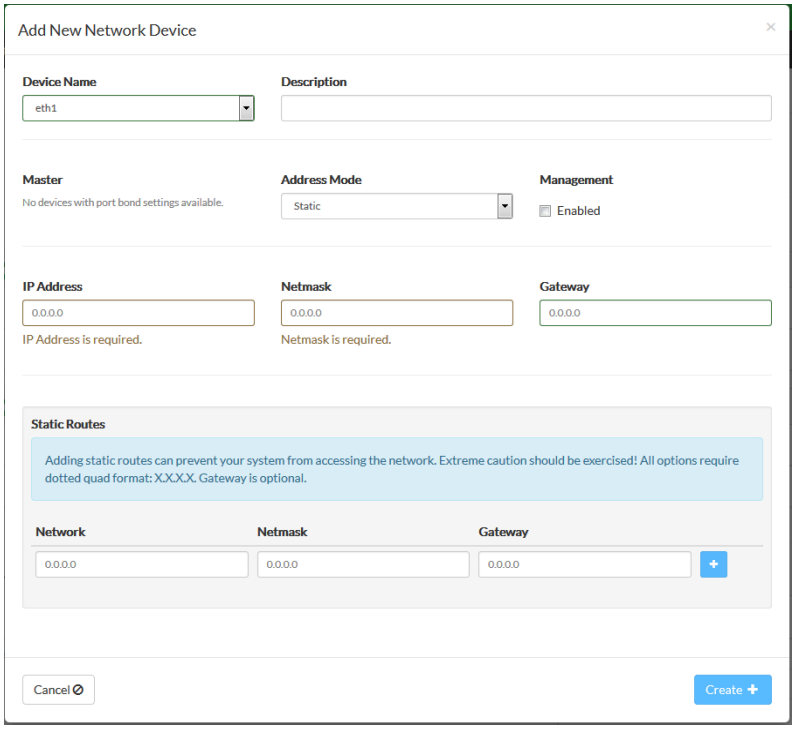

4. Click Create. The new device appears in the Network Devices section.

2.1.4 Ethernet Devices with Bonding

Note:	Conductor redundancy – High Availability (HA) – must be disabled in order to configure network devices. If performing a first-time configuration, configure devices before you enable HA (page 65). If you are configuring network devices after initial configuration, you must first disable HA (page 66).

Note:	We recommend that when setting 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.

Step A: Set up eth0

1. The eth0 device was automatically set up during installation (see the relevant install guide).
2. If you did not set up eth1 at that time, do so now, as described above.
   - Because you are creating a bond, the eth0 and eth1 devices should be set up with static IP addresses.
   - The eth0, eth1, and bond0 devices should all be on the same subnet.

Step B: Create bond0

1. In the Network Devices section, click Add Device and choose Bonded. The Add New Network Device dialog appears, specifying fields appropriate to a bond.

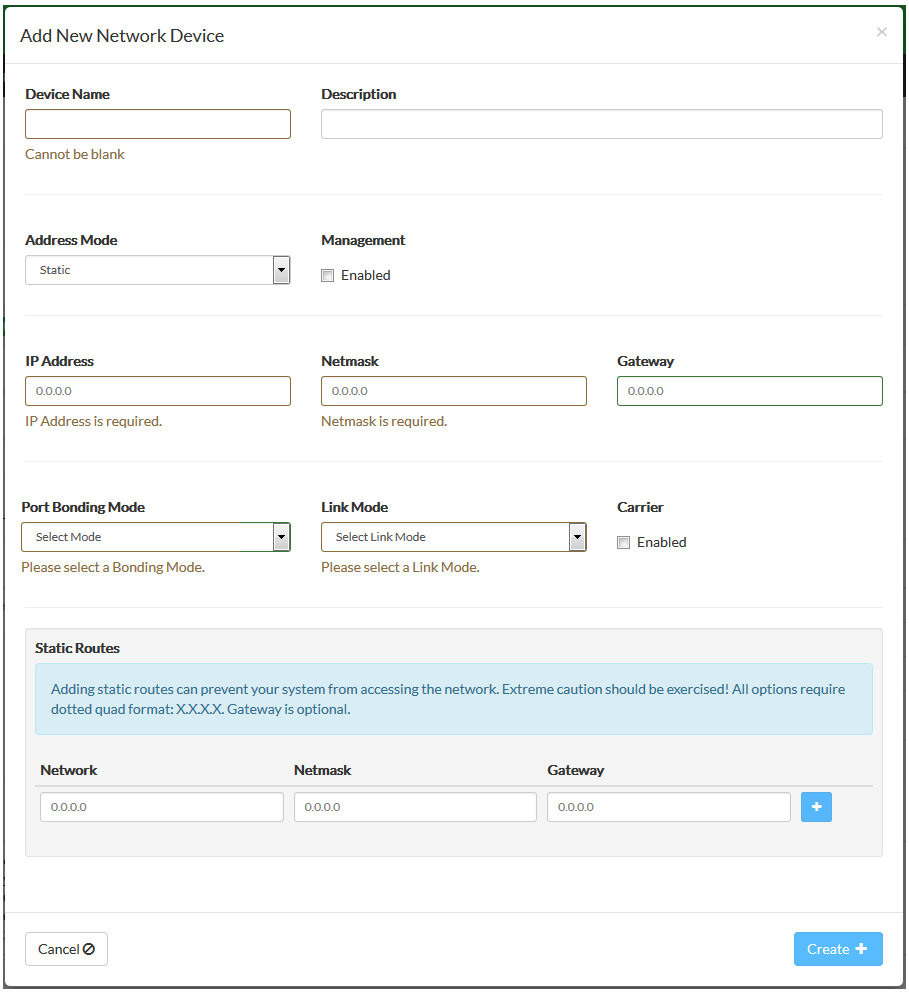

2. Complete the fields as follows:

Field	Description
Device Name	Specify “bond0.”
Description	Optional.
Address Mode	Select DHCP or Static or None. We recommend that you select Static for a bond. If you choose Static, then IP Address, Netmask, and Gateway fields appear.
Management	Typically unchecked because eth0 is usually set up as the management interface.
Port Bonding Mode	Choose the desired mode. See the table on page 17. If you are setting up bonding in order to mitigate false failovers in the case of output listening (for AWS Elemental Live and/or AWS Elemental Statmux and as described in the user guides for those products), then choose mode 1 (Active-Backup).
Link Mode	Choose the appropriate mode. Different supplementary fields appear for each mode. See the table on page 18.
Carrier	Check if appropriate.
IP Address, Netmask, Gateway	The fields appear only if you set Address Mode to Static. The eth0, eth1, and bond0 devices should all be on the same subnet.
Static Routes	Optional.

3. Click Create. The bond appears in the Network Devices section.

Bonding Modes

Mode ID	Mode	Description
0	Round Robin	Sets a round-robin policy for fault tolerance and load balancing. Transmissions are received and sent out sequentially on each bonded slave interface, beginning with the first one available.
1	Active Backup	Sets an active-backup policy for fault tolerance. Transmissions are received and sent out via the first available bonded slave interface. The other bonded slave interface is only used if the active bonded slave interface fails.
2	Balanced XOR	Sets an XOR (exclusive-or) policy for fault tolerance and load balancing. Using this method, the interface matches up the incoming request's MAC address with the MAC address for one of the slave NICs. Once this link is established, transmissions are sent out sequentially, beginning with the first available interface.
3	Broadcast	Sets a broadcast policy for fault tolerance. All transmissions are sent on all slave interfaces.
4	IEEE 803.ad Dynamic Link Aggregation	Sets an IEEE 802.3ad dynamic link aggregation policy. Creates aggregation groups that share the same speed and duplex settings. Transmits and receives on all slaves in the active aggregator. Requires a switch that is 802.3ad compliant.
5	Adaptive Transmit Load Balancing	Sets a Transmit Load Balancing (TLB) policy for fault tolerance and load balancing. The outgoing traffic is distributed according to the current load on each slave interface. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed slave.
6	Adaptive Load Balancing	Sets an Active Load Balancing (ALB) policy for fault tolerance and load balancing. Includes transmit and receive load balancing for IPV4 traffic. Receive load balancing is achieved through ARP negotiation.

MII Link Mode Fields

Field	Description
MII Monitoring Frequency	Specifies the MII link monitoring frequency in milliseconds. The frequency determines how often the link state of each slave is inspected for link failures. 100ms is a good starting point.
Down Delay	Specifies the time, in milliseconds, to wait before disabling a slave after a link failure has been detected. Only applies to the MII Link Mode and should be a multiple of the MII Monitoring Frequency (will be rounded to nearest multiple). Defaults to 0.
Up Delay	Specifies the time, in milliseconds, to wait before enabling a slave after a link recovery has been detected. Only applies to the MII Link Mode and should be a multiple of the MII Monitoring Frequency (will be rounded to the nearest multiple). Defaults to 0.
Carrier	Used in conjunction with the MII Link Mode. If checked, then MII will use MII or ETHTOOL io Ctls (less efficient, and uses deprecated kernel calling sequences), instead of netif_carrier_ok. Relies on the device driver to maintain link state.

ARP Mode Fields

Field	Description
ARP Interval	Specifies the ARP link monitoring frequency in milliseconds. Periodically checks slave devices for traffic, generates regular interval traffic via ARP probes for ARP IP Target.
ARP IP Target	Specifies the IP address to use for ARP probes in ARP Link Mode.

Step C: Assign eth0 and eth1 to bond0

1. Click Edit at the far right of eth0. The Manage Network Device dialog appears.

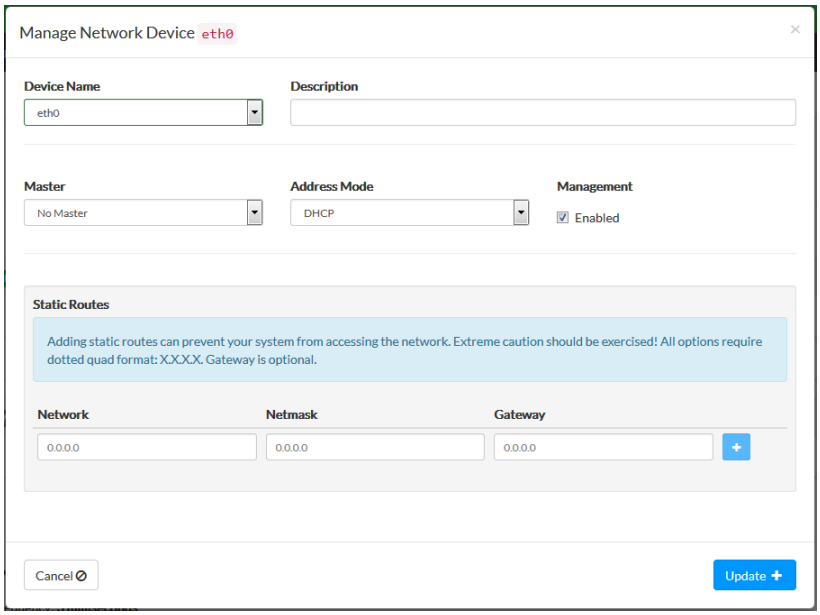

2. In the Master field, choose bond0. (Notice how the old wording has disappeared because a bond is now available.)
3. Click Create. This device now shows with “Master” unchecked to indicate it is bonded to bond0.
4. Repeat for eth1.

1.9. Setting up Additional DNS Servers

Perform this procedure on	The primary Conductor node and worker nodes

The DNS server for the node is usually set up initially during installation. After installation, you can use the web interface to add more servers, as required.

1. From the web interface for the node, choose Settings>Network screen.
2. Click Add DNS in the Domain Name Servers section.
3. Enter the IP address and click the checkmark button to add the address to the list.

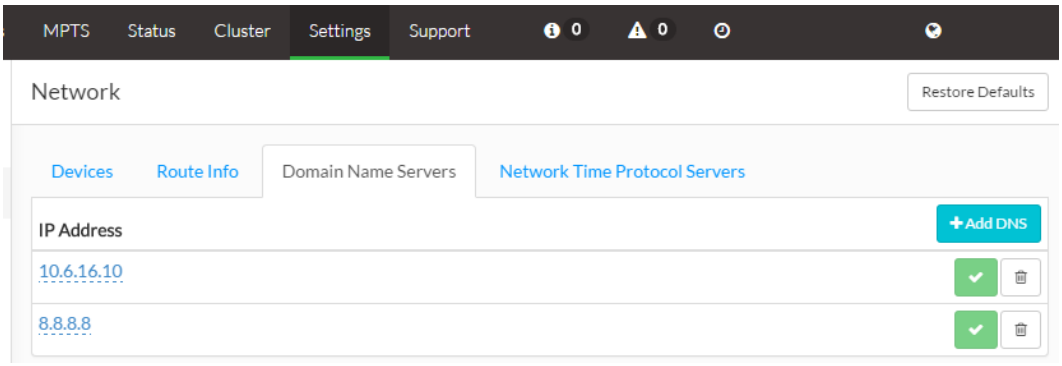

1.10. Setting up Additional NTP Servers

Perform this procedure on	The primary Conductor node and worker nodes

The NTP server for the node is usually set up initially during installation. After installation, you can use the web interface to add more servers, as desired.

1. From the web interface for the node, choose Settings>Network screen.
2. Under the Network Time Protocol Servers tab, click the Add NTP button.
3. Enter the IP address and click the checkmark button to add the address to the list.

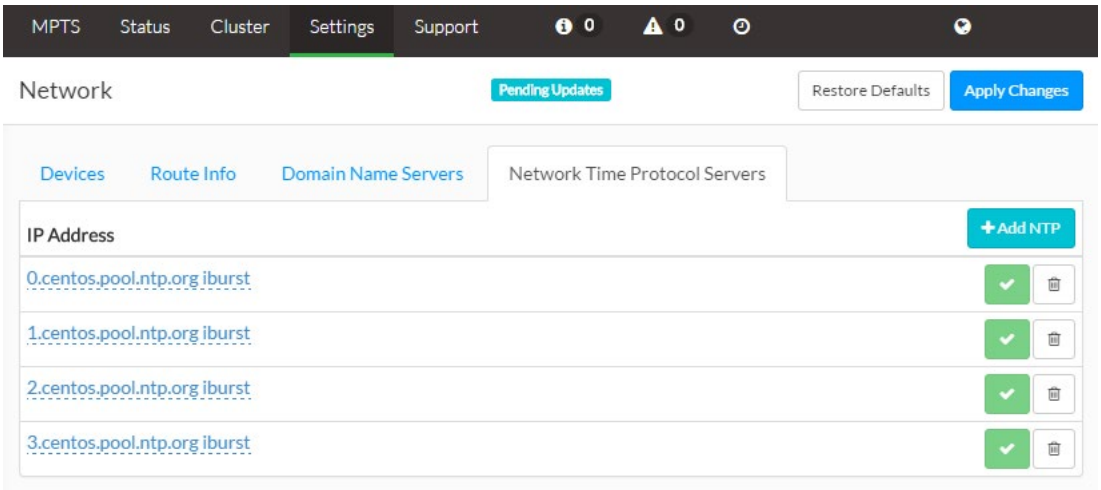

1.11. Configuring the Firewall

Perform this procedure on	The primary Conductor node and worker nodes

This section describes how to enable or disable the firewall around individual hardware units.

Whether the individual firewalls are enabled or not, we recommend that the cluster always be installed behind a customer firewall on a private network.

1.11.1.1. On the Conductor Node

Warning: If you are setting up for Conductor redundancy, you must *disable* the firewall on each Conductor node.

1. From the web interface for the primary Conductor node, choose Settings>Firewall.
2. Make sure that the firewall is disabled. If not, click Stop Firewall.

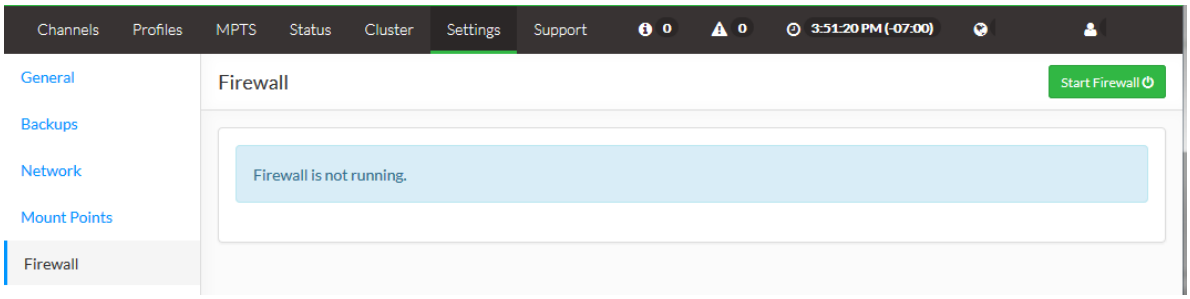

1.11.1.2. On the Worker Node

You can enable or disable the firewall. By default, the firewall is enabled.

The installer configures the ports on your firewall that must be open for incoming and outgoing traffic to and from each node. You can open more ports if required for any reason.

1. From the web interface for the worker node, choose Settings>Firewall.
2. Click Firewall On. A list of ports appears.
3. Add or delete ports as desired.

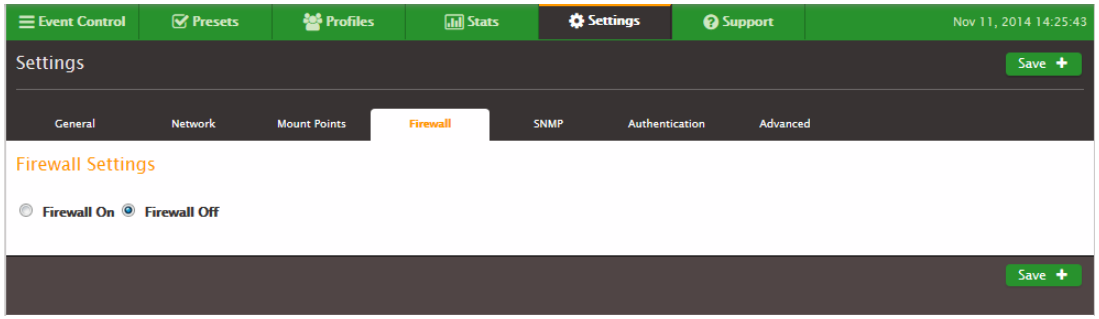

1.12. Setting up Nodes in the Cluster

Perform this procedure on	The primary Conductor node

To set up the cluster and allow the worker nodes to be managed by the Conductor, you must add the worker nodes. If you have a backup Conductor node, you must also add that node.

Warning: You must do this work from the Conductor node that is designated as your (initial) primary node.

1.12.1.1. Adding Nodes to the Cluster

Prior to adding nodes to the cluster, ensure you have addressed SSL configuration. The cluster must be 100% SSL or non-SSL. No mixing of security within the cluster is supported. Errors and/or degradation of service (nodes taken offline) will result if a node does not conform.

1. From the web interface for the primary Conductor node, choose Cluster>Nodes. The Nodes screen appears and will be blank.
2. Click Add Node. The Add Nodes to Cluster dialog appears.

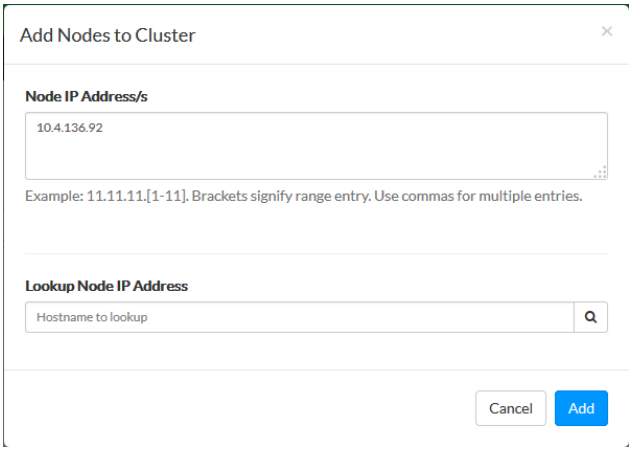

3. Either:
   - Enter the IP address of the node and click Add.
   - Or, if your network has a DNS server, you can do a lookup: type the hostname of the machine in the Lookup Node IP Address field. The hostname is set during installation of the AWS Elemental software and should be in the form “elae-12345678”.

     Click the Search icon. The corresponding IP address will appear below the field. Click the plus icon beside the address to add it to the Node IP Address/s list, then click the Add button.
4. The node is added to the list on the Nodes screen:
   - The Status column will specify “Online.”
   - The product column will specify “Conductor Live” or “Live” or “Statmux”, depending on which node you added.
5. Repeat to add all worker nodes and the backup Conductor node, if applicable.
6. If the node contains SDI cards, you must “import” the devices so that the Conductor node knows about the devices.

Click the downward facing arrow on the far right of the node and select “Import Devices.” The device is detected on the specified worker node and its configuration is automatically added to the Conductor database.

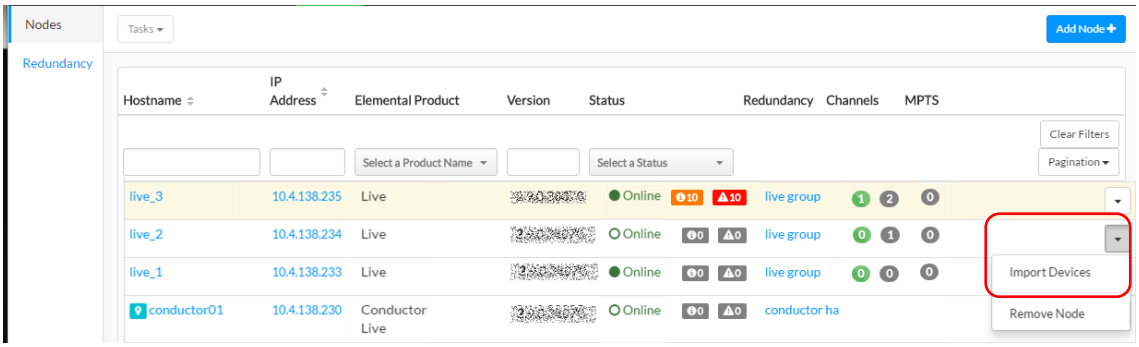

If you do not import devices, they will not appear in the Settings>Devices screen (page 32) and you will not be able to specify these devices as video sources in a channel.

1.12.1.2. Removing a Node from the Cluster

Generally, you remove a node only if you are reconfiguring the cluster for SSL or you are moving a node to another cluster (controlled by another Conductor Live 3). When moving a node, first remove it from its current cluster, then add the node to the new cluster.

1. For an AWS Elemental Live node, make sure there are no channels associated with the node:
   - Display the Channels screen and filter the channels list to the node you want to remove. Make a note of these channels.
   - Either wait for each channel to complete or stop the channel that is running on the node (click the red Stop button beside the channel).
   - Modify each channel so that the node it is associated with is “None”: click the Edit button (pencil icon) and change the node.
2. For either an AWS Elemental Live or Statmux node, make sure there are no MPTS outputs associated with the node:
   - Display the MPTS screen and show the Configuration. Verify which node each MPTS output is using.
   - If an MPTS output is using the node you are deleting, set the MPTS output to use a different node. See the Conductor Live 3 User Guide for details.
3. One the Nodes page, click the Delete button. At the prompt, click Remove.

1.12.1.3. Running a Node in “Stand-alone” Mode

Note that there is no need to remove a node in order to run the AWS Elemental Live node in “standalone” mode (without Conductor Live 3 controlling the activity). You can always run events from AWS Elemental Live, even when the node is part of a cluster. However, note that we do not recommend running events from AWS Elemental Live; see the Overview chapter of the Conductor Live 3 User Guide for more information.

Similarly, there is no need to remove a Statmux node in order to run it in stand-alone mode: you can create and run an MPTS output directly in the Statmux node, even when the node is part of a cluster, but, again, it is not recommended.

1.12.1.4. Moving a Worker Node to another Cluster

To move a node, first remove it from its current Conductor Live 3, then go to the web interface for the new Conductor Live 3 and add it.

1.13. Conductor Redundancy

Perform this procedure on	The primary Conductor node

1. From the web interface for the primary Conductor node, choose Cluster>Redundancy. The Redundancy screen appears and will be blank.

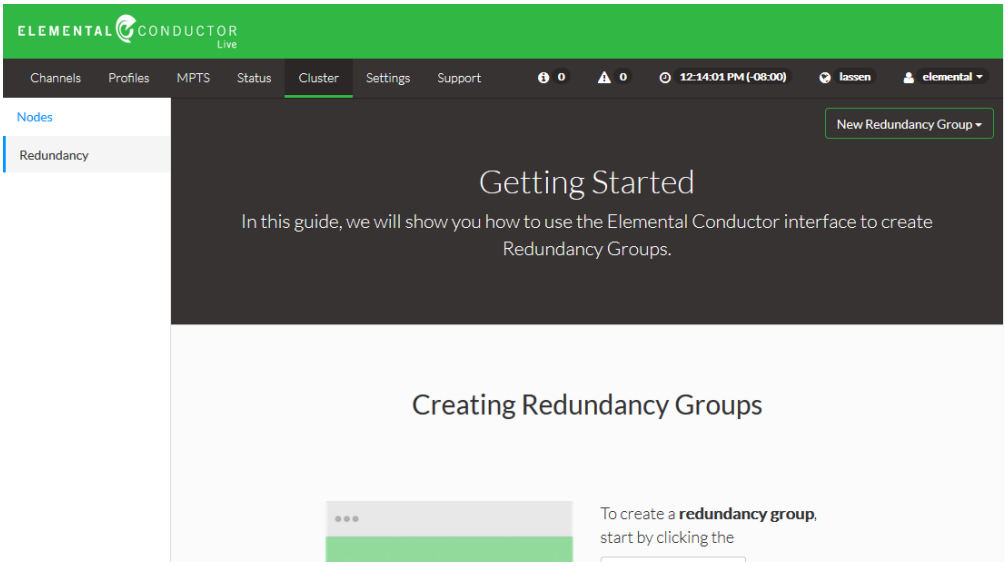

2. Click New Redundancy Group>Elemental Conductor Live. The Add dialog appears.

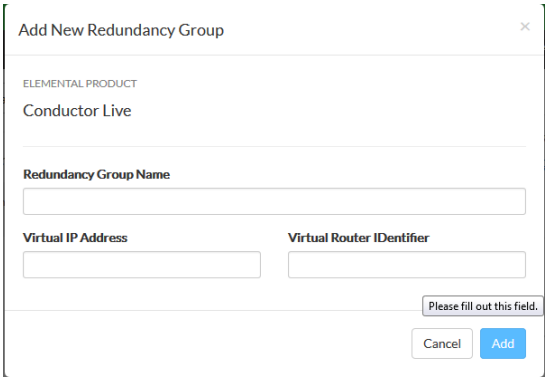

3. Complete the dialog as follows

Field	Description
Redundancy Group Name	Any name, for example, “Conductor”.
Virtual IP Address	A valid IPv4 address. This address must be: • An address on your network that will never be allocated to any other host. • An address on the same subnet as the Conductor nodes. This address serves as the “cluster ID” for the primary and backup Conductor nodes. When Conductor redundancy is enabled (page 65), one of the two Conductor nodes “registers” with this VIP as the primary node. Initially, the node that registers is the one that you are working on when you enable Conductor redundancy. So, if you are working on the web interface for the node at 10.4.136.90 when you click Enable HA, then that node automatically gets tagged as the primary Conductor. After that, any time a Conductor failover occurs, the node that is promoted to primary re-registers with the VIP, effectively indicating “I’m the primary now.”

Warning:	You have now set up for Conductor redundancy, but do not click the Enable HA button yet. You will click this button as the very last step in deploying the cluster, page 65.

1.14. Worker Redundancy: Setting up Nodes for Failover

Perform this procedure on	The primary Conductor node

Worker node failover is implemented in the Conductor Live 3 cluster by creating a "redundancy group" in the cluster, assigning each node to one redundancy group, and assigning a role to each node – active or reserve. Active nodes run activity, reserve nodes are backups that are idle until an active node fails.

Any nodes in the cluster that are not in a redundancy group will have their failure detected but will not fail over.

Multiple Redundancy Groups

You can create more than one redundancy group:

- You must put AWS Elemental Live and Statmux nodes in separate redundancy groups.
- You can optionally create more than one group in order to have one set of reserve nodes to service one set of active nodes and another set to service another set of active nodes. For example, one redundancy group for half of the AWS Elemental Live nodes and another group for the other AWS Elemental Live nodes.

Failover on AWS Elemental Live Nodes

- An AWS Elemental Live node can fail over to a reserve node that is running AWS Elemental Live or running Elemental Live with the Statmux option.
- An AWS Elemental Live node with the Statmux option can fail over to a reserve node that is running AWS Elemental Live or running AWS Elemental Live with the Statmux option. It cannot fail over to an AWS Elemental Live node or to a Statmux node. However, any channels on the failed node that involve Statmux MPTS outputs will not restart on the failover node.
- Given these two rules, if you have both AWS Elemental Live nodes and AWS Elemental Live with the Statmux option nodes in the cluster, you should set up redundancy groups so that the reserve nodes are running AWS Elemental Live with the Statmux option.

Failover on Statmux Nodes

If your deployment includes standalone Statmux nodes and you want to set them up for node failover, you must create a redundancy group consisting of all the active Statmux nodes and at least one reserve node.

Warning: The Statmux nodes must not be put in the same redundancy group as the AWS Elemental Live nodes.

Redundancy Types

The distribution of active vs reserve determines the redundancy type for the entire redundancy group. (There is no idea of assigning a different redundancy type to each node in the group.) The types are:

- One-to-one. There is one active node and one backup node in the group.
- Many-to-one. There are several active nodes but only one backup node in the group.
- Many-to-many. There are several active nodes and several backup nodes in the group.

The current redundancy type can be viewed in Conductor Live 3 when you have finished assigning a role to each node.

SDI Direct Input and Failover

If your deployment includes SDI direct inputs (page 32), failover will occur if a node fails and it has SDI inputs, so long as the reserve node also has SDI inputs and preferably the same number of SDI inputs.

Warning: To avoid problems with failover, ensure the reserve node is using the same SDI card make and manufacturer as the active nodes.

SDI Router Inputs and Failover

If your deployment includes SDI router inputs (page 32), failover will occur if a node fails and it has SDI inputs, so long as the reserve node also has SDI inputs (preferably the same number of SDI inputs) and those inputs have been configured on the Router screen (page 32).

Assuming this setup has been performed, then when the active node fails and node_Y (for example) takes over, node_Y requests the input from the router. Assuming that the router has a path to enough SDI input on node_Y, then the router will be able to fulfill this request.

2.1.5 Initial Setup

This procedure applies to setting up a redundancy group of either AWS Elemental Live nodes or Statmux nodes.

Step A: Select Backup Hardware Units

- You must make sure that a backup node (called a “reserve” node) is as powerful as the most powerful active node. If it is not and that most powerful active node fails, then it is possible that some of the moved channels will fail to restart.

  We recommend that all the nodes in a cluster use the same hardware model to ensure that they all have the same capacity.

  This uniformity recommendation includes having the same number of SDI cards in each hardware unit, particularly if your deployment includes a router handling the SDI input.
- The reserve node does not need to be more powerful than the most powerful active because a reserve node never takes on the work of two (or more) failed nodes.
- Keep in mind that the nodes are not dedicated to being either “active” or “reserve” and there is no idea of a node remembering that it is “really” a reserve node or an active node. For example, once a reserve node takes over the work of a failed active node, it becomes an active node itself and has no history of having once been a reserve node.

Step B: Create Redundancy Group

To set up for failover, you set up at least one redundancy group and assign Active and Reserve nodes to that group. You must set up a redundancy group even for one-to-one mode.

From the web interface for the primary Conductor node, choose Cluster > Redundancy.

On the Redundancy screen, click New Redundancy Group. If the cluster contains both Elemental Live and Statmux nodes, a dropdown menu appears. Select the node type.

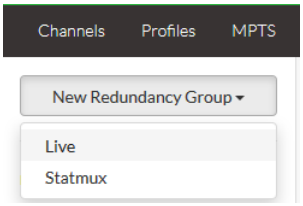

Complete both fields and click Add.

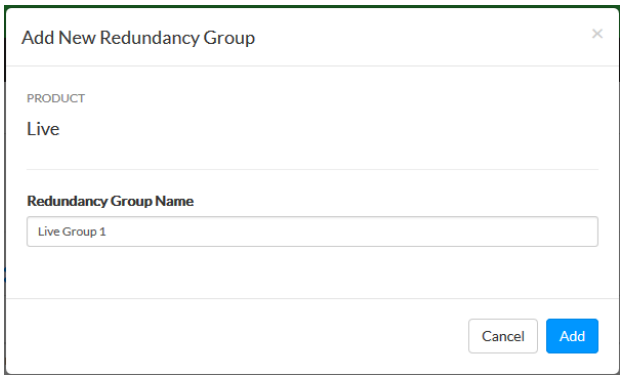

The group is added to the list on the left side of the Redundancy screen

Notice that there are no nodes yet in the group.

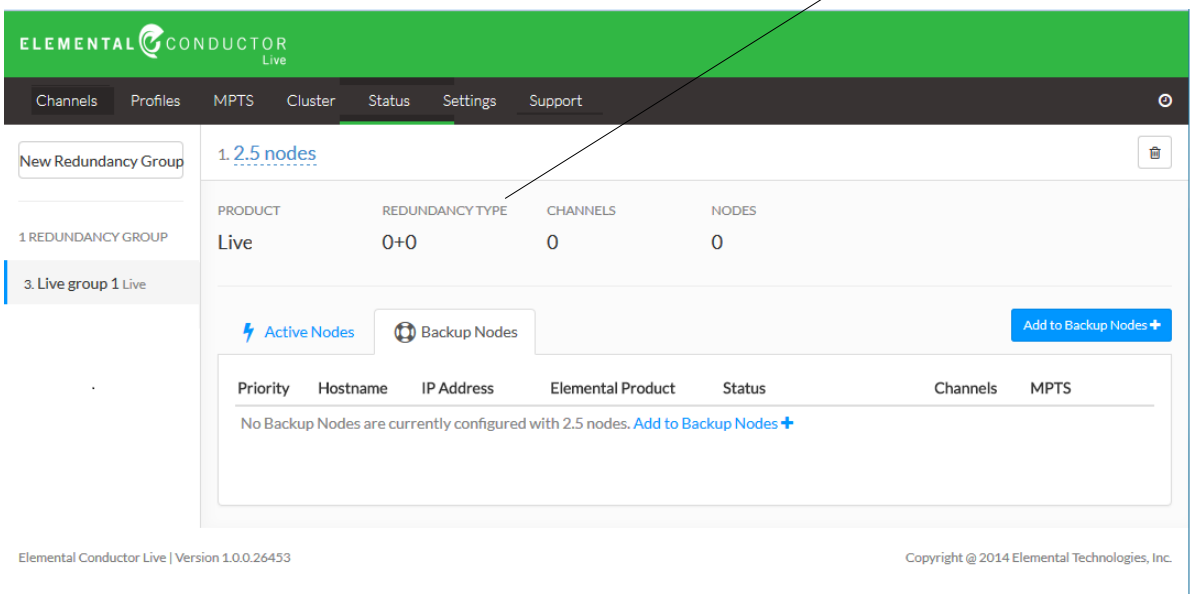

Step C: Add Active Nodes

On the Redundancy screen, select the group where you want to add nodes.

Click Add to Active.

The nodes list shows only nodes that are not currently in a group.

Specifies the channels currently associated with this node. Do not worry if it specifies “0 channels.”

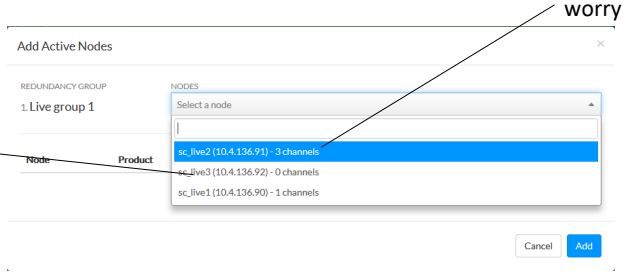

Select a node to add to the group. Only the applicable nodes show (only Elemental Live or only Statmux).

The node is added to the list.

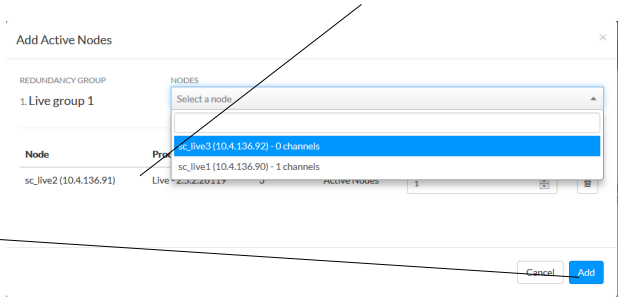

Repeat to add more nodes. Click Add when you’ve added all the active nodes.

The nodes are listed in the Active Nodes tab.

N+0. The “N” indicates that you’ve added more than one active node.

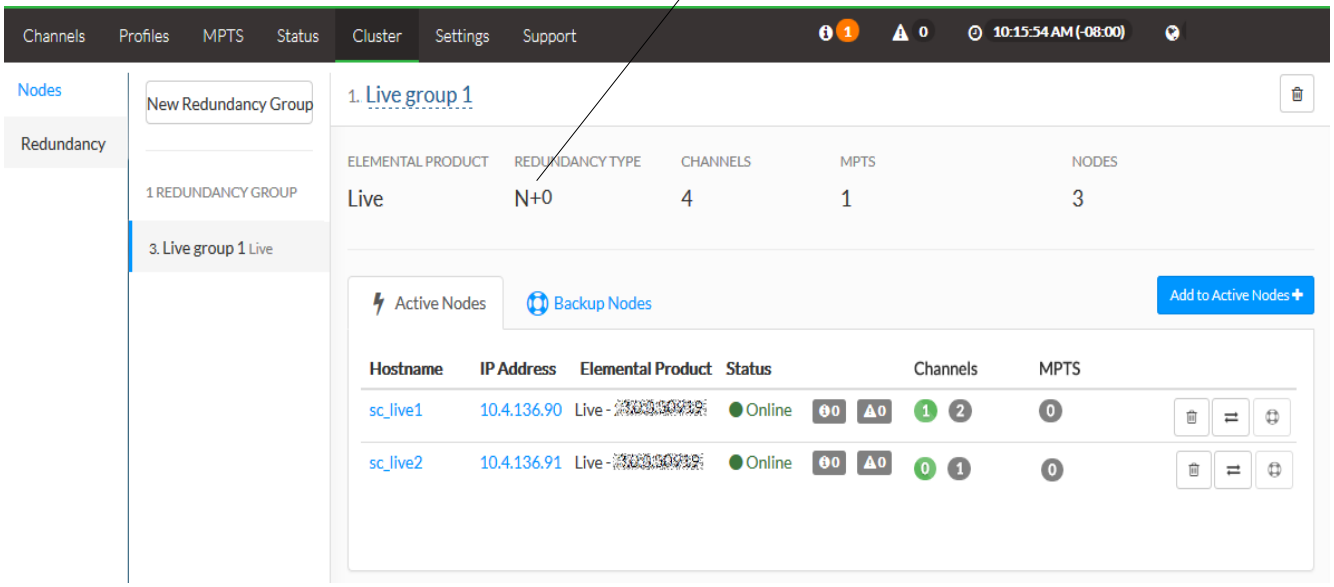

Step D: Add Reserve Nodes

Click the Reserve nodes tab and repeat the steps above to add at least one reserve node.

Note: Make sure you add nodes in order: add the backup node with the highest priority first. The priority of reserve nodes determines which reserve node gets used first in a failure when there are two or more reserve nodes.

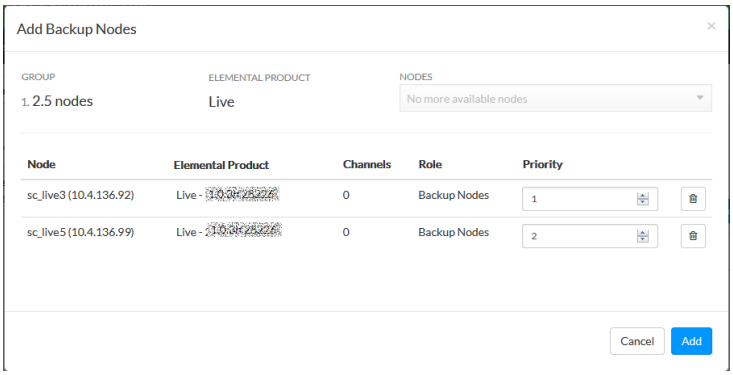

The final result: In this example, the group has been set up with two Active nodes and 1 Reserve node, which is type “N+1”.

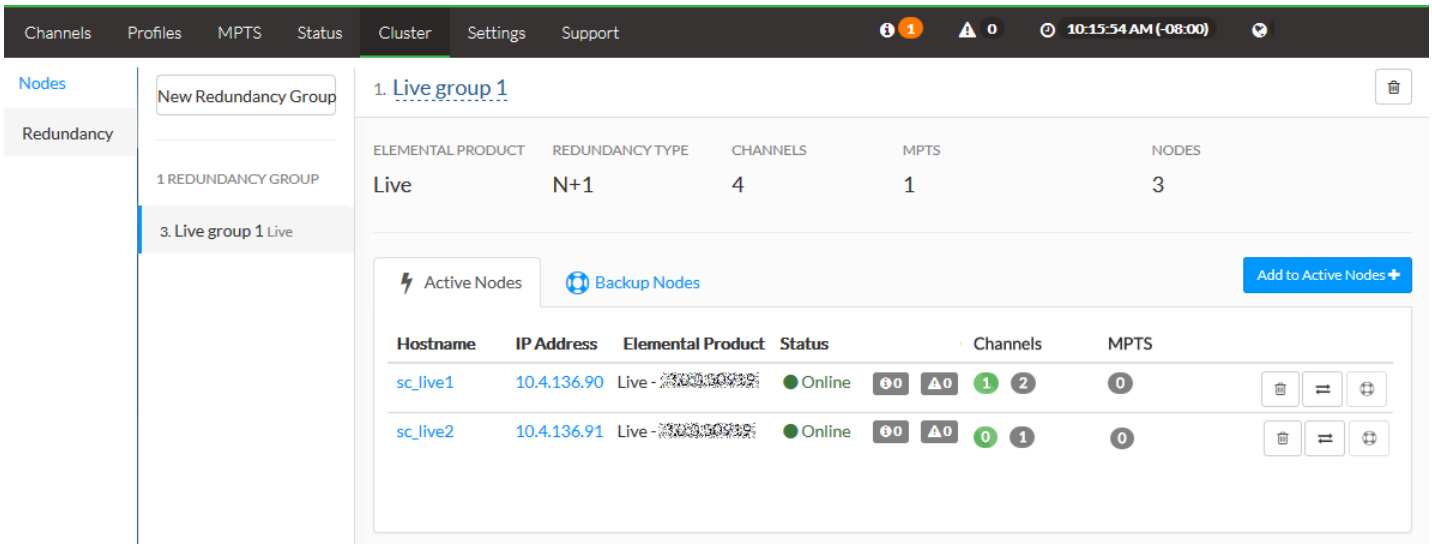

Step E: Testing Failover

Once you have at least one active node and one reserve node, you can test the failover sequence using the Initiate Failover feature. This feature lets you test your worker node redundancy setup without taking the drastic step of unplugging an AWS Elemental Live hardware unit or shutting the unit down from the command line.

Note:	Using the Initiate Failover feature has absolutely no effect on the status of Conductor redundancy and does not disable Conductor redundancy.

1. From the web interface for the primary Conductor node, display the Channels screen.
2. Set up a channel to use the active node; see the section on channels in the Conductor Live User Guide. Run at least one of the channels associated with that node.
3. Display the Redundancy screen and show the Backup Nodes tab. Make a note of the name of the reserve node that has the highest priority (priority closest to 1).
4. Switch to the Active Nodes tab.
5. Click the Initiate Failover button (double-arrow icon) beside the active node.
   - The active node will become Idle and its channel count will change to 0.
   - The node you noted on the Reserve Nodes will now appear on the Active Nodes tab. It will be Online and will show the channel count that the active node previously had. This node has picked up the channels that the active node was running.
   - If you only had one backup node, the Redundancy Type changes to N+0. Note that this is not a valid configuration and you will receive alerts until you have at least one backup node.

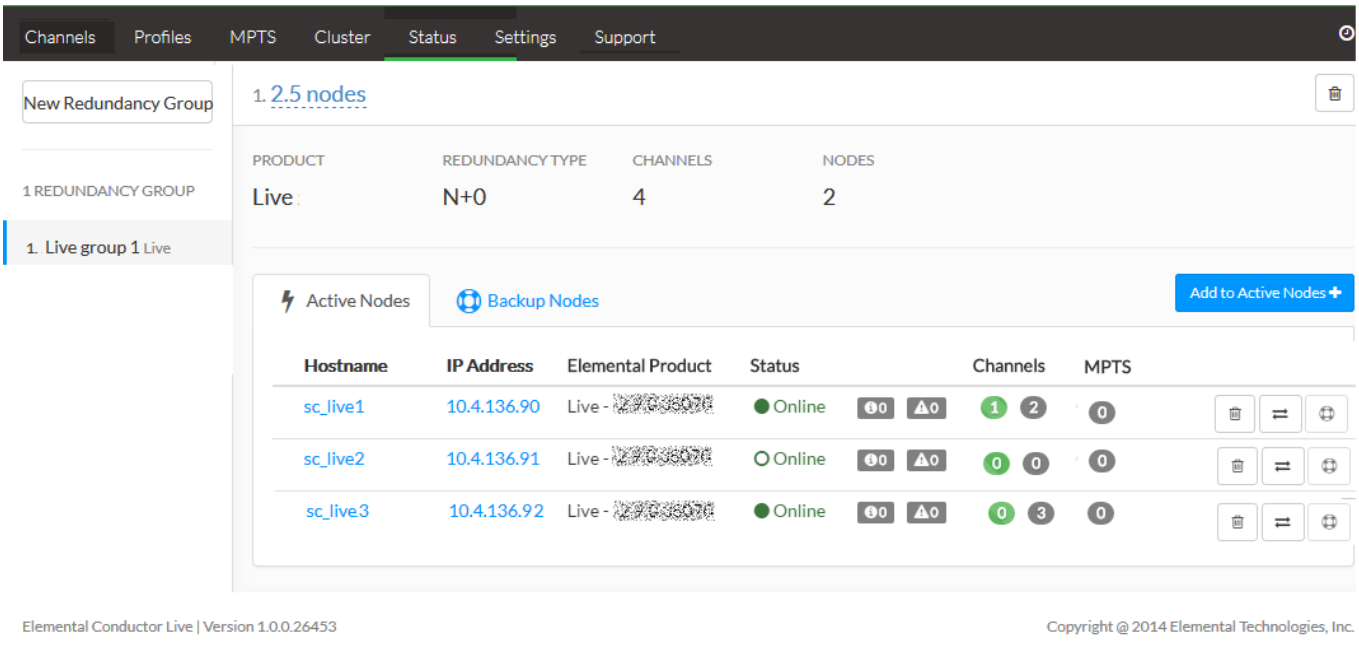

You now know that you have set up redundancy correctly.

2.1.6 Changing the Redundancy Setup

You can change the worker redundancy setup.

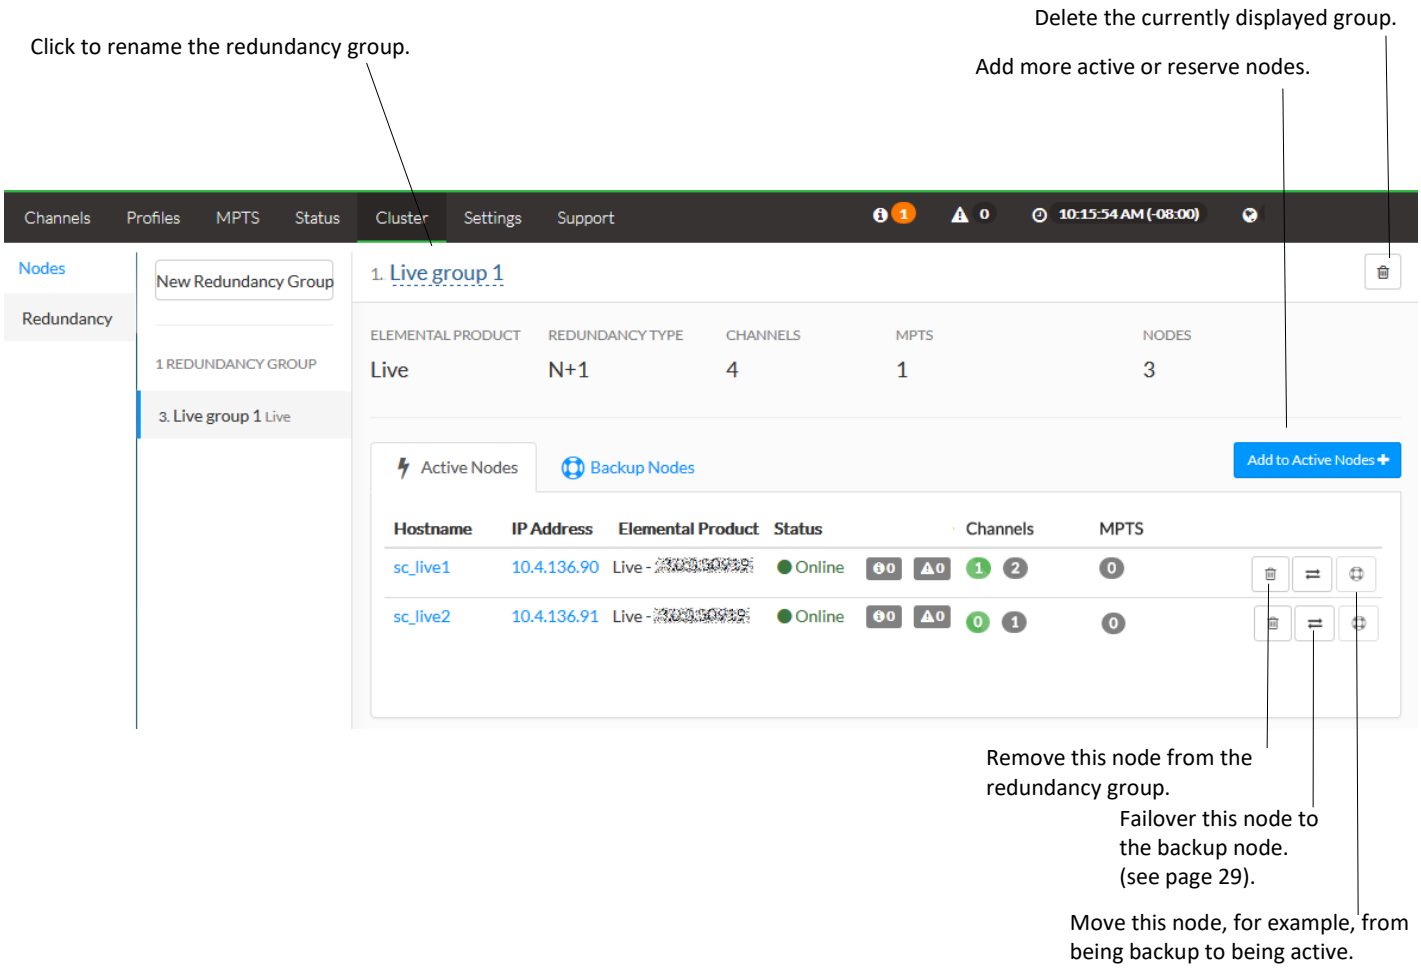

1.15. Adding Mount Points

Perform this procedure on	Both Conductor nodes

To make remote assets—such as scripts, image files, or video source files—accessible to your Conductor cluster, create mounts as described in this section. When you mount a remote folder to a local folder, all the contents of the remote folder will appear at this mount point as if the contents were actually in the local mount folder. In this way, you can view the folder and verify that backup files have been created. And you can, for example, copy or delete a file from the remote folder by copying it or deleting it from this mount folder.

Mount point controls initiated from the primary Conductor node are automatically synced to clustered worker nodes and any secondary Conductor node. The sync occurs within 1 hour for nodes already part of the cluster when the mount is created and within 3 minutes for nodes added to a cluster that already has mount points set up.

Note: Do not make mount point changes locally on worker nodes via the user interface; they will be overwritten the next time the sync task runs.

Mounts created through command line on a worker node are not impacted by the sync process.

1.15.1.1. Mounting on the Conductor Node

1. On the web interface for the primary Conductor node, go to Settings>Mount Points.
2. Click Add Mount Point. The Add New Mount Point dialog appears:

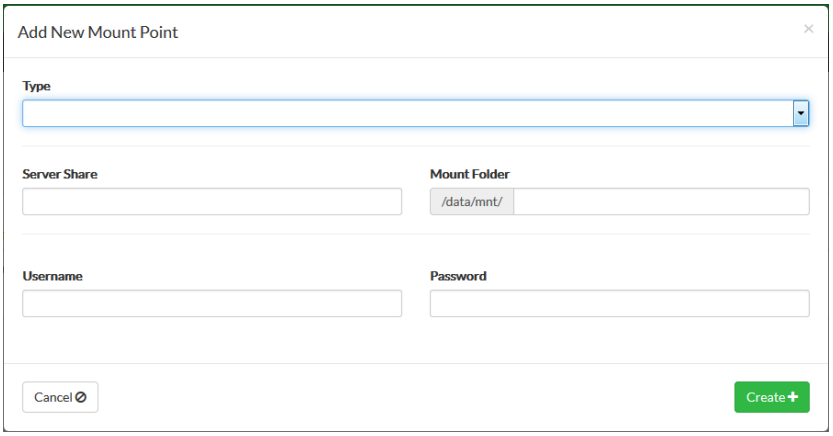

3. Complete the dialog as follows and click Create.

Field	Description
Type	Choose the type of remote server: • cifs: Choose this for a Windows CIF server or for a Windows or Mac SMB server. • nfs: Choose this for a Linux server. • webdav: Choose this for a DAVFS server.
Server Share	The address of the folder on the remote computer that you want to make available on this node.
Mount Folder	The folder on the Conductor node where the remote folder will be mounted. As shown, this folder must be under /data/mnt. You can specify a subfolder within that; if that folder does not exist, it will automatically be created.
Username	If the remote server folder is protected with a username/password, enter the username here.
Password	If the remote server folder is protected with a username/password, enter the password here.

4. Wait a few minutes. The newly mounted folder appears on the screen.

1.16. Configuring for SDI Direct Inputs

Perform this procedure on	The primary Conductor node

Your deployment can include SDI video inputs via a cable connected directly to the AWS Elemental Live hardware unit.

If your deployment includes SDI video inputs that are direct connects, make sure you import the devices when you add the node to the cluster. See page 21. Once this is done, you will be able to select “SDI Direct Input” as the Input type when creating a Profile.

1.17. Configuring for SDI Router Inputs

Perform this procedure on	The primary Conductor node

If your deployment includes SDI video inputs connected to a router, set up the SDI video router in Conductor Live 3 to provide information about your router configuration. Once this is done, you will be able to select “SDI Router Input” as the Input type when creating a Profile.

2.1.7 Configuration Overview

To set up the router information in Conductor Live 3, you will need to provide the following pieces of information:

- **Information on the router** itself: the name you want to use for it, its IP address, and any information required by your router’s protocol, such as User, Level, or MatrixID.
- The **router input** port numbers being used for streams that will go to AWS Elemental Live nodes.

  These are the numbers the router uses for its input ports, not numbers generated by AWS Elemental software. In the diagram below, they are represented by the green circles on the left. Called “Inputs” in Conductor Live 3.
- The **router output** port numbers being used by the cables between the router and Live nodes.

  These are represented in the diagram by green squares numbered 1 through 7 on the right side of the router. Called “Outputs” in Conductor Live 3.
- The **inputs to each SDI card** on each node. These correspond to the port where cables are plugged into the node from the router. Called “Connected to” in Conductor Live 3.

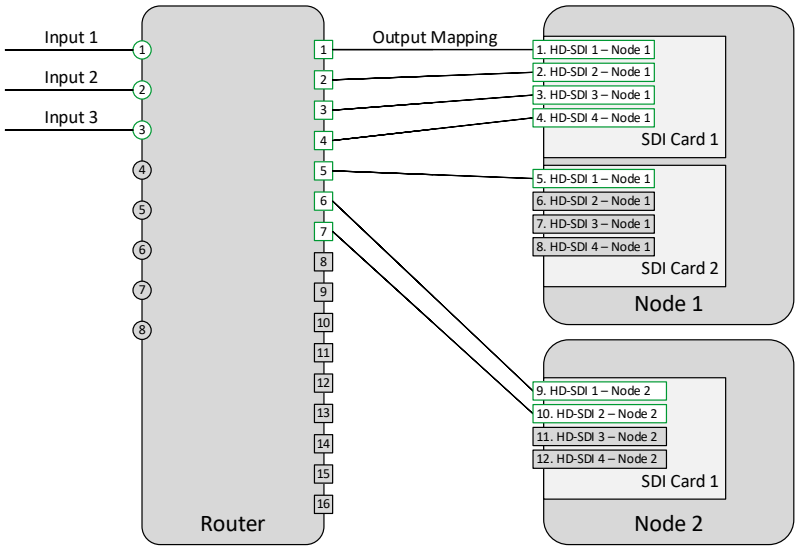

Example of an SDI Router Configuration

Step A: Getting Ready

Step B: Before you begin, make sure that you are using one of the following router protocols that are supported by Conductor Live 3.

Router	Protocol
Miranda nVision	NV9000
Snell Aurora	SW-P-08
BlackMagic VideoHub	Blackmagic Videohub Ethernet Protocol
Harris Panacea	Harris terminal protocol

1. Identify all the SDI inputs on every node that you want to use. One SDI card may have several inputs; you are not required to use them all. Make sure to include inputs on any reserve nodes (used for failover, page 24).
2. On the router, identify the outputs that you want to use with any AWS Elemental Live node. You must have one router output for each SDI input (on all the nodes) you want to use.

   For example, on Node 1 you have an SDI card on which you want to use 4 inputs and another SDI card on which you want to use only 1 input. On Node 2, you have an SDI card on which you want to use 2 inputs, for a total of 7 inputs. This means you need 7 outputs on the router. This is illustrated in the figure on page 33.
3. On the router, identify the inputs that you plan to use.

Step C: Create the Router

1. From the web interface for the primary Conductor node, choose Settings>Routers.
2. Click Add Router and choose the type of router. The Add New Router Dialog appears.

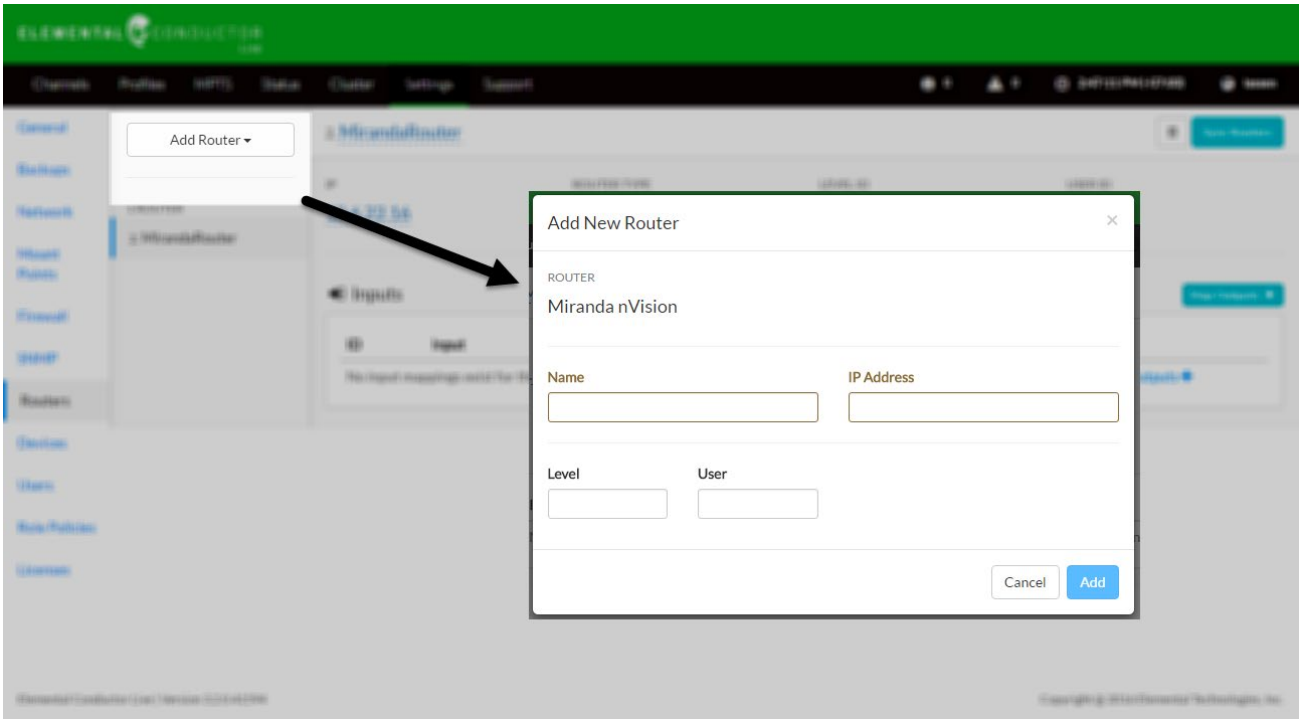

3. Complete the dialog as follows and click Create.

Field	Description
Name	This name will appear in the list of routers on the Routers screen after you click Add.
IP Address	The IP address without any protocol.
Level	Appears only for Harris Panacea and Miranda nVision.
User	Appears only for Miranda nVision.
Matrix Id	Appears only for Snell Aurora.

The router is added to the list of routers and is selected.

Step D: Complete the Input Mappings

Next, provide Conductor Live 3 with the port numbers for the router inputs that will end up at a Live node. Doing so maps the router's input number to an ID that is automatically generated by the software.

Note: An "Input Mapping" in Conductor 3 Live does not associate an input with an output. It only specifies which inputs to the router are used. The mapping from router inputs to router outputs is managed internally by the router.

1. Click Map Inputs. The Inputs section expands.

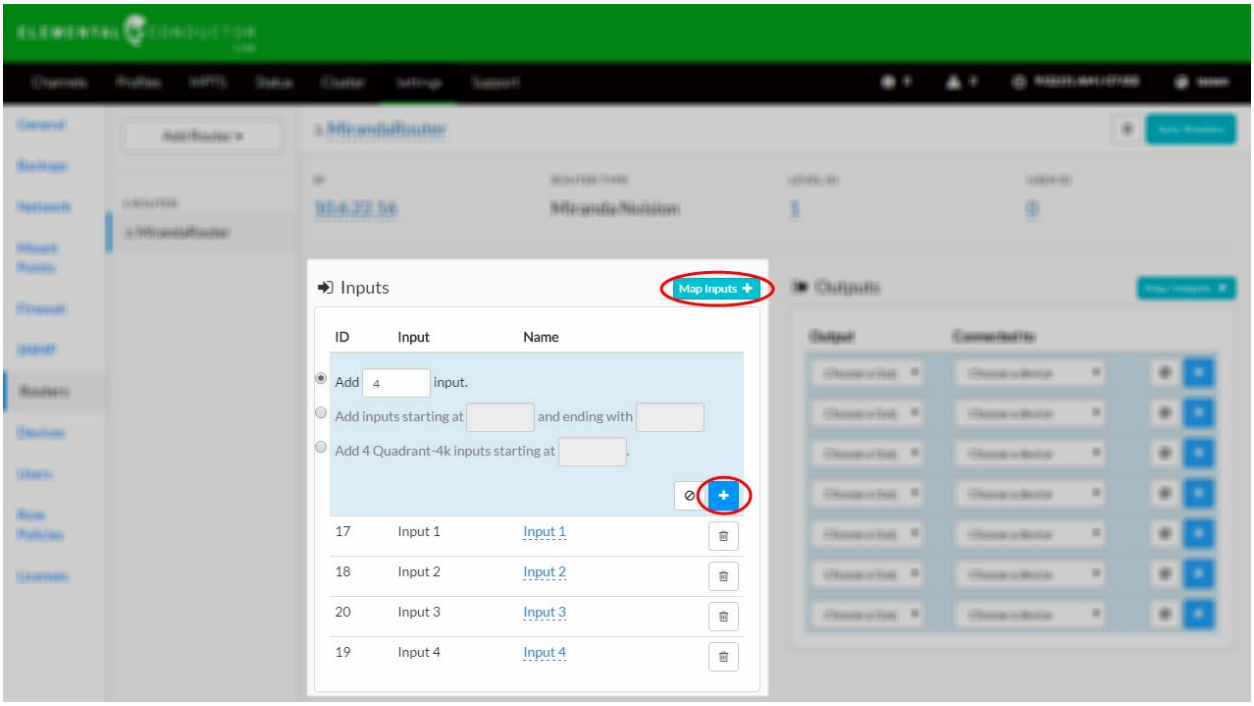

2. Complete the fields to identify the specific inputs you want to enable (those that have cabling):
   - Add: To add one or several individual inputs. Enter the number of inputs to be created. The system will assign input numbers, beginning with the first number after the largest one already created. In the image above, four inputs were created at once, numbered 1 through 4 because there were no existing inputs to start with.
   - Add inputs starting at: To add a range of individual inputs. Enter the first and last number in the range. Use this to create inputs that do not start at 1 and are not consecutive with the existing input numbers.
   - Add 4 Quadrant-4k inputs starting at: To add 4 inputs grouped together to create an HEVC input.

Note:	These input numbers must be the same as the identification of each input on your router. For example, you must know that, on the router, the second input from the left is "input 2." Conductor Live 3 cannot detect information about the disposition of input IDs.

3. Click the Add (+ icon) to add the inputs. The inputs you create will appear below the expanded Inputs section.
4. Repeat to add all the inputs you require.

   A line appears for each input. In the example above, four inputs are configured: the inputs with the router IDs 1, 2, 3, and 4. These router inputs have been assigned the AWS Elemental IDs 17, 18, 20, and 19 respectively. The name is automatically generated based on the input number you specify.

Step E: Complete the Output Mappings

Finally, complete the Output Mappings to map each router output to each desired SDI input (that is, to each SDI input on each AWS Elemental Live hardware unit that you plan to use). The mappings must reflect the actual cabling from the output side of the router to the input side of the SDI card.

The Output Mappings are represented in the figure on page 33 by the lines between the outputs on the router and the inputs on the SDI cards on the nodes. In the figure, the four inputs on the SDI card at the top have a path into the router. The one and only input on the second card has a path to the router. And two of the four inputs on the bottom SDI card have a path to the router.

1. Click Map Outputs. A new line appears. Repeat until you have enough lines for all the planned inputs on all SDI cards. Following from the above example, this would be 7 lines.

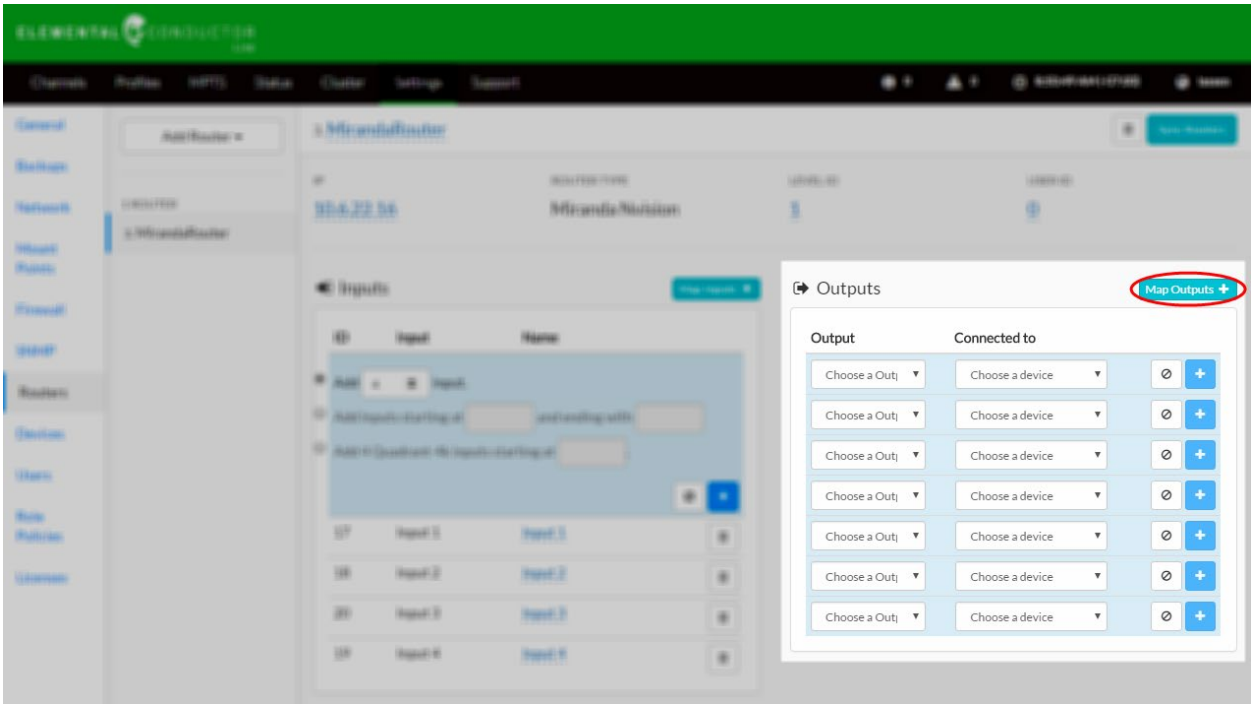

2. Complete the first line as follows:

   - Output: Click to display the dropdown menu. The choices in this dropdown list have the form "Output X," where X is a number that should correspond to the appropriate router output port.

     Select the appropriate Output from the list. For example, if your cabling comes from the router's output port 20, select "Output 20" from the dropdown list.

     Note: The correct number for the output is determined by the router, not by AWS Elemental software.

   - Connected to: Select the card and node that the router output is cabled to. The choices in this dropdown list are auto-detected from the node hardware by AWS Elemental software.

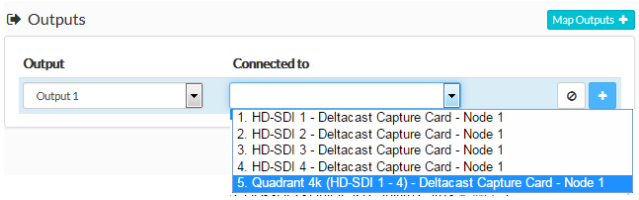

     If you added a 4 Quadrant-4k input in Step C and want to map those four inputs, select the Quadrant 4k (HD-SDI) card. This will map all four inputs to the one output.

   - Click Add (+ icon).

3. Repeat for each line to create all necessary Output Mappings. They will appear in a list at the bottom of the Outputs section.

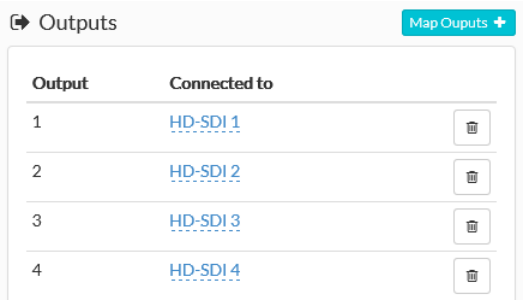

In the above example, only some of the physical SDI inputs are configured–those that are actually attached to the router.

Step F: Sync Routers

When the list of Inputs and the list of Output Mappings correctly represents the physical setup of the router and nodes, click Sync Routers (at the top right of the screen) to push relevant router information down to all AWS Elemental Live nodes.

1.18. Supporting RTMP Inputs

Perform this procedure on	AWS Elemental Live nodes

1. From the web interface for the AWS Elemental Live, choose Settings>Advanced. The Advanced screen appears.
2. Verify these fields:
   - Enable RTMP input: Checked.
   - RTMP input port: Specifies the desired port. The default port (1935) is already enabled on the node. If you specify a different port, you will have open on the firewall; see page 20.

3. Click Save.

Note:	The other fields on these Advanced tabs do not relate to initial configuration. They relate to fine-tuning of the load balance on the nodes. Therefore, they are not described in this guide.

1.19. Backing up the Databases

Perform this procedure on	The primary Conductor node

Backup copies the data related to your framework (channels, profiles, nodes, MPTS outputs, redundancy groups) from the Conductor node to another server in order to be able to restore the data to the Conductor node in case of a major hardware failure.

Backup files are named:

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

1.19.1.1. Setting up for Backup on Conductor Nodes

Backup and Conductor Failover

You only need to set up backup on the primary Conductor node. Keep in mind that the database on the backup Conductor node is already an exact copy of the primary database.

During normal operation, the primary Conductor node backs up its database according to the schedule.

In the event of a failover, the failed primary node will stop backing up to the remote storage and instead the other node (which is now primary) will start backing up. It will back up to the same location as the other Conductor node, so you will not have to worry about managing two backup files.

The Procedure

Note	Conductor Live 3 is configured by default to back up to /home/elemental/database_backups on the local node. We strongly recommend that you change this default to point to a mounted remote server as described in the following steps.

1. Designate a server and directory on your network as the backup server for the database for both Conductor nodes. Decide on a name for the backup directory, for example, Conductor_backups.
2. On the primary node, mount that remote server so that the Conductor node can write to it. See page 31.
3. Go to Settings>Backup. The Backups screen appears. Complete the table as shown below.

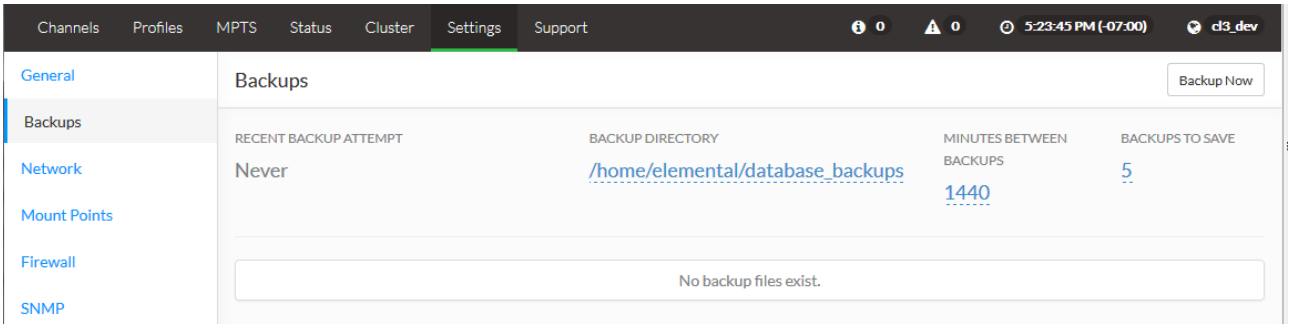

Field	Description
Recent Backup Attempt	Shows the date and time of the most recent backup – either succeeded (tagged with a green checkmark) or attempted (tagged with a red x).
Backup Directory	Click and specify the "Mount Folder" value of the server you mounted. For example, /data/mnt/Conductor_backups.

Field	Description
Minutes between Backups	The frequency of backups. The first backup will be made this number of minutes after you click Save. You can also create a backup at any time by clicking Backup Now on this screen. This backup does not affect the scheduled backup; the next scheduled backup will take place as planned.
Backups to Save	The number of backups to save. Once this number of files has been created, older files are deleted at the next scheduled backup.

4. On the backup node, mount the remote server so that the backup Conductor node can write to it. Note that you must mount the same directory as you mounted on the primary Conductor node. In this way, for which node is primary, backups are always being written to the same directory on the remote server. See page 31.

   Do not configure backup on the backup node; mounting the remote server is the only step you need to perform.

Disabling Backup

To disable backing up, change the Minutes Between Backups to 0.

1.19.1.2. Setting up for Backup on Worker Nodes

The node is automatically configured to back up its database to the local disk to:

```
/home/elemental/database_backups
```

Note:	When a worker node is part of a cluster, you can back up to the local disk, there is no need to back up to a remote server.

To view the backup setup:

1. From the web interface for the desired node, choose Settings>General.
2. Review the management database fields. +

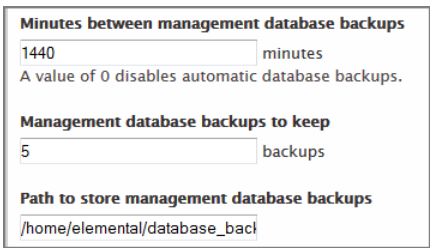

1.20. Restoring a Database

Perform this procedure on	The primary Conductor node and worker nodes

If you are restoring the Conductor database, you should restore to the hardware unit that is the primary Conductor node or that should be the primary Conductor node (once you have finished recovering from the failure.)

Note:	This procedure describes how to restore on a Conductor node. But the procedure is nearly identical on both AWS Elemental Conductor and a worker node: only the filename changes.

1.20.1.1. The Procedure

Run the install script with the restore option.

1. At your workstation, start a remote terminal session to the applicable Conductor hardware unit. Log in with username “elemental” and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.
2. Type the following command to identify the version of AWS Elemental Conductor that is currently installed:

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

   Several lines of information appear, including the version number, for example:

   ```
   Elemental Conductor Live (3.0.0.12345)
   ```

3. Run the install script with the restore option:

   ```
   [elemental@hostname ~]$ sudo sh ./elemental_conductor_live_3.0.n.nnnnn.run
   --restore-db-backup <path><backup-file>
   ```

Where:

- <product> is the product installer: elemental_production_conductor_live_n.n.n.nnnnn.run or elemental_production_conductor_file_n.n.n.nnnnn.run.

  n.n.n.nnnnn is the software version, which you obtained in the previous step.
- <path> is the path to the backup file. This path could simply be the remote directory where backups were originally stored.
- <backup-file> is the file you want to restore. The file is unzipped and copied to the appropriate directory. Do not unzip the file manually before restoring it!

1.21. Setting Up Alerts and Messages

Perform these procedures on	The primary Conductor node only

2.1.8 About Alerts and Messages

AWS Elemental systems provide information about the status of the systems and the channels via alerts and messages. Alerts can be sent to you; messages must be actively retrieved.

Manage alerts through the CL3 node only. CL3 systems will send operational status for all systems and all transcoding channels.

	Alerts	Messages
Access Options	Web UI SNMP poll REST calls Automatic email notification Web callback notification SNMP trap	Web UI SNMP poll REST calls
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 unexpected system error has occurred.
Active/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 "recent" when they are less than 24 hours old.
Visibility (Web UI 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. Visibility does not affect the return of SNMP and REST requests. Visibility does not affect email notifications, web callback notifications, or the emission of SNMP traps.	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.

See the CL3 User Guide for details on viewing alerts and messages via the web interface.

See the CL3 API Guide for information about using the REST interface to get alerts and messages.

2.1.9 Setting up for Email or Web Server Notification

You can set up your CL3 system to notify you when alerts occur. The notification can be an email or an HTTP POST to a web server. Do not configure worker nodes to send email notifications. CL3 systems will send operational status for all systems and all transcoding channels.

AWS Elemental systems use open relay to send email notifications. Before subscribing to notifications, make sure that your network allows receipt of open relay email. To receive messages from an AWS Elemental system in a network that does not allow receipt of open relay email, configure a sendmail relay with another mail server, as described in the section Configure Sendmail Relay Server on page 46.

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

Subscribing to All Alerts

To subscribe to all alerts generated by the cluster:

1. Go to Settings, which defaults to the General tab.

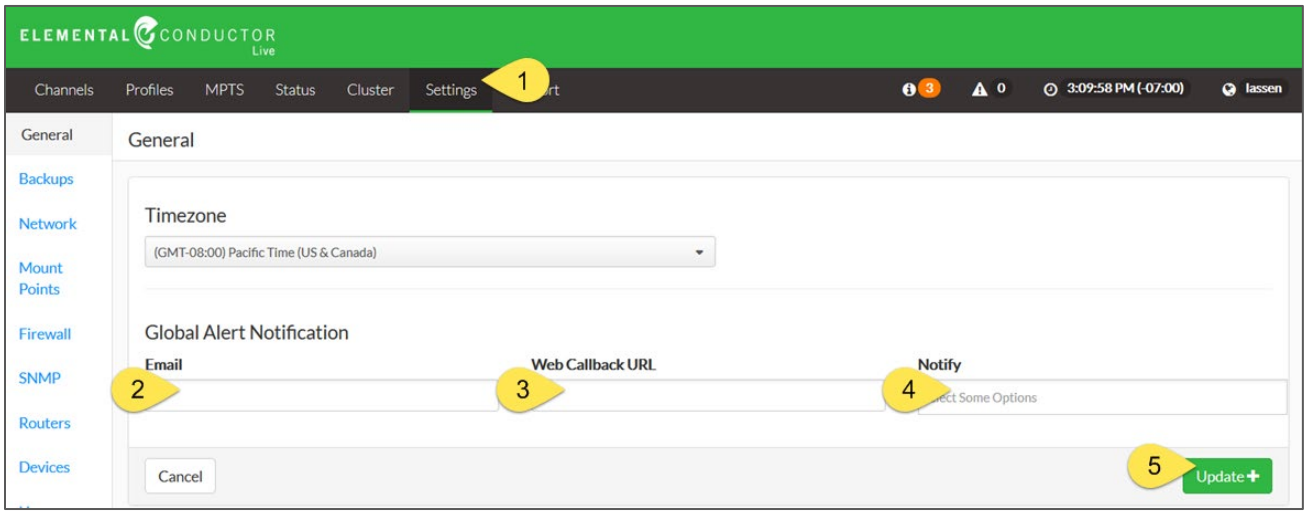

2. Complete the dialog as follows

	Field	Instructions
1	Settings	Click here to go to the Settings>General tab.
2	Email	Enter the email address of the alert recipient. Required if you do not enter a URL in the Web Callback field.
3	Web Callback URL	Enter the URL of the appropriate .php file on your web server in the Web Callback URL field. Required if you do not enter an email address. For instructions on how to configure your web server to receive notifications, see the section, Configure a Web Server for Notifications, page 45.
4	Notify	Select one or more of the options in the Notify dropdown box. The selections represent the type of change to the alert, for example, On Started means “when the alert first appears”. You can select several options. Some options only apply to some alerts.
5	Update	Click here to finish subscribing to alerts.

Subscribing to Individual Alerts

To subscribe to only specific alerts:

1. Go to Status>Notifications to bring up a list of all alerts.
2. In the row for the alert you want to subscribe to, click the + icon to bring up the Create a New Notification dialog.

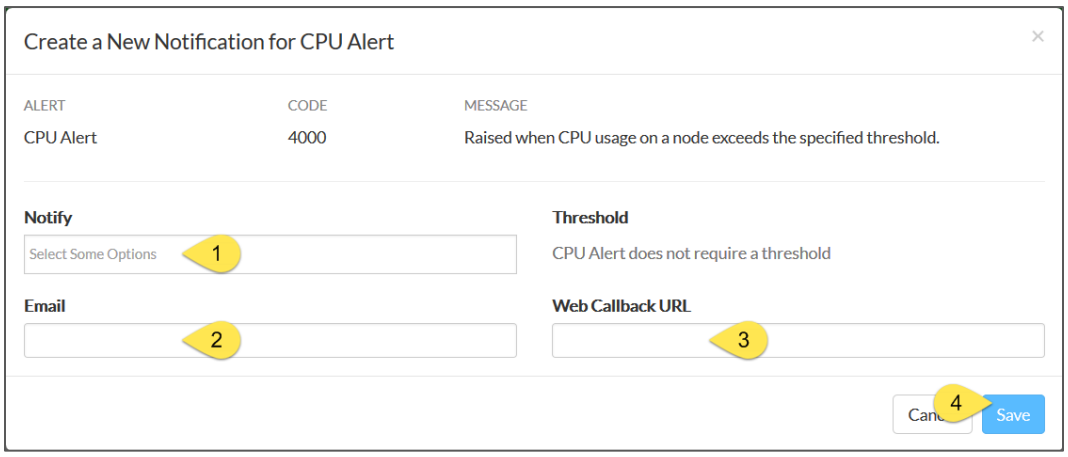

3. Complete the dialog as follows:

	Field	Instructions
1	Notify	Select one or more of the options in the Notify dropdown box. The selections represent the type of change to the alert, for example, On Started means "when the alert first appears." You can select several options. Some options only apply to some alerts.
2	Email	Enter the email address of the alert recipient. Required if you do not enter a URL in the Web Callback field.
3	Web Callback URL	Enter the URL of the appropriate .php file on your web server in the Web Callback URL field. Required if you do not enter an email address. For instructions on how to configure your web server to receive notifications, see the section Configure a Web Server for Notifications, page 45.
4	Save	Click here to finish subscribing to this alert.

4. Repeat to subscribe to multiple alerts individually.

1.21.1.2. Configure a Web Server for Notifications

To receive web callback notifications, you must have a web server that supports php scripting. You can configure this server to receive alert notifications from AWS Elemental systems as follows:

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. In this example, the file is called **notification.php** and is saved in the directory /webcallback.
3. Subscribe to global or individual alerts as described on pages 43 and 44. In the Web Callback URL field, enter the URL to your web server, e.g. http://example.com/webcallback/notification.php
4. Test your setup by typing the following (in a single line) at the command line of your CL3 system:

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

5. Open your notify.php file 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=value1&param2=value2
```

6. Enter your web callback URL (e.g. http://example.com/webcallback/notify) into a web browser to see the HTTP post.

1.21.1.3. Configure Sendmail Relay Server

Use this procedure to set up a sendmail relay server if your network does not accept open relay messages.

This procedure involves editing a file using a text editor at the Linux command line. These instructions are for using nano, which is already installed on all AWS Elemental systems.

Gather the Mail Server Information

To configure your CL3 system to relay the notification emails through a mail server, you will need the following information about the mail server:

- hostname
- IP address, only required if DNS is not configured on the network

Install the Sendmail Configuration Tool

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

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

   You will see a caution message asking you to confirm that you want to run the command.

   ```
   ##############################################################
   ##############################################################

    CAUTION: Updating system configurations can interfere with
    the proper operation of Elemental software. Continue only
    if instructed by Elemental Support. Please contact Elemental
    Support at techsupport@elementaltechnologies.com with any
    questions or concerns. Additionally, these custom changes
    may be reverted or overwritten during a standard Elemental
    software upgrade.

   ##############################################################
   ##############################################################
   ```

2. Type `yes`.

   The system will work for a while, then prompt you: `Is this ok [y/N]:`

3. Type `y`

   The system will work for a while, then return the message `Complete!`

Edit the Sendmail.mc File

1. Open the sendmail.mc file:

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

   This will open the file in nano. You will see your cursor at the top of the screen.

2. Use the down arrow key to go to the line that defines SMART_HOST. It is just past halfway down the page.

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

3. Uncomment this line by using the delete key and the arrow keys to delete the “dnl” at the beginning and end of the line.
4. Change the text “smtp.your.provider” to the hostname of the mail server that will be performing the relay.
5. Type Ctrl+o to save the file, then Ctrl+x to exit nano.

Check the Hosts File

If your network is not configured with DNS, you also need to add a static entry to the hosts file on your CL3 system. To do so:

1. Open the file /etc/hosts in nano:

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

   You will see something similar to this:

   ```
   132.0.0.1 SYS-1 localhost localhost.localdomain localhost4 localhost4.localdomain4
   ::1 SYS-1 localhost localhost.localdomain localhost6 localhost6.localdomain6
   ```

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. For example:

   ```
   192.0.2.0 ExampleMailHostname
   ```

3. Type Ctrl+o to save the file, then Ctrl+x to exit nano.

Apply the Changes

1. Type the following command to apply the changes:

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

   The system will respond as follows:

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

2. Restart sendmail by typing:

   ```
   sudo service sendmail restart
   ```

Test the New Configuration

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

1. Subscribe to global alert notifications on the CL3 web interface on the <CL3 IP address>/settings page. 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, but provide a bogus input address, such as “udp://1.1.1.1:1111.
3. Check your email for the notifications message.
4. If necessary, return to <CL3 IP address>/settings and re-subscribe the original recipient.

2.1.10 Setting Up SNMP Traps

Perform this procedure on	The primary Conductor node and worker nodes

Nodes can be configured to generate SNMPv2 traps as follows:

- All alerts generated by the AWS Elemental Live and CL3 nodes.
- When a Conductor node or worker node fails.

SNMP traps are generated for the following events:

Notification	Event	Contents
ELEMENTAL-MIB::alert	Any alert generated by the Live and Statmux nodes within the cluster	• ELEMENTAL-MIB::alertSet: 1 if the alert is being set, 0 if the alert is being cleared. • ELEMENTAL-MIB::alertMessage: Message describing the alert that was set or cleared.

1. On the web interface for the node, choose Settings>SNMP. The SNMP screen appears. The screenshot below is for Conductor Live 3. The screens for AWS Elemental Live and Statmux are nearly identical.
2. Complete the fields as described in the table below.

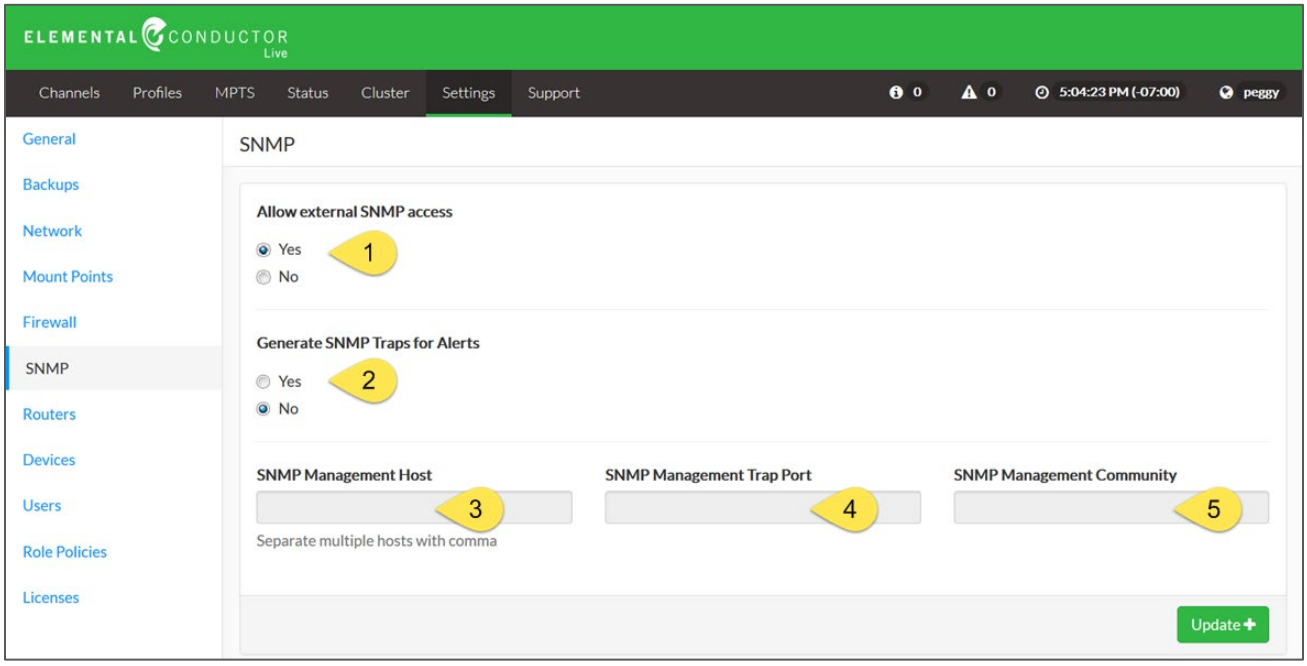

	Field	Instructions
1	Allow external SNMP access	Click Yes to open the SNMP port on the firewall. The port must be open in order to do a snmpwalk.
2	Generate SNMP traps for alerts	Click Yes to generate traps.
3	SNMP Management Host	Enter the IP address of the trap destination.
4	SNMP Management Trap Port	Enter “162”.
5	SNMP Management Community	Enter “Public”.

3. Click Update.

2.1.11 Setting Up SNMP Polling

Perform this procedure on	The primary Conductor node.

Rather than passively receiving traps sent to you by the systems, you can actively poll the SNMP interface. Polling the Conductor node with a snmpwalk, as described on page 50, will provide you with information about all the systems and encoding channels in the cluster.

You can interact with CL3 system using a variety of network management systems. AWS Elemental products ship with the Net-SNMP (http://www.net-snmp.org/) command-line tools to access the SNMP interface while logged into the system directly or over SSH. Examples in this document are given using net-snmp commands.

In order to access the SNMP interface externally, either the firewall must be disabled or access to just the SNMP interface must be enabled. Default settings allow external access to the SNMP interface. See page 20 for instructions on enabling and disabling the firewall. The setting for enabling access to the SNMP interface through the firewall is on the Settings>SNMP tab.

1.21.1.4. Management Information Bases

AWS Elemental provides the following Management Information Bases (MIBs) for use with CL3 clusters:

- ELEMENTAL_MIB.txt - Base MIB for all AWS Elemental systems
- ELEMENTAL_CONDUCTOR_MIB - Objects specific to nodes in the cluster Conductor Live 3 channels

Details on each MIB are provided in the subsections below. These MIBs are installed on the system by default, located in /opt/elemental_se/web/public/mib/.

You can use the MIBs with the net-snmp tools to query individual variables as follows:

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

returns

```
ELEMENTAL-MIB::serviceStatus.0 = INTEGER: 1
```

All AWS Elemental: ELEMENTAL_MIB

All AWS Elemental systems ship with the base MIB, which can provide information on the system itself. The following variables are available via this MIB:

Variable	Values
serviceStatus	0 if the elemental_se service is not running, 1 if the service is running.
firewallStatus	0 if the system's firewall is off, 1 if on.
networkSettings	Will always return 1. Required for some network management systems.
mountPoints	Number of user-mounted filesystems in /mnt.
version	Product version.
httpdStatus	0 if the httpd service is not running, 1 if the service is running.
databaseBackup	1 if writes (starting backups) are allowed. 0 if writes are not allowed.

CL3: ELEMENTAL_CONDUCTOR_MIB

Conductor Live 3 systems ship with this MIB, which provides the following variables:

Variable	Description
channelId	System-assigned ID of the channel.
channelName	User-defined name of the channel.
channelRunning	Indicates whether the channel is currently running. 0 is not running, 1 is running.
channelError	Indicates whether the channel is in the error state. 0 is no error, 1 is error.
channelLiveEventId	Live-assigned ID of the event associated with the channel.
channelStartTime	Start time of the channel. Provided only if the channel is currently running.
channelDuration	The duration of time the channel has been running. Provided only if the channel is currently running.
channelAlerts	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.
channelMessages	The text bodies of any messages generated in the last 24 hrs related to the channel, including the time the message was last set. Each message is separated by semicolons.
nodeId	CL3-assigned ID of the node that the channel is running on.
nodeHostname	Hostname of the node that the channel is running on.

1.21.1.5. Polling the Entire SNMP Interface

You can gather information on their events by doing a snmpwalk on the CL3 system. This will provide information on all the channels running in the cluster.

The entire Conductor Live 3 interface can be queried via snmpwalk as follows.

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

1.22. Managing User Authentication

1.22.1.1. Purpose of User Authentication

User authorization with Conductor Live 3 is intended to:

- Allow managers to track activity throughout the cluster on a per-user basis.
- Limit accidental access or changes to a node by allowing distinct login credentials for each node. This way, an operator with access to multiple nodes must enter the credentials for a specific node prior to sending any commands.

Whether authentication is enabled or not, we recommend that the cluster always be installed behind a customer firewall or a private network.

1.22.1.2. Types of User Authentication

We offer two types of user authentication:

- Local authentication: user information is created and managed from the node.
- PAM authentication: user information is created and managed from an external LDAP server.

See the configurations on page 52 for how these authentication options can be deployed.

1.22.1.3. User Authentication and User Types

User authentication involves using the credentials of several types of users:

User Type	How created	Log-in Credentials		Use
		Username	Password	
Default, remote terminal user	Built-in	elemental	Default, or as changed by admin.	Users manually enter this information when: • Logging in to a remote terminal session on the Conductor. • Enabling user authentication on worker nodes from the Conductor web interface. See page 54 for more information. • When using PAM, when logging in to the Conductor the first time after authentication is enabled.
Admin API user	Created in command line when authentication is enabled on the Conductor.	Customer created. Must *not* be the username of another administrator (i.e., not the name of a real person).	Customer created.	• The Conductor node presents this information when communicating with another node. This authentication automatically takes place when a user submits a Live node command through the Conductor node. • The person configuring authentication uses this information the first time they access a node's web interface after local authentication is enabled. This is typically the only time a user manually uses this information.
People and third-party clients	Created through the web interface, or through an LDAP server (for PAM).	Customer created. Varies by user.	Customer created. Varies by user.	Users manually enter this information when accessing a node through the web interface or REST API.

1.22.1.4. User Authentication Configuration Options

Option 1: Local Authentication on Conductor Nodes

In this configuration, users logging in to the Conductor Live node through the web interface or REST API are validated against log-on credentials housed locally on the Conductor. Users can access worker nodes without providing usernames and passwords.

The following steps apply to option 1:

1. Disable high availability (if applicable). See page 66.
2. Enable local authentication on the primary Conductor node. See page 53.
3. Set-up users as administrators, viewers, or operators on the primary Conductor node. See page 58.
4. Re-enable high availability. See page 65.

Option 2: Local Authentication on All Nodes

In this configuration, users logging in to the Conductor Live node and worker nodes through the web interface or REST API are validated against log-on credentials housed locally on each node in the cluster.

The following steps apply to option 2:

1. Disable high availability (if applicable). See page 66.
2. Enable local authentication on the primary Conductor node. See page 53.
3. From the Conductor web interface, enable local authentication on the worker nodes. See page 54.
4. Set-up users as administrators, viewers, or operators on the primary Conductor node. See page 58.
5. Set-up users as administrators, viewers, operators, or managers on each worker node. See page 59.
6. Re-enable high availability. See page 65.

Option 3: PAM Authentication, Local Authentication on Worker Nodes

In this configuration, users logging in to the Conductor Live node through the web interface or REST API are validated against log-on credentials housed on an external LDAP server. Additionally, users logging in to worker nodes are validated against log-on credentials housed locally on each worker node.

The following steps apply to option 3:

1. Disable high availability (if applicable). See page 66.
2. Enable PAM authentication on the primary Conductor node. See page 60.
3. From the Conductor web interface, enable local authentication on the worker nodes. See page 54.
4. Set-up users as administrators, viewers, operators, or managers on each worker node. See page 59.
5. Modify role policies (as needed). See page 63.
6. Re-enable high availability. See page 65.

Option 4: PAM Authentication Only

In this configuration, users logging in to the Conductor Live node through the web interface or REST API are validated against log-on credentials housed on an external LDAP server. Users can access worker nodes without providing usernames and passwords.

The following steps apply to option 4:

1. Disable high availability (if applicable). See page 66.
2. Enable PAM authentication on the primary Conductor node. See page 60.
3. Modify role policies (as needed). See page 63. Re-enable high availability. See page 65.

2.1.12 Enabling Local Authentication

Use the following procedures when enabling local authentication on the Conductor node (for options 1 and 2) and on the worker nodes (for options 2 and 3).

Perform this procedure on	The primary Conductor node and worker nodes (as needed)

1.22.1.5. Enabling Local Authentication on a Conductor Node

Perform the following steps at the Linux command line.

1. Disable Conductor redundancy (if applicable) as described on page 66.
2. At your workstation, start a remote terminal session to the primary Conductor node.
3. At the Linux prompt, log-in with the remote default username “elemental” and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.
4. Change to the directory where the configuration script is located:

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

5. Run the configuration script as follows:

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

   Where:

   -a specifies to show only the user authentication prompts

   --https enables SSL (see SSL Configuration section on page 7 for more information)

   You can enable authentication without SSL, but passwords will be transmitted without encryption.

Warning: All nodes must have the same SSL configuration (enabled or disabled). If you are making any changes to SSL, all nodes must be removed from the cluster and updated.

6. This prompt appears:

```
Do you wish to enable authentication?
```

   Enter Y.

7. If you did not use the --https option for SSL, the following prompt appears:

```
If you wish to enable authentication, please re-run with the '--https' option.
If SSL isn't enabled, any usernames/passwords entered here including LDAP passwords would
be transmitted in plain text without encryption. This poses a significant security risk.
Accept the risk and continue without SSL?
```

If you intended to enable authentication without SSL enabled, enter Y to proceed. Otherwise, enter N to re-enter the configuration script with --https.

Warning: If you already had SSL enabled and do not include the --https option, SSL will be disabled.

8. This prompt appears:

```
Do you wish to enable PAM?
```

Enter N. If you are using PAM authentication, see configuration options 3 and 4 on page 52.

9. You are prompted to enter a username, email address, and password for the "admin API user," as described on page 51.

Note that the only time you will log-in with this information is upon initial access to each node web interface after authentication is enabled.

10. You are notified that authentication is enabled and the following prompt appears:

```
Httpd must be restarted, which may interrupt REST commands. Restart now?
```

Enter Y.

The AWS Elemental service starts and the Conductor node is ready.

11. If applicable, re-enable Conductor redundancy as described on page 65.

1.22.1.6. Enabling Local Authentication on a Worker Node

1. Make sure you have followed the procedure in the Enabling Local Authentication on a Conductor Node section.
2. Go to the Conductor web interface by entering the IP address of the primary Conductor node in a web browser.

For example:

```
10.4.136.90
```

The Login screen appears.

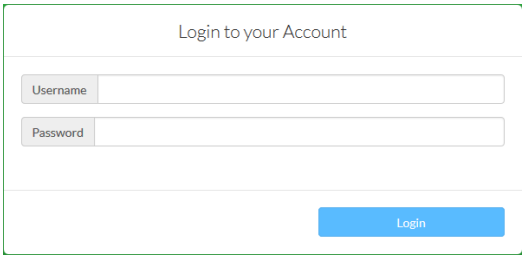

3. Log into the web interface using the "admin API User" credentials you created via the configure script.

Note:	Typically, you would never log into the web interface using the "admin API User" credentials. But initially, this is the only user that exists. In a later step in this procedure, you will create credentials for your administrators, including yourself, and use those credentials for subsequent logins.

You cannot log in using the built-in, remote terminal "elemental" user credentials!

The Conductor web interface appears.

4. Choose Cluster>Nodes.
5. Click and hold down Tasks and choose Enable Node Authentication.

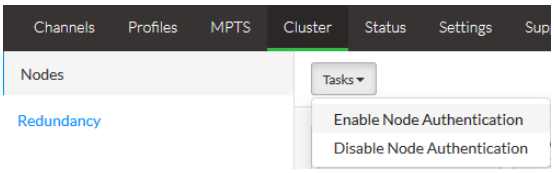

The Enable Node Authentication screen appears.

6. Enter the credentials for the “admin API User”. Then click Next.

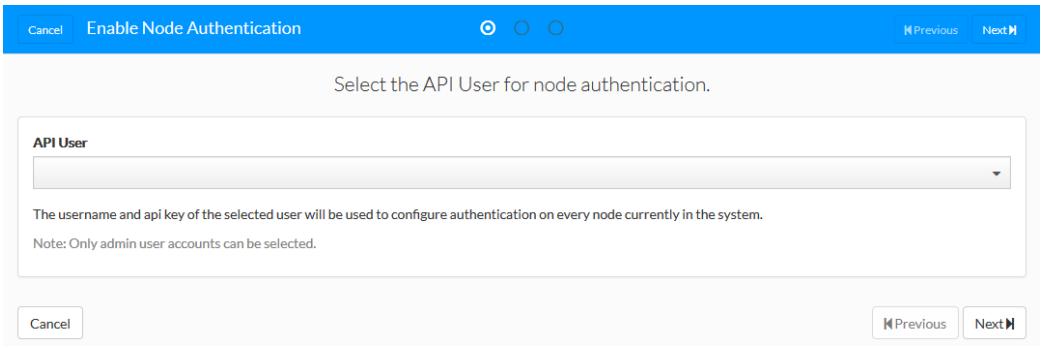

7. On the second dialog, create a new password for the admin API user to use when communicating with worker nodes to enable authentication. The password must be a minimum of 4 characters.

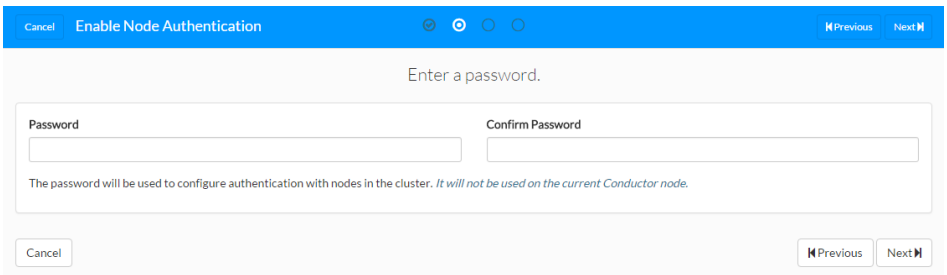

8. On the third dialog, enter the built-in, remote terminal “elemental” username and the password set by your system administrator. Then click Next.

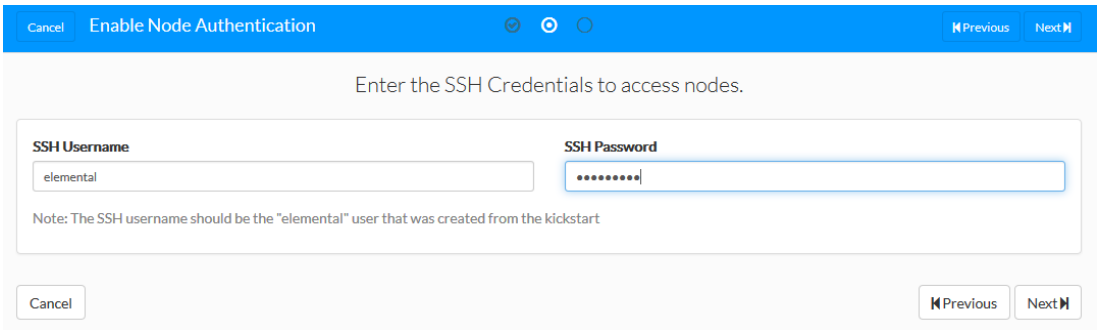

9. Click Configure Now.

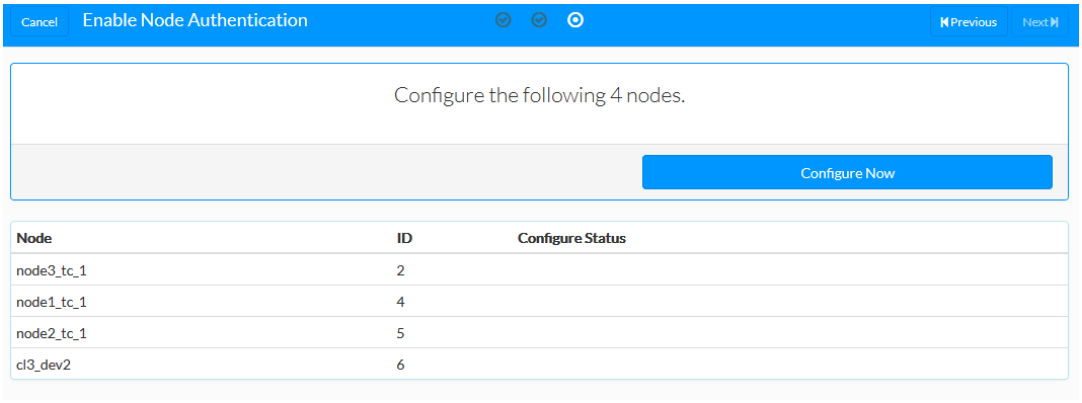

All the nodes in the cluster will be configured with user authentication enabled. Once user authentication is enabled, the Cluster>Nodes screen appears.

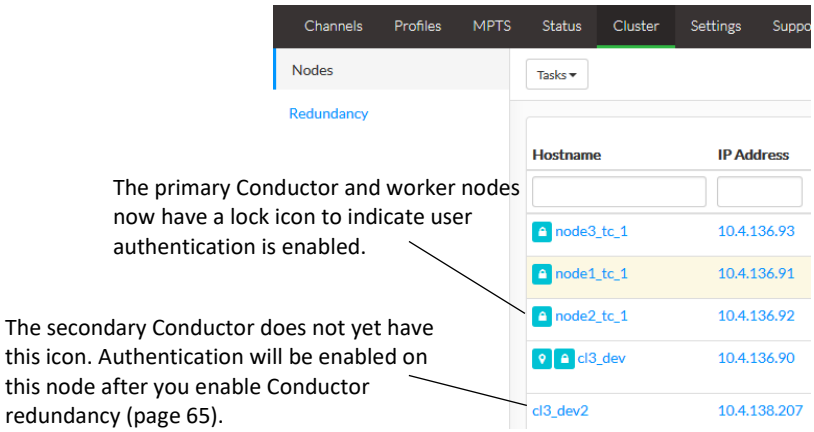

2.1.13 Creating Users

Perform this procedure on	The primary Conductor node and worker nodes (as needed)

When local authentication is enabled, the following procedures are performed in the web interface of each applicable node. Remember that with PAM authentication, users are maintained in the LDAP server and are not added to Conductor Live.

User Roles

Access is defined by the role assigned to the user (either locally or on the LDAP server):

- Admin can access all functionality on the node.
- Operator can access all functionality except the ability to add or remove users and enable or disable Node Authentication.
- Viewer can access all screens but cannot perform actions.
- Manager can access event functionality on the worker node. This role is not available on the Conductor.

When PAM authentication is enabled, users cannot be added or edited, but the roles can be managed through the Role Policies screen. See section 2.1.15 on page 62 for more information.

Users and Node API Keys

If a user has access to multiple worker nodes, a separate API key is needed for each node. This is because, as discussed in Purpose of User Authentication on page 51, users are set up locally on each node to avoid erroneous changes to the wrong nodes. When you set up a user on the Conductor node and five worker nodes, you are setting up six different users whose credentials do not span multiple nodes.

Adding Local Users with Option 1

If you are enabling local authentication only on Conductor nodes:

1. Add administrators to the primary Conductor node. See page 58.

   Remember to add an admin user for yourself, so you are not using the API User credentials (which should not be used by a regular administrator except the very first time you log in).
2. Add other users as operators or viewers to the primary Conductor node.
3. Provide other users with their username and password.

Adding Local Users with Option 2

If you are enabling local authentication on all nodes in the cluster:

1. Add administrators to each node: both the primary Conductor node (page 58) and all of the worker nodes (page 59).

   Remember to add an admin user for yourself, so you are not using the API User credentials (which should not be used by a regular administrator except the very first time you log in).
2. Add other users as operators or viewers to the primary Conductor node.
3. If you want users to be able to access the worker nodes directly, add those users as operators or viewers on the relevant worker nodes.

   If you don’t to add a user to a particular node, that user will not be able to use the web interface or REST API interface on that node.

Notify users:

- Provide users with their username and password.
- For users and clients who will use the REST API, advise the users to log into the web interface for each node and click their name in the menu bar to obtain their personal API key for that node:

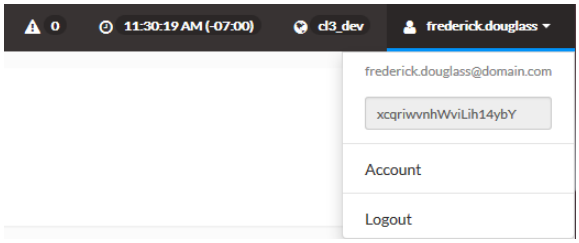

- Notify REST API users that REST commands must include additional HTTP headers that identify the user, as described in the introductory sections of the Conductor Live 3 API Guide.

Adding Local Users with Option 3

If you are enabling PAM on the Conductor node and local authentication on worker nodes, follow the steps in the Adding Local Users with Option 2 section.

Users for the Conductor node are added to the LDAP server and are not addressed in the web interface.

1.22.1.7. Adding Users on a Conductor Node

1. Display the web interface for the primary Conductor node. The login dialog appears (because you enabled user authentication in the Enabling Local Authentication on a Conductor Node section).
2. Log in as the “admin API User” you created via the configure script.
3. Choose Settings>Users. The Users screen appears. The API User appears in the user list.
4. Click Add User.
5. The Add New User screen is displayed. Complete all fields and click Save.

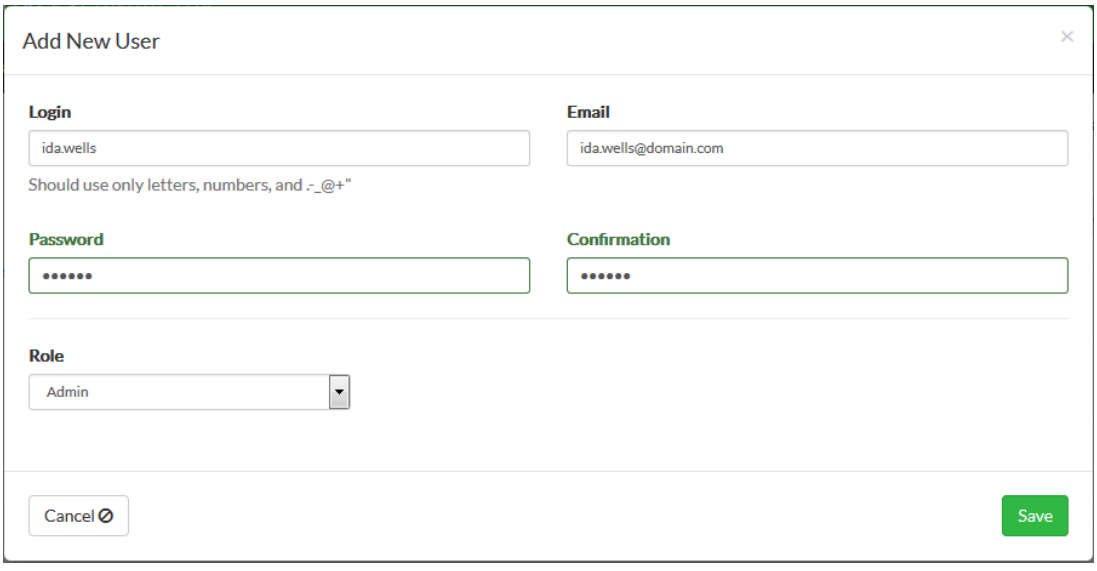

The new user is added to the list of users:

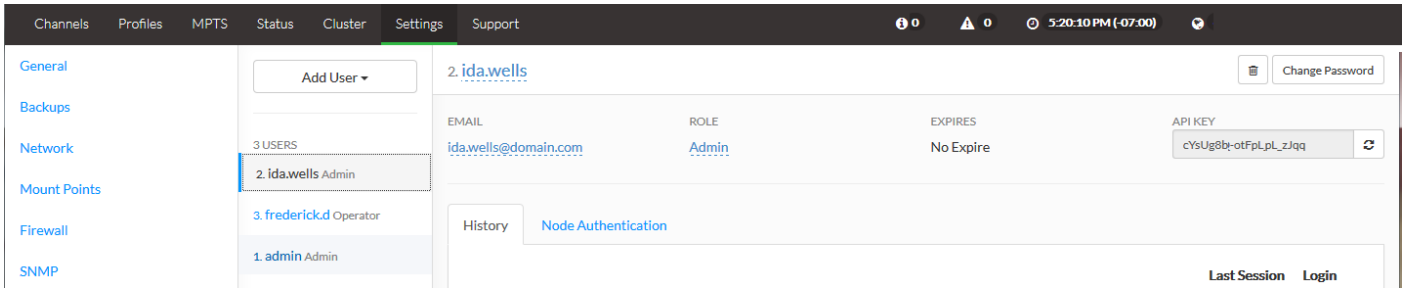

1.22.1.8. Adding Users on a Worker Node

1. Display the web interface for the worker node. The login dialog appears (because you enabled user authentication in the Enabling Local Authentication on a Worker Node section).
2. Log in as the "admin API User" you created via the configure script.
3. Choose Settings>Users. The Users screen appears. The API User appears in the user list because it was pushed from the Conductor node during the setup in Section Adding Users on a Conductor Node, above.

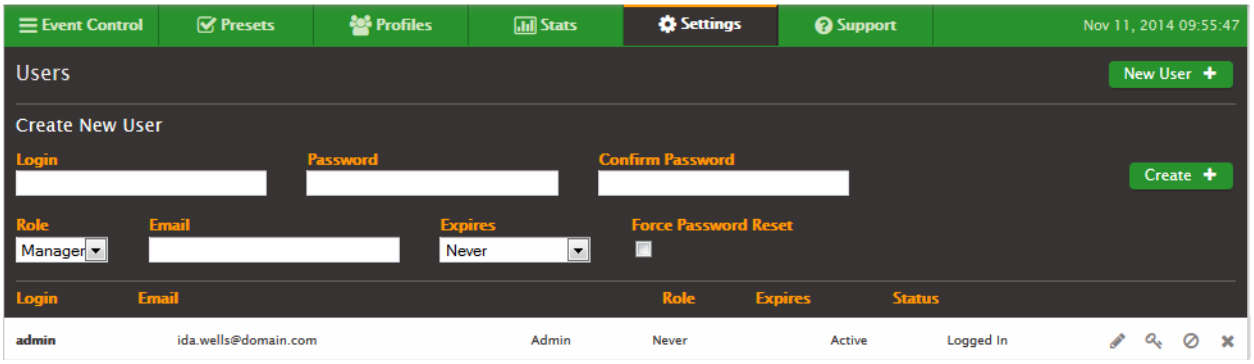

4. Complete all fields and click Create. Some notes:
   - Expires: If checked, the user will automatically expire after the specified period of time.
   - Force Password Reset: If checked, user must reset their password the first time they login.
   - Role: Select a role: Admin, Manager, User, and Viewer.

Note	If your organization uses the REST API, make sure to tell each user to log into the web interface on the worker node and choose Settings>User Profile in order to make a note of their personal API key.

User Roles

When creating users, choose the appropriate role, based on the following table:

		Roles			
Action	Meaning	Admin	Manager	Operator	Viewer
Manage Users	Create and edit users and roles.	✓			
Manage Live Events	Create and edit jobs.	✓	✓		
Control Live Events	Control the state of jobs (Start, Stop, Archive, etc.).	✓	✓	✓	
Manage Presets	Create and edit Presets, Preset Categories, and Audio Remixing Presets.	✓	✓		

		Roles			
Action	Meaning	Admin	Manager	Operator	Viewer
Manage Profiles	Create and edit Profiles.	✓			
Manage Schedules	Create and edit Schedules.	✓			
Manage System Settings	Update the system settings (any tab under Settings)	✓			
Manage Alerts	Update alert thresholds and to update alert notification settings.	✓			
Read-only access	View all screens	✓	✓	✓	✓

2.1.14 Enabling PAM Authentication

Use the following procedure when enabling PAM authentication on the Conductor node (for options 3 and 4).

Perform this procedure on	The primary Conductor node

Perform the following steps at the Linux command line.

1. Disable Conductor redundancy (if applicable) as described on page 66.
2. At your workstation, start a remote terminal session to the primary Conductor node.
3. At the Linux prompt, log-in with the remote default username “elemental” and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.
4. Change to the directory where the configuration script is located:

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

5. Run the configuration script as follows:

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

   Where:

   -a specifies to show only the user authentication prompts.

   --https enables SSL (see SSL Configuration section on page 7 for more information)

   You can enable authentication without SSL, but passwords will be transmitted without encryption.

Warning: All nodes must have the same SSL configuration (enabled or disabled). If you are making any changes to SSL, all nodes must be removed from the cluster and updated.

6. This prompt appears:

   ```
   Do you wish to enable authentication?
   ```

   Enter Y.

7. If you did not use the --https option for SSL, the following prompt appears:

   ```
   If you wish to enable authentication, please re-run with the '--https' option.
   If SSL isn't enabled, any usernames/passwords entered here including LDAP passwords would
   be transmitted in plain text without encryption. This poses a significant security risk.
   Accept the risk and continue without SSL?
   ```

   If you intended to enable authentication without SSL enabled, enter Y to proceed. Otherwise, enter N to re-enter the configuration script with --https.

Warning: If you already had SSL enabled and do not include the --https option, SSL will be disabled.

8. This prompt appears:

```
Do you wish to enable PAM?
```

   Enter Y.

9. You are notified that authentication is enabled and the following prompt appears:

```
Httpd must be restarted, which may interrupt REST commands. Restart now?
```

   Enter Y.

   The AWS Elemental service starts and the Conductor node is ready.

10. If applicable, re-enable Conductor redundancy as described on page 65.

2.1.15 Managing Role Policies

Perform this procedure on	The primary Conductor node

As described in User Roles on page 57, role policies dictate what kind of access a user has when working on the Conductor node.

Conductor Live comes with 6 default role policies. You can use the role policies as-is and define groups in the LDAP server to match, or modify the default roles in Conductor (except for the elemental/Admin role) to match your environment.

1.22.1.9. Adding Role Policies

1. Display the web interface for the primary Conductor node. The login dialog appears (because you enabled user authentication in the Enabling Local Authentication on a Conductor Node section).
2. Log in with the built-in, remote terminal “elemental” username and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.
3. Choose Settings>Role Policies. The Role Policies screen appears.
4. Click Add Role Policy.

   You cannot add, edit, or delete role policies unless PAM authentication is enabled.
5. The New Role Policy screen is displayed. Complete all fields and click Save.

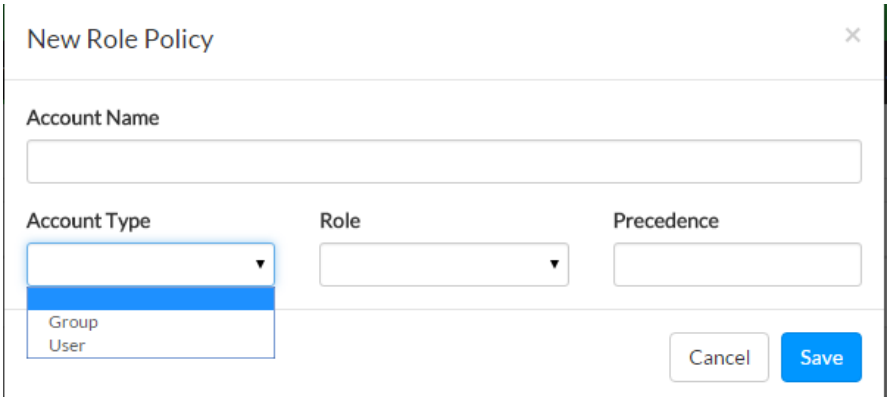

The new role policy is added to the list:

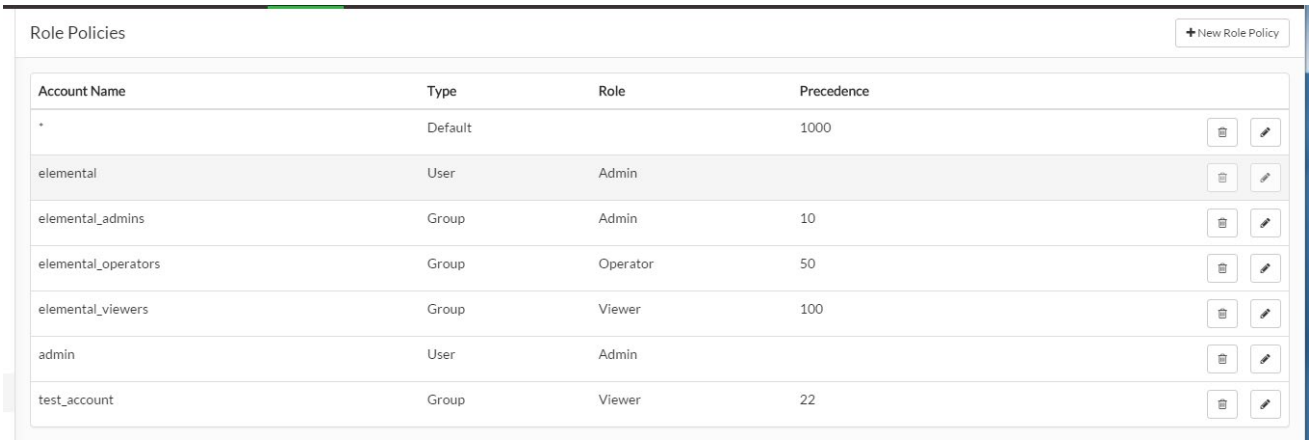

1.22.1.10. Modifying Role Policies

1. Display the web interface for the primary Conductor node. The login dialog appears (because you enabled user authentication in the Enabling Local Authentication on a Conductor Node section).
2. Log in with the built-in, remote terminal “elemental” username and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.
3. Choose Settings>Role Policies. The Role Policies screen appears.
4. Click the edit button (pencil) to the right of the policy to be modified.
5. The Update Role Policy screen is displayed. Modify applicable fields and click Save.

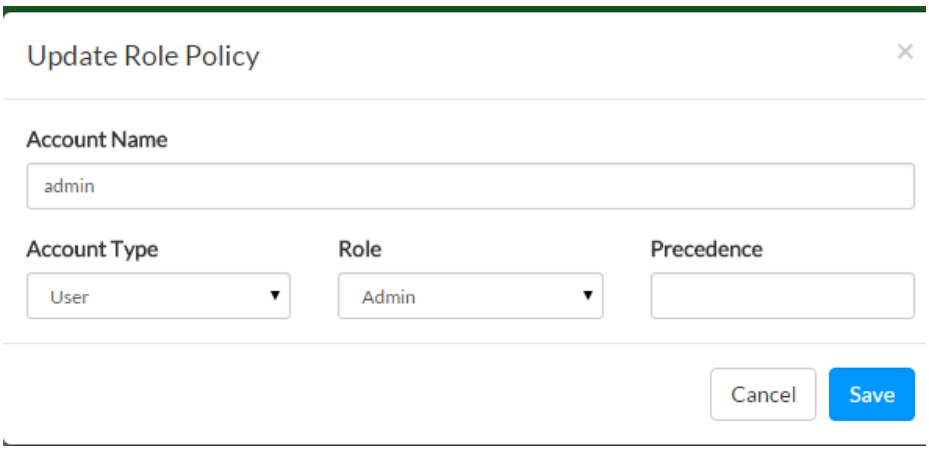

2.1.16 Disabling Authentication

1.22.1.11. Disabling Authentication on the Cluster

1. Disable Conductor redundancy (if applicable) as described on page 66.
2. To disable user authentication on the entire cluster, run the configure script through the command line again. This prompts appears:

```
Authentication is enabled. Do you wish to disable authentication?
```

3. Enter Yes.

When you return to the web interface, you will not be prompted to log in, and there will be no menu item to enable or disable node authentication (as described in the Enabling Local Authentication on a Worker Node section on page 54).

4. If applicable, re-enable Conductor redundancy as described on page 65.

1.22.1.12. Disabling Authentication on Worker Nodes

1. From the primary conductor web interface, choose Cluster>Nodes, then and click and hold down Tasks and choose Disable Node Authentication.

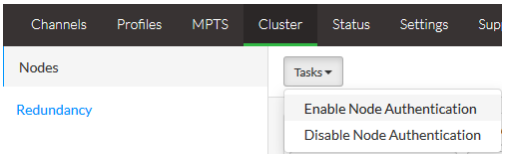

   The Disable Node Authentication screen appears.

2. Enter the built-in, remote terminal “elemental” username and the default password (if not changed by admin). See Appendix A. Default Password Information on page 68 for more information.

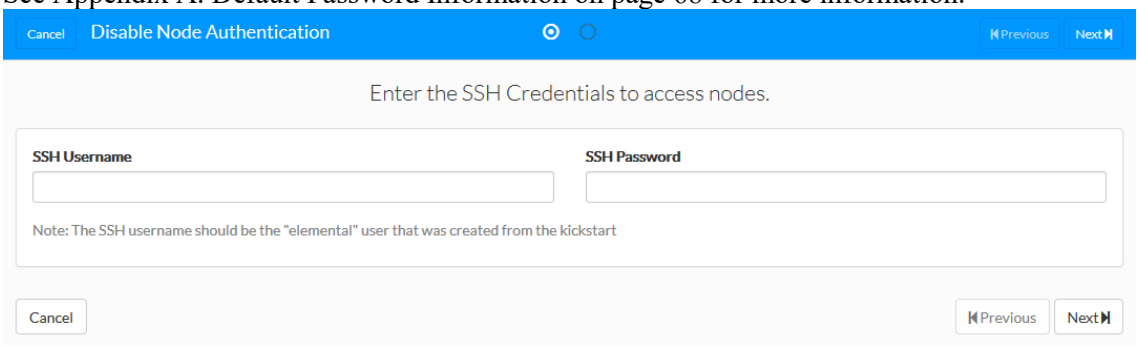

3. Click Configure Now.

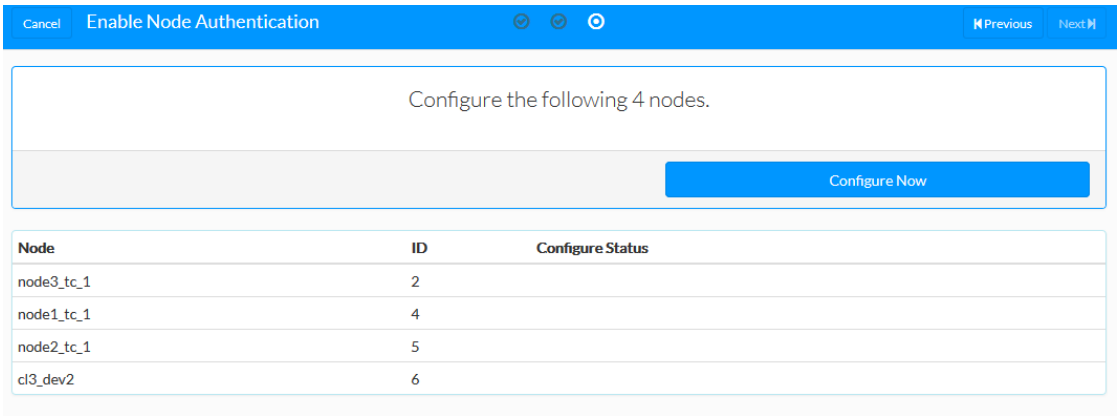

   User authentication is disabled on all the worker nodes in the cluster.

1.22.1.13. Changing and Deleting Users

1.22.1.14. User Self-management on Worker Nodes

1.22.1.15. Adding New Nodes to the Cluster

1.23. Enabling Conductor Redundancy

Perform this procedure on	Primary Conductor node

As the last step in the initial deployment, and only if you are implementing Conductor redundancy, you must enable Conductor redundancy.

When redundancy is enabled, note that:

- The Conductor node you were accessing when you clicked the Enable button is set up as the primary Conductor. The other Conductor node is set up as the backup Conductor.
- The virtual IP address (shown on the screen) is activated. The primary Conductor registers with this VIP as the primary node. This means that:
  - The primary Conductor node becomes inaccessible through its local IP address. The web page you are using is automatically redirected to the virtual IP address, but you will still be accessing the primary Conductor node.
  - The backup Conductor node becomes completely inaccessible: if you enter the IP address of the Conductor node, you will be automatically redirected to the virtual IP. In other words, you will be accessing the primary Conductor node; it is not possible to access the backup Conductor node.
  - Any time a Conductor failover occurs, the node that is promoted to being the primary re-registers with the VIP, effectively indicating “I’m the primary now”.
- The backup Conductor database synchronizes itself with the primary Conductor database. When redundancy is enabled, information is copied over from the primary to the backup.
- Both Conductor nodes must be on the same software version. If they differ, a validation error is received and redundancy will not enable.

Note:	Mount points are not copied over to the backup. Remote servers must be mounted individually on each Conductor.

1.23.1.1. Enabling Redundancy

If you are using a VM, take a snapshot before enabling high availability. See the VMware VSphere help text for more information.

1. On the primary Conductor web interface, access Cluster>Redundancy.
2. In the High Availability field, click the Enable button.

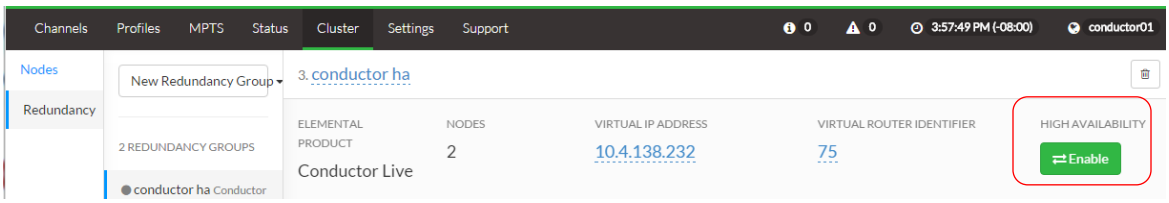

3. Verify that high availability is enabled. From Linux prompts, access the primary and secondary Conductor nodes with username “elemental” and the password as described in Appendix A. Default Password Information on page 68.
4. In the remote terminal session for each Conductor, enter the following command to verify that the service is running:

   ```
   [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 out for the following INFO lines on the primary Conductor node:

   ```
   CONDUCTOR: Initializing environment
   I, [2015-11-13T04:37:54.491204 #4978]  INFO -- : Configuring the HA environment
   I, [2015-11-13T04:37:54.660644 #4978]  INFO -- : configuring keepalived

   .
   .
   .
   I, [2015-11-13T04:38:03.905069 #4978]  INFO -- : Elemental Conductor is ready
   ```

   Ensure the secondary Conductor is also ready.
5. Enter Ctrl+C to exit the tail command.
6. 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.
7. Confirm that you see this line on both Conductors:

   ```
   database system is ready to accept connections.
   ```
8. Enter Ctrl+C to exit the tail command.
9. Type the following command to exit the session as the sudo user:

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

The cluster upgrade is now complete.

1.23.1.2. Disabling Redundancy

If you are using a VM, take a snapshot before disabling high availability. See the VMware VSphere help text for more information.

Warning:	Disabling occurs immediately, but enabling always involves a wait, so do not disable redundancy without a good reason!

1. On the primary Conductor web interface, access Cluster>Redundancy.
2. In the High Availability field, click the Disable button.

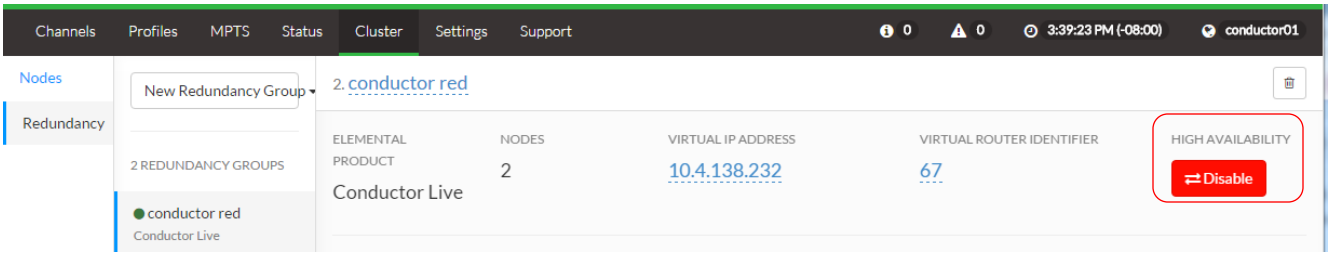

3. Verify that high availability is disabled. From Linux prompts, access the primary and secondary Conductor nodes with the username "elemental" and the password, as described in Appendix A.
4. In the remote terminal session for each Conductor, enter the following command to verify that Conductor high availability 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 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
```

Ensure the secondary Conductor is also ready.

5. Enter Ctrl+C to exit the tail command.
6. 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.

7. Confirm that you see this line on both Conductors:

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

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

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

If you are separating the Conductors into two separate clusters, proceed to the next step. Otherwise, exit the session as the sudo user on the secondary Conductor as well.

Warning: You must perform the next step if the secondary Conductor is being moved to a different cluster. This ensures that the settings are in-sync across all systems.

10. On the secondary Conductor, run the configure script with the clean database option:

```
[elemental@hostname ~]$ sudo ./configure -c --skip-all -xeula
```

11. When complete, exit the session as the sudo user:

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

APPENDIX A. DEFAULT PASSWORD INFORMATION

The default password for the “elemental” username varies based on the version of software that you are running. We highly recommend that you change the default password to provide further security.

1.23.1.3. Changing the Password

There are no specific password criteria but remember:

- Passwords are case-sensitive. Using a combination of numbers, punctuation, and upper and lowercase letters will increase the security of your password.
- AWS Elemental has no record of your password. Anyone contacting Support must be aware of the password in case AWS Elemental needs to access your system for troubleshooting.
- The password must be changed on each node individually. Updates to the Conductors or any other node do not push to the entire cluster.

1. From a Linux prompt, log in with username “elemental” and default password. Run the following script:

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

2. At the prompt, enter and confirm the new password.

   ```
   Changing password for user elemental.
   New password:
   Retype new password:
   ```

   The following response verifies the update.

   ```
   passwd: all authentication tokens updated successfully.
   ```

3. Update the Samba password with the following command.

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

4. At the prompts, enter password information:

   ```
   Old SMB password:
   New SMB password:
   Retype new SMB password
   ```

   The update is verified.

   ```
   Password changed for user elemental
   ```

5. Update all applicable nodes separately and provide the new password(s) to users who are responsible for contacting Support.