aws elemental

AWS ELEMENTAL CLOUD CONFIGURATION GUIDE

AWS ELEMENTAL CLOUD VERSION 2.4.8

AWS Elemental
1320 SW Broadway
Portland, Oregon, 97201

+1 503 222 3212
www.elemental.com

Contents

About this Manual

Phase 2 of Implementation

This guide provides detailed information about the second implementation phase of AWS Elemental Cloud. Topics in this guide focus on deployments and services and include:

- Creating deployments to provision services within.
- Enabling, creating, starting, and configuring services to support video processing and delivery in the cloud.
- Performing common management functions such as upgrading deployments, changing deployment settings, and viewing deployment usage information.

Prerequisite Knowledge

It is assumed you are familiar with:

- Connecting to the AWS Elemental Cloud web interface using your web browser.
- Your processing and delivery requirements.

1 OVERVIEW

1.1. AWS Elemental Cloud Charges

Charges are incurred in AWS Elemental Cloud when:

- Features packages (such as Advanced Audio, Audio Decoding, and HEVC) are added to a service
- A service is started

  Services are charged from the time they are started, until they are stopped, regardless of if content is being processed or not.

1.2. Deployments versus Services

A *deployment* represents a virtual data center, holding services that are being used in a video processing and delivery workflow. Multiple deployments can be configured to support a variety of workflows, and can be configured to support various node limits and storage terabytes for each service.

A *service* represents a virtual cluster of AWS Elemental video processing or delivery solutions that is enabled within a deployment. Available services are:

- AWS Elemental Live
- AWS Elemental Server
- AWS Elemental Delta
- Network Address Translation (NAT)

AWS Elemental Live and AWS Elemental Server are each managed by a Conductor, while AWS Elemental Delta runs as a collection of independent nodes.

1.3. Where to Work

When you are initially creating the deployments and services, you are working in either the AWS Elemental Cloud web interface, or through REST API calls to AWS Elemental Cloud.

After services are created in their deployments, you work through a variety of management systems to configure events, jobs, profiles, and so on. This guide provides introductory information about accessing those management systems in the Configuring Services section starting on page 15. The main focus of the document, however, is how to work in the AWS Elemental Cloud web interface.

1.4. Lifecycle of a Service

The stages of a service in AWS Elemental Cloud are:

The services are *enabled* in a deployment (Enable Services on page 8).

1. The services are *created*, or "spun up" (Creating Services on page 12). They show as *running* in the deployment and are ready to be configured. Content cannot start to be processed until this step is complete.

   This is when charges start to be incurred.
2. The services are *configured* in their management systems (Configuring Services on page 15).
3. The services are active and *running* (and are able to manage processing resources and associated job queue or events) once they have been started.
4. Services remain *running* in the deployment until they are stopped and show as *inactive*.

   Charges are no longer incurred when the services are inactive.

1.5. Note about Regional Data Transfer

Your data crosses regions in the following circumstances:

- You intentionally created your deployments across different regions
- When you log-in to the web interface or any services, your username and hashed password cross regions to be communicated back to AWS Elemental Cloud

2. CREATING A DEPLOYMENT

Step A: Create the Deployment

1. From the Deployments home page, click the **Add Deployment** button at the bottom right of the page.

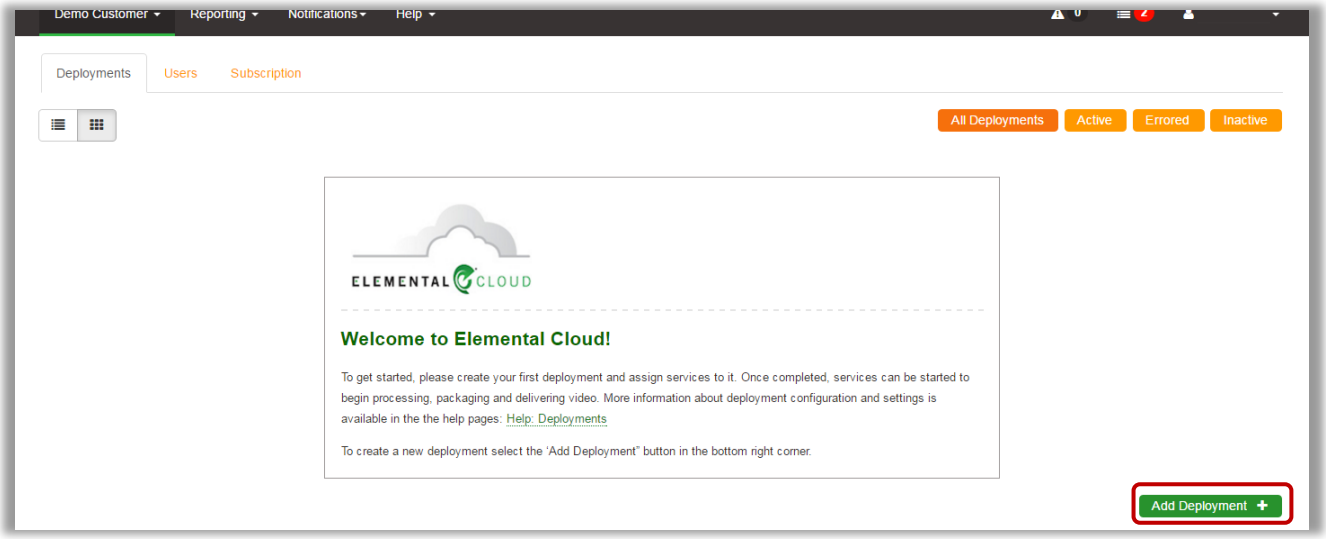

   The Add Deployment page opens.

2. Complete the Deployment Info and Deployment Options sections:
   - Create a unique, descriptive **Deployment Name.** We recommend that you don't use sensitive identifying information when naming your deployment since this information may be later included in diagnostic logs.
   - Indicate the **Time zone** that services will display.
   - Select the **Region** where services will be provisioned.
   - In **Mode**, indicate if the deployment will be high availability or reduced redundancy.

     See the Mode section on page 10 for more information.

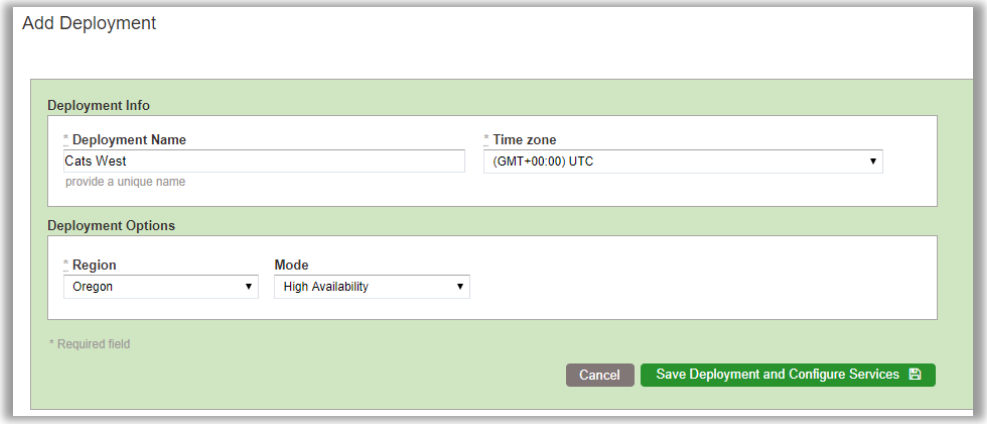

3. Click the **Save Deployment and Configure Services** button.

   You are taken to the Services tab of the deployment.

Step B: Enable Services

Enable services to make them available in the deployment. Note that the services are not created until they are started as described on page 12.

To enable services:

1. For your required services, click to toggle services from **OFF** to **ON**.

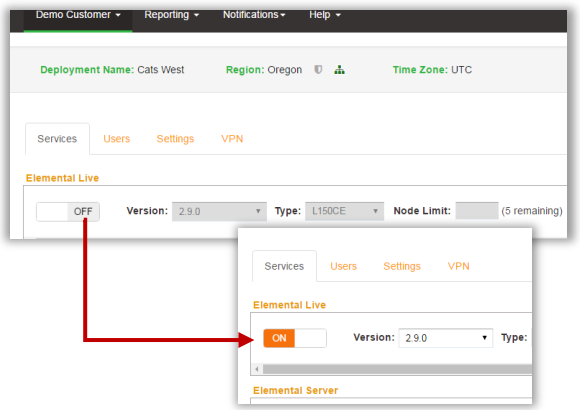

2. Complete the additional information as needed for each enabled service (such as versions, limits, and service options).

   Note that:

   - You cannot select End of Life versions.
   - Within the same deployment, you cannot select AWS Elemental Server and AWS Elemental Live services with different database versions. Once a software version is selected on a service, only compatible versions are available for selection on other services.
   - The **Node Limit** fields on AWS Elemental Live and AWS Elemental Server indicate how many nodes are available for the service. The maximum amount you can select is either 25 or dependent on the available nodes enabled in the subscription.
   - Clustered mode Deltas have 2, and *only* 2, nodes. You can indicate 1-6 nodes in Independent mode. Note that the 6-node maximum is for the entire subscription so you might be allowed to deploy fewer than 6 nodes in this deployment.
   - Additional options on AWS Elemental Live, Server, and Delta are selectable based on your subscription.
   - **Increased Node Storage** on AWS Elemental Server and AWS Elemental Live adds 500GB to the nodes' hard drive. This storage is accessible by the nodes only. Content is deleted when services are stopped.
   - If you do not enter an IP range when enabling Network Address Translation, all non-VPC traffic will be routed through the NAT.

3. Click the **Update** button to save your changes and move on to address deployment users, settings, or VPN connection. Go to:

   - For users, see Step C: on page 9,
   - For settings, see 0 on page 10, or
   - For VPN access, see Step E: on page 11

   Otherwise, click the **Save and Exit** button to be returned to the Deployments page. Skip to Creating Services on page 12.

Step C: Set Users

The Users tab displays all users who have access to this deployment. By default, users with an Administrator or Owner main role and/or Admin service role are populated on this tab.

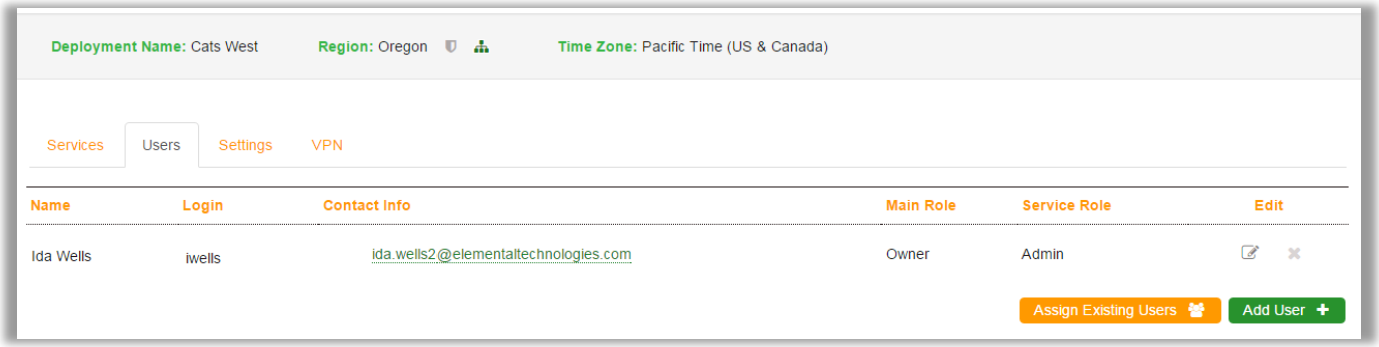

You can edit, delete, assign, and add users from this page.

Warning	Edit, delete, and add user actions take place at the subscription-level. If you change a user's role, delete a user, or add a new user, the updates take place throughout your subscription (example: if you delete a user, it removes their access to AWS Elemental Cloud.

If a user does not have an Administrator/Owner or Admin role, they are not automatically added to new deployments. To assign an existing user to this deployment:

1. Click the deployment's **Users** tab (as pictured above).
2. Click the **Assign Existing Users** button at the bottom right of the page.

   A list of existing users is displayed alphabetically by last name.
3. Locate the user(s) to add and place a check in the box to the left of their name(s).

4. Click the **Save Selected Users** button.

Note:	A user's Service Role persists throughout all deployments in the subscription. Example: You cannot have a user be Manager for one deployment and Viewer for the rest. When you add a user to a deployment, you cannot indicate the Service Role because that is managed at the subscription-level and does not vary between deployments.

Step D: Specify Settings

Required settings are populated when the deployment was created. The deployment's Settings tab allows you to further specify non-required settings: high availability, VPC peering, and IP whitelisting.

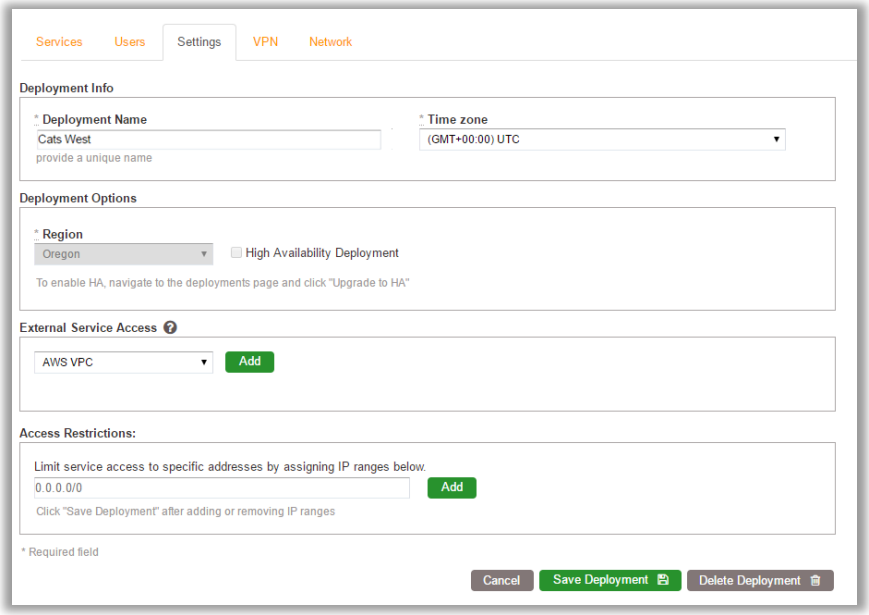

See the following sections for more information.

Mode

In the Mode field, **High Availability** enables resiliency for the service managers in your deployment by transparently replicating resources across multiple data centers (availability zones) in the same region. This way, if a hardware or data center failure occurs, your services remain running in a geographically diverse location.

Alternatively, you can select **Reduced Redundancy** to deploy a single service manager for the deployment.

High availability is selected by default when a deployment is created (whether through the web interface or REST API). The Settings tab allows you to switch modes up until services have been added to the deployment. High availability cannot be disabled after services have added to the deployment. If you need to *enable* high availability after services are added, see Upgrade to High Availability on page 29.

VPC Peering

VPC peering establishes a private and secure connection between and AWS Elemental Cloud deployment and a virtual private cloud (VPC) within another AWS account. The External Service Access section of the Settings tab allows you to enable VPC peering with your AWS VPC. VPC peering is not enabled by default. Contact your AWS Elemental Sales Manager for information about enabling VPC peering.

To enable VPC peering from the web interface, see below. To use REST API to enable VPC Peering, see the REST (API) tab of the AWS Elemental Cloud Help pages.

1. Click the **Add** button on the deployment Settings tab.
2. Complete the fields:
   - **Target VPC**: your AWS VPC ID
   - **AWS Account Number**: ID for your AWS account
   - **CIDR**: IP address for your AWS VPC (address must not overlap with 10.0.0/16; e.g. 172.16.0.0/16)

3. Once the fields are complete, save the peering request by clicking the **Save** button at the bottom of the page.
4. The connection request is initiated when services in the AWS Elemental Cloud deployment are started (page 12). No further VPC peering actions take place until then.
5. From your AWS console, accept the peering request from AWS Elemental Cloud and enter the 10.0.0/16 route to map your VPC network to the AWS Elemental Cloud deployment.

IP Whitelisting

The Access Restrictions section of the Settings tab allows you to restrict what IP addresses have access to the services in your deployment (whitelisting).

To indicate what IP addresses have access, enter the range of IPs and click the **Add** button.

Step E: Configure VPN

Configure VPN access to establish a connection between the deployment and a private network with an OpenVPN endpoint (such as a key server in a private datacenter).

To configure a VPN connection:

1. Generate server and client keys for the VPN server.
2. On the deployment's VPN tab, enter the IP address of the VPN server and the certificates and key information you generated in step 1.

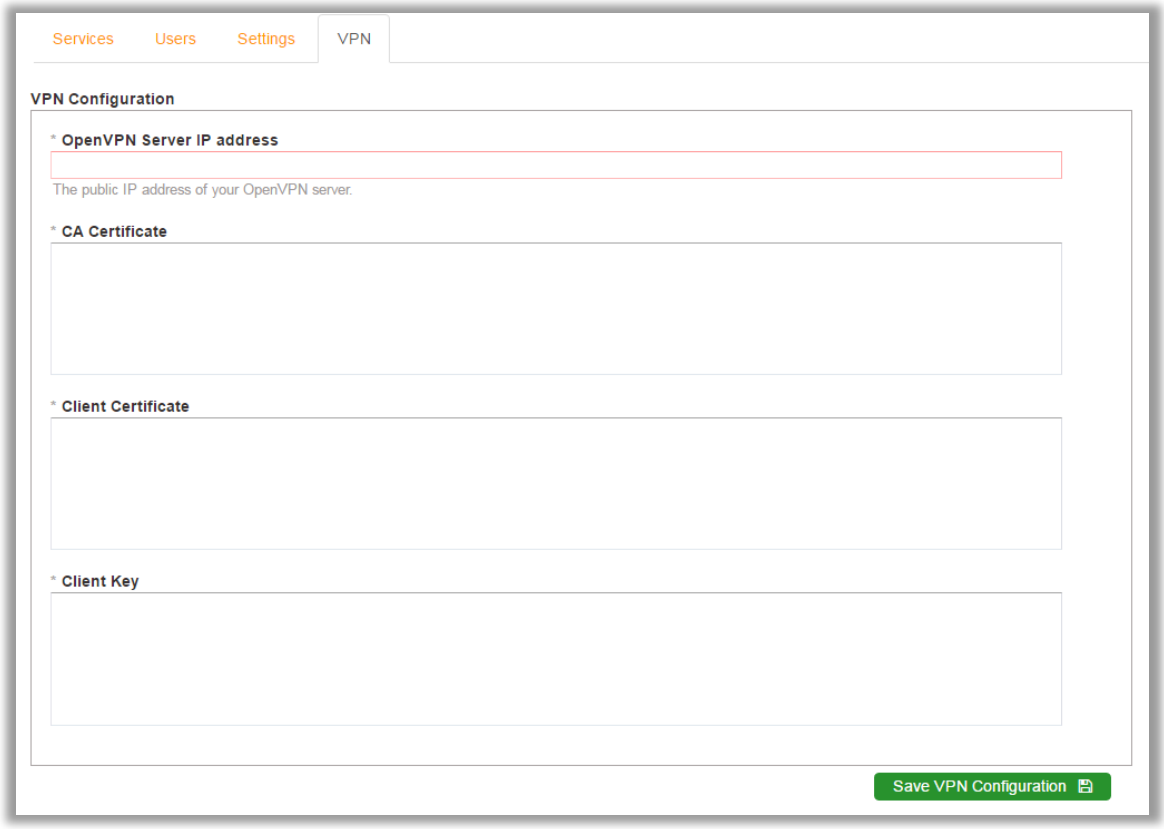

3. Click the Save VPN Configuration button.

   When the services are started, they will automatically connect to the VPN tunnel.

3. CREATING SERVICES

Once the deployment is started and services are added, create the services. Note that until services are configured, processing will not take place. See Configuring Services on page 15 for more information.

To:

- Start all services in the deployment at the same time, see page 12.
- Start services individually or at different times, see page 13.

Also see page 13 for information about stopping services.

3.1. Create All Services

To start all inactive services in the deployment, click the **start deployment** button in the header of the deployment.

Note	A service incurs charges as soon as it is started, up until it is stopped.

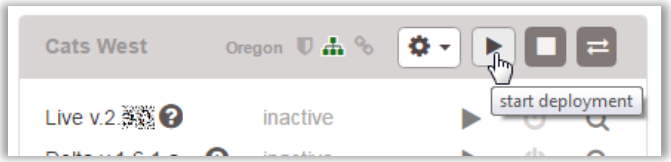

The services statuses progress from *creating* to *starting*.

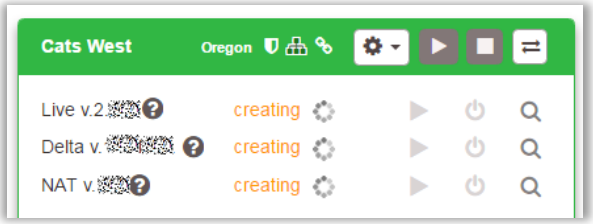

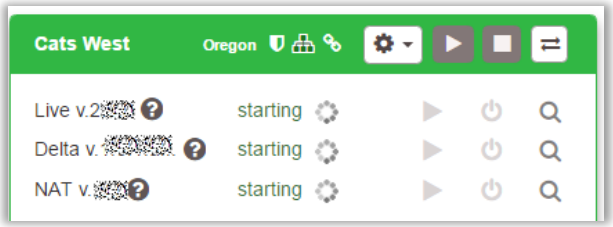

After 10 to 15 minutes (45 for high availability deployments), the services are created and can be configured. See Configuring Services for further information.

Reminder:	Content is not being processed until after the services are configured. At this stage, the *running* status indicates that the services have been created and can be accessed for configuration.

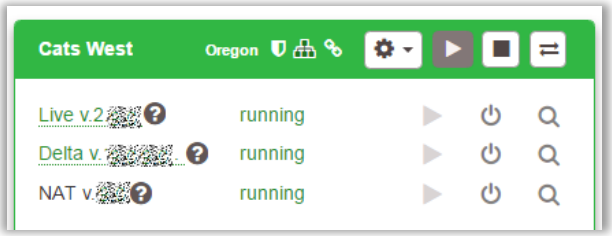

3.2. Create Individual Services

To start services separately, click the **create service** button for the service being started.

Note	A service incurs charges as soon as it is started, up until it is stopped.

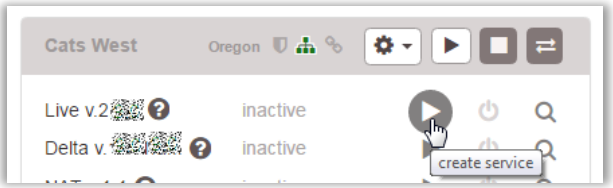

The service status progresses from *creating* to *starting*.

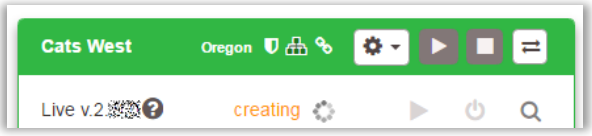

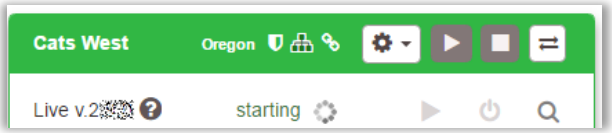

After 10 to 15 minutes (45 for high availability deployments), the service is running and can be configured. See Configuring Services for further information.

Reminder:	The service is not active until after it is configured. At this stage, the *running* status indicates that the service has been created and can be accessed for configuration.

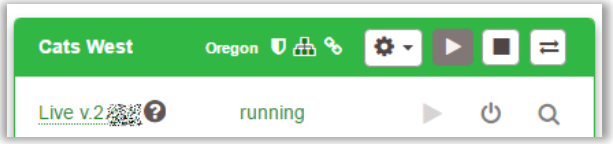

3.3. Stop Services

Stopping – or shutting down - a service temporarily suspends the processing on that service. Additionally, the settings on the services are removed, including:

- Cloud settings (such as failover, Autoscaler, and number of nodes)
- Timezone (defaults back to the timezone indicated on the AWS Elemental Cloud web interface)
- Auto archive and poll rates

Events, jobs, and filters are paused but not removed when the service is stopped.

Stop All Services

To stop all services in the deployment, click the **stop deployment button** in the deployment header.

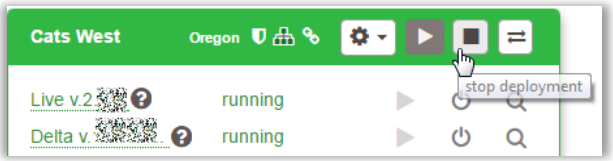

Stop Individual Services

To stop a specific service, click its corresponding **shutdown service** button.

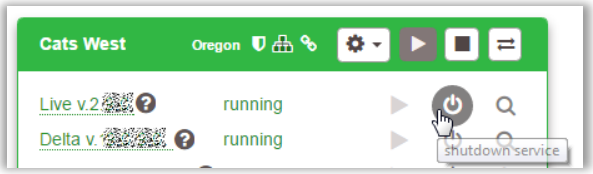

4. CONFIGURING SERVICES

4.1. AWS Elemental-branded Services

The AWS Elemental-branded services in the deployment are similar to “core” AWS Elemental products on the ground, but with the added benefits of being more resilient and dynamic, as afforded by the AWS Elemental Cloud platform. When you created services in the deployment, most of the configuration was done through the Settings and other screens. The remaining configuration is done from the node or the service’s management system.

See product-specific documentation for getting started with jobs, events, and filters:

- AWS Elemental Conductor File API and User Guide
- AWS Elemental Conductor Live API and User Guide
- AWS Elemental Delta User Guide

AWS Elemental Server

Autoscaler

Autoscaler monitors how many jobs are in the queue and determines how many Server nodes are required to handle the workload. To configure Autoscaler:

1. Click the **Server** link in the deployment. The service must be running for the link to be available.

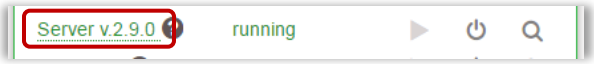

   Conductor File is opened in a new tab.

2. Log-in with your AWS Elemental Cloud credentials.
3. Go to **Configuration** > **Autoscaler**.

   The Autoscaler Settings page opens.

4. In the **Maximum nodes** field, enter the highest amount of Server nodes that can be running at once. Note that this number cannot exceed the amount of Server nodes authorized when the service was enabled.
5. Optionally, use the **Minimum nodes** to set the fewest nodes to be running at any time. If any number greater than 0 is entered, at least 1 node will always be running and incurring charges, regardless of the jobs in queue.

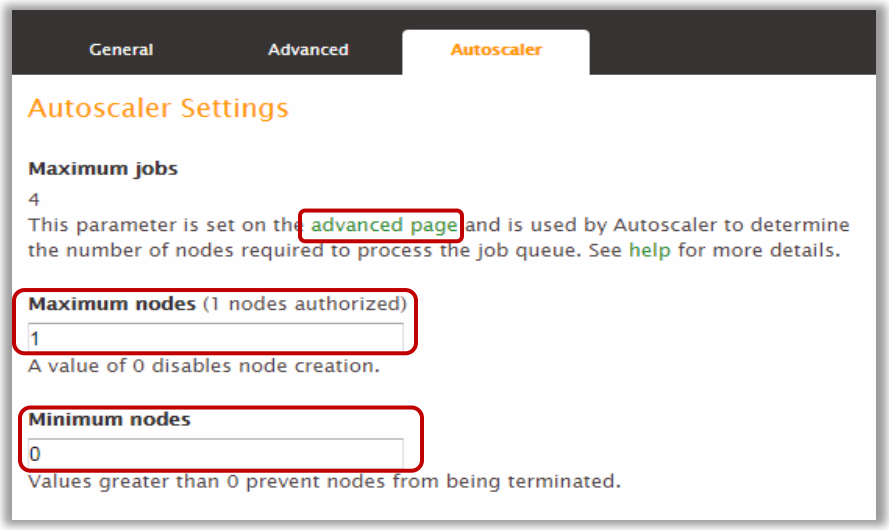

6. If you need to adjust the maximum number of jobs allowed per node, click the **advanced page** link in the Maximum jobs section.

   The Cluster Sequencer Settings page opens to the Server Nodes tab.
7. In the **Maximum jobs that can be running** field, enter the preferred value.

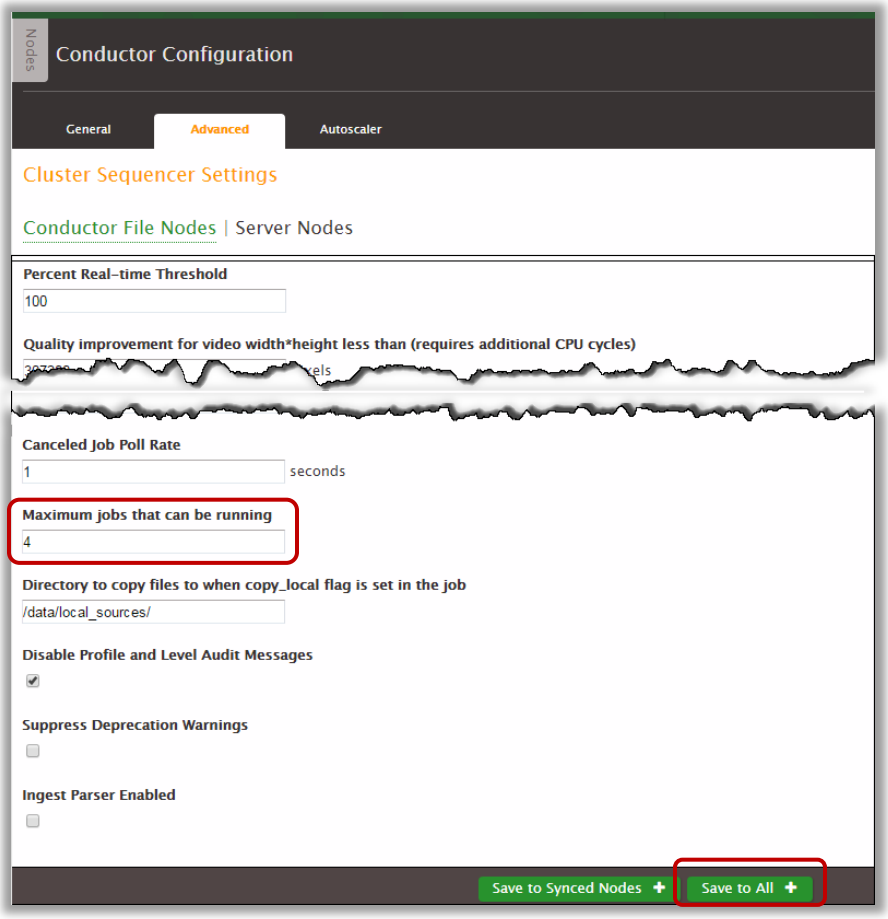

8. Click the **Save to All** button.

AWS Elemental Live

Number of Nodes

Conductor Live manages the AWS Elemental Live events and profiles. Prior to configuring events in the Conductor, you must indicate the number of nodes it will be managing. To set the nodes:

1. Click the **Live** link in the deployment. The service must be running for the link to be available.

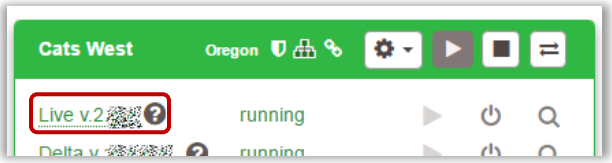

   Conductor Live is opened in a new tab.

2. Log-in with your AWS Elemental Cloud credentials.
3. Go to **Configuration** > **Cloud**.

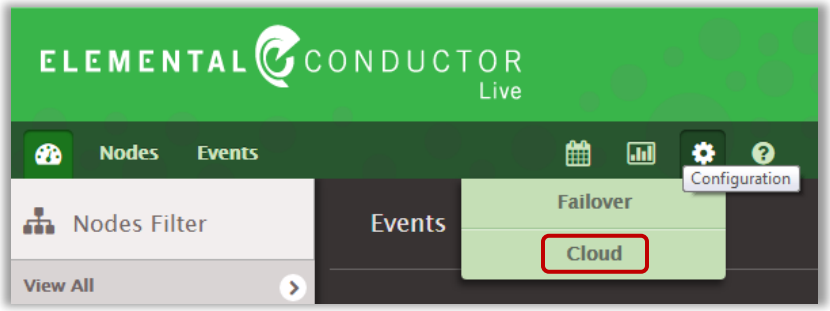

4. In the **number of nodes** field, set the nodes to be added to the cluster being controlled by the Conductor. If the number of active nodes falls below this number, Autoscaler will start a new node.

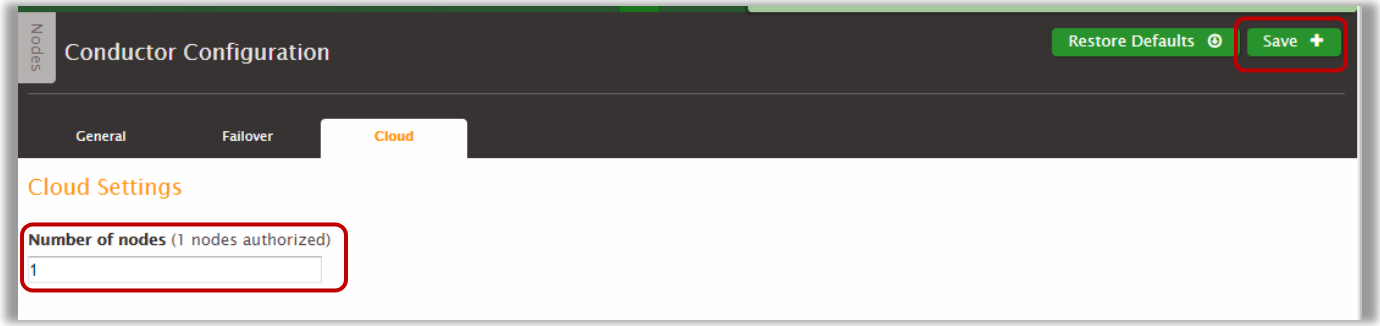

   The parenthetical text shows how many Live nodes you authorized when the deployment was created.

5. Click the **Save** button.
6. Configure events as needed.

Inputs

While AWS Elemental Live in the cloud supports the same input formats as AWS Elemental Live on the ground, we recommend you use either:

- RTP with FEC,
- RTMP, or
- HLS

IP Addresses

Generally, AWS Elemental Live nodes in the cloud are assigned a new IP address whenever the node is stopped and restarted. To support certain workflow requirements, Live node IP addresses can be reserved so that assignments persist even when a node is restarted. Note:

- Assigned IP addresses can be found on the AWS Elemental Live service details Node List tab.
- Reserved IPs are set on new services. If an existing service is running when reserved IPs are enabled on the subscription, the service must be restarted to receive an IP assignment.
- IPs are automatically released when:
  - Services failover to a different region. New IPs are assigned when the services are running in the new region.
  - Node limits are reduced.
- If node limits are increased, IPs are allocated to the new services when they are restarted.

To enable reserved AWS Elemental Live IP addresses, contact your AWS Elemental Sales Manager.

4.2. Network Address Translation (NAT)

NAT provides a single, persistent public IP address for all of the services in a deployment.

When NAT is running with other services, non-VPC traffic is routed through a single port, thereby producing one IP address. This is beneficial because services in the deployment are transitory: a new IP address is assigned when a service is restarted. This IP reassignment presents a problem when you rely on the addressing for such purposes as IP whitelisting on servers holding content. Traffic that is routed through the NAT can then be referenced by the NAT route's IP address and not by the IP address assigned directly to the worker node.

Note:	As all indicated traffic must flow through the NAT, performance may be inhibited. Analyze your performance and security needs to determine if NAT will work in your deployment.

NAT does not require further configuration. When the services are started, traffic is automatically routed through the NAT. If you did not specify IP ranges on the Services tab when you enabled NAT on the deployment, all non-VPC traffic will go through the NAT.

To update IP information, see Change Services Settings on page 26.

To gain the public IP address, click the **question mark graphic** next to the NAT service in your deployment. Make note of the public IP. If high availability is enabled, two IPs are shown because subnets are split between two NAT routes. Make note of both IP addresses.

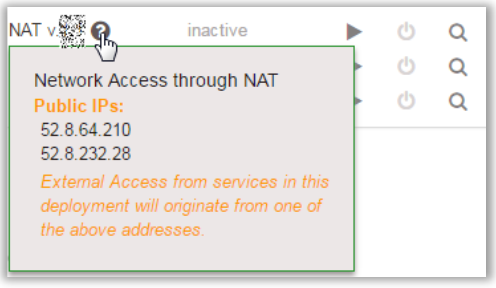

5. MANAGING THE DEPLOYMENT

This section details common management functions that are performed after services are created and started.

5.1. View Service Details

To view the details of a service, click the **magnifying glass** button that corresponds to that service to display the service detail page. On it, you will find basic information entered when the service was created, as well as tabs for additional details. See the following sections for information about the tabs provided.

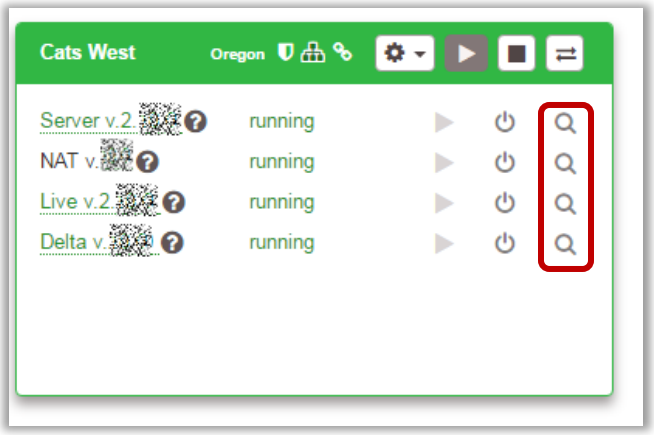

Alternatively, click the question mark next to a service to view its access information.

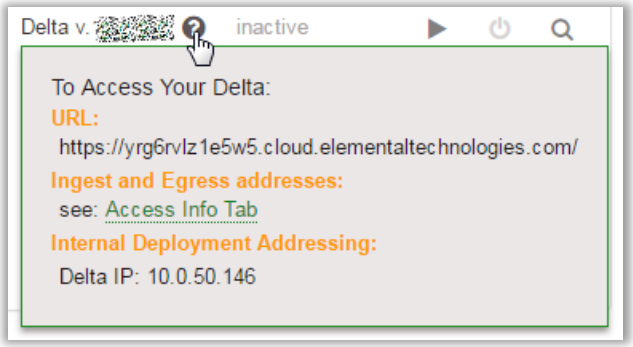

Usage Tab

Available on: all services

Purpose: provides a graph or detailed list about when the service was used and the statuses it encountered.

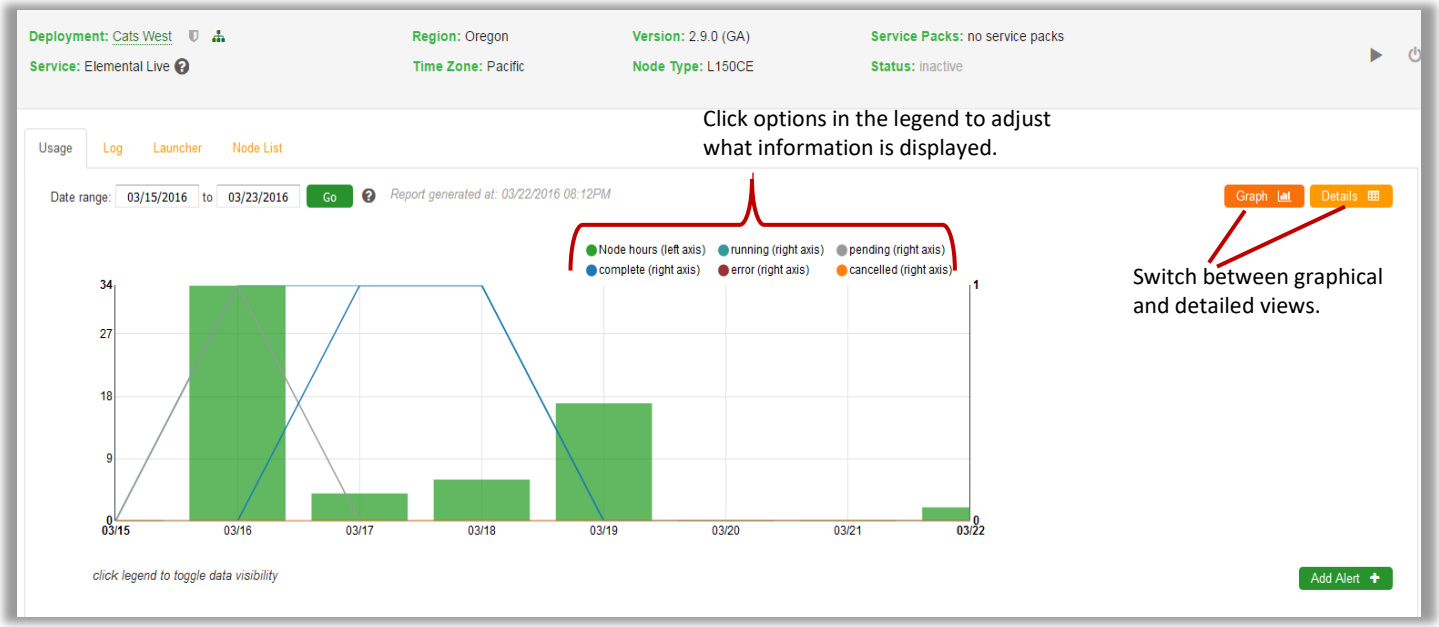

Log Tab

Available on: all services

Purpose: displays XML logs for all activity on the selected service.

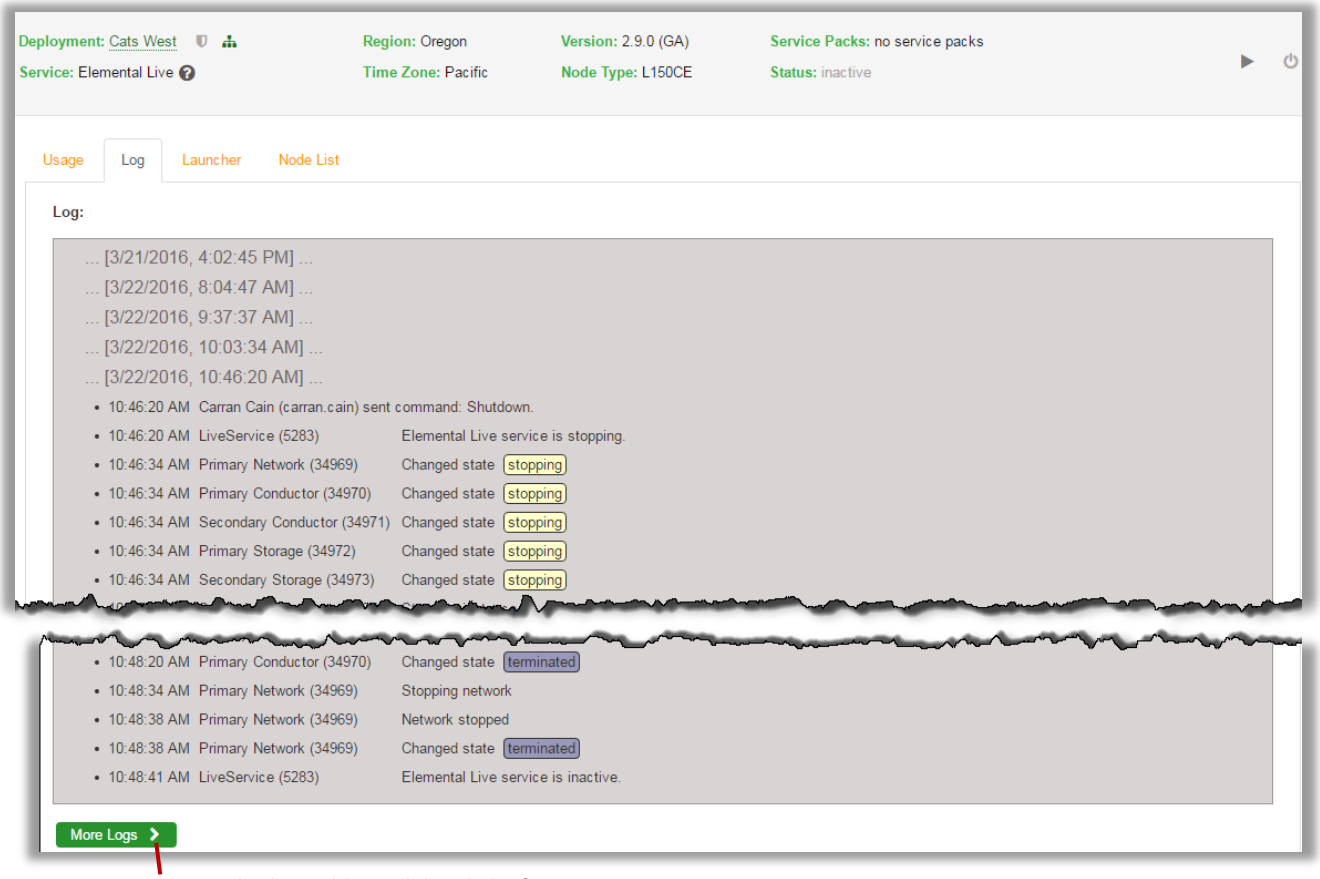

More Logs displays additional detailed information.

Launcher Tab

Available on: AWS Elemental Live, AWS Elemental Server

Purpose: creates a sample video to test the service's configuration. Note that you can only launch jobs/events for services that are currently running.

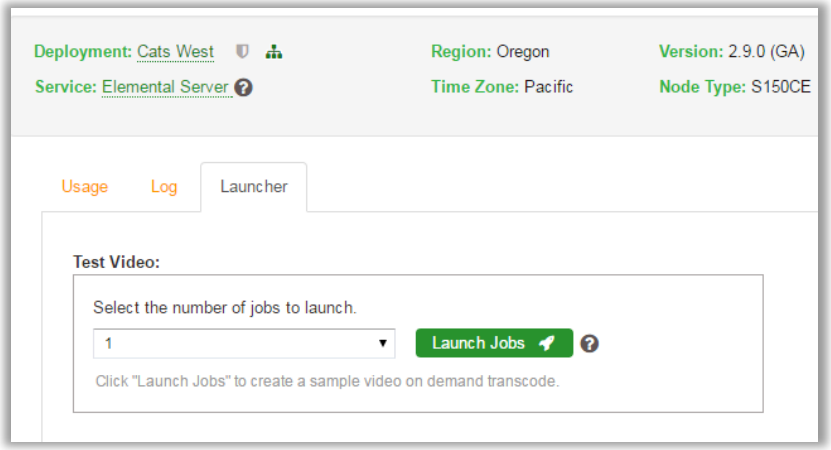

Node List Tab

Available on: AWS Elemental Live

Purpose: displays the IPs reserved for the Live nodes.

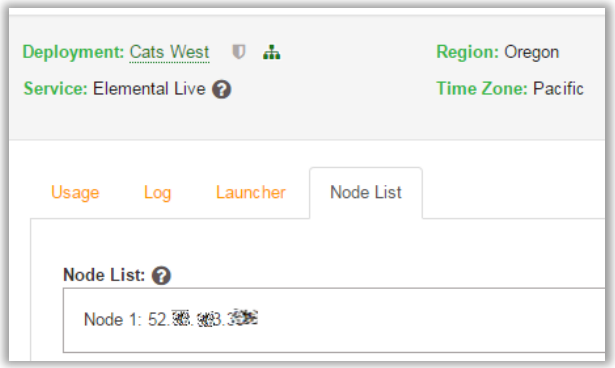

Access Info Tab

Available on: AWS Elemental Delta

Purpose: displays the Delta properties, such as external and internal content ingest addresses.

Note that different information is displayed, based on if you are using Deltas in Independent mode or Clustered mode. See Working with Multiple Deltas in AWS Elemental Cloud for more information.

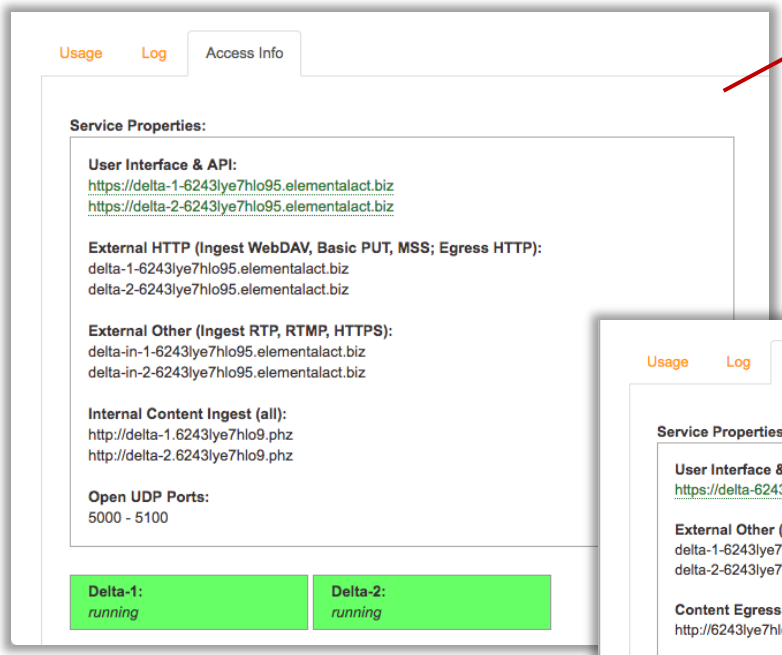

Independent Mode access info

Note that each Delta node has its own set of IP addresses (delta-1 and delta-2 in this image).

Additionally, the state of each Delta node is captured at the bottom of the page ("running" in the image).

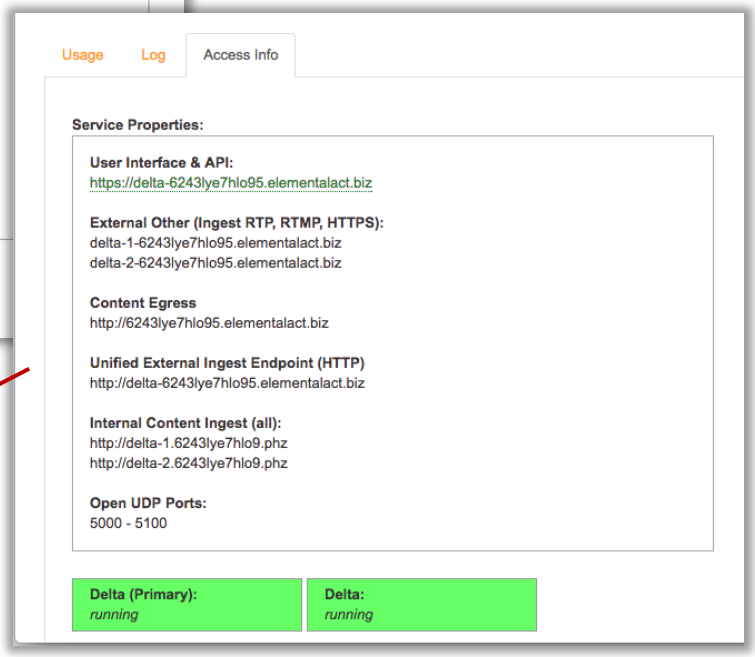

Clustered Mode access info

Note that the ingress and egress load balancer addresses are listed where applicable, rather than addresses for each node.

Additionally, the state of each Delta node is captured at the bottom of the page ("running" in the image).

5.2. Failover the Services

Test the deployment's failover capabilities by manually initiating failover on running services:

In the deployment's header, click the **toggle failover** button.

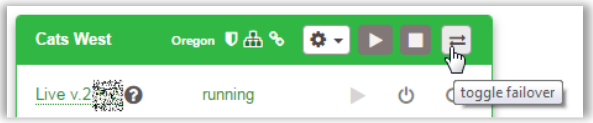

All active services are stopped and the banner at the top of the page indicates that services are being relocated.

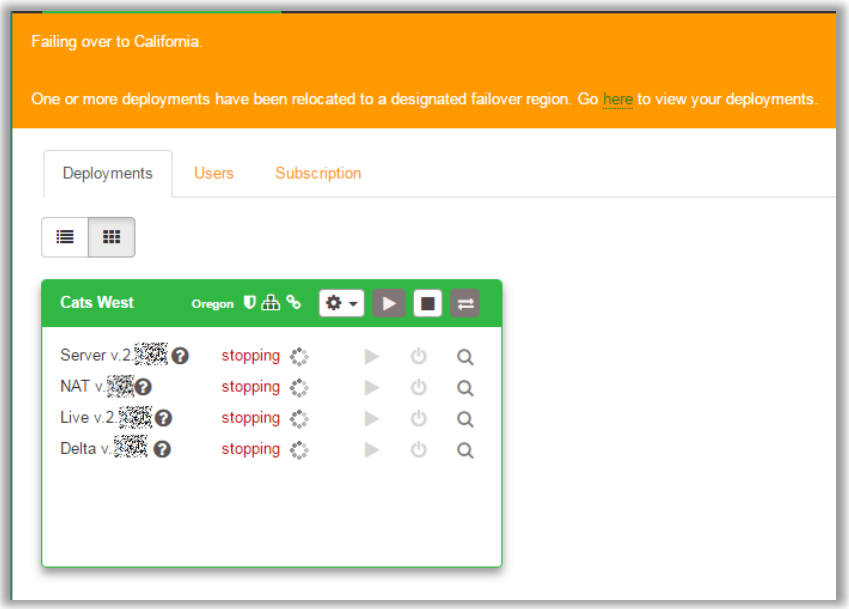

Similar to when services are initially started, they go through a creating and starting process during failover.

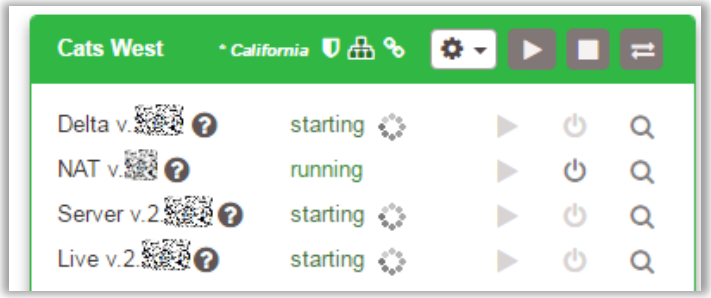

When failover is complete, all previously active services are running and the region is updated in the deployment header. If services are running in a failover region, an asterisk (*) is shown beside the region.

Note that you cannot failover inactive services.

5.3. Disable a Service

Disabling a service removes it from the deployment. This means that all settings, configurations, and database backups are deleted. Perform this action with care. Note that if you need to temporarily stop a service, it is better to stop it, as described in Stop Services on page 14.

To disable a service:

1. Stop the applicable service as described on page 14.
2. Select **Settings** from the **configuration** drop-down in the deployment header.

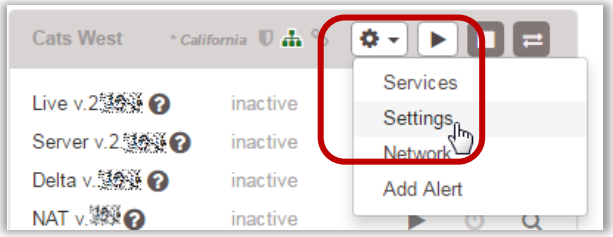

   You are taken to the Settings tab of the deployment.
3. Locate the service to be disabled and click to toggle services from **ON** to **OFF**.

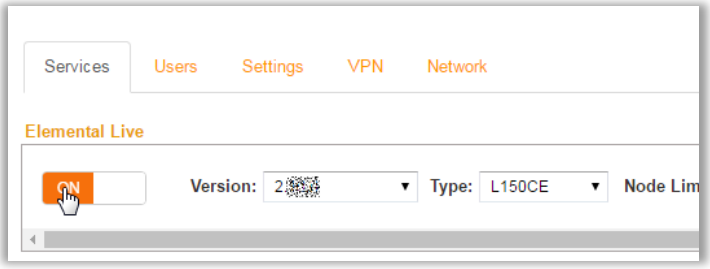

4. The service has been started before, a warning is displayed. Click OK on it to proceed.

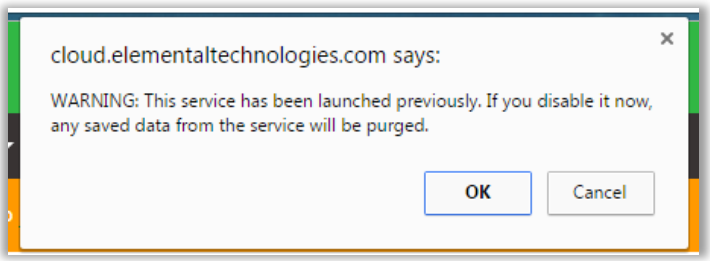

5. Click the **Save and Exit** button.

   You are returned to the Deployments page and the disabled service is no longer displayed.

5.4. Delete Deployment

A deployment can be removed from your subscription to free up resources. Note that when a deployment is deleted, all database information, profiles, presets, and so on are also deleted. Perform this action with care.

To delete a deployment:

1. Stop all active services in the deployment as described on page 14.
2. Select **Settings** from the **configuration** drop-down in the deployment header.

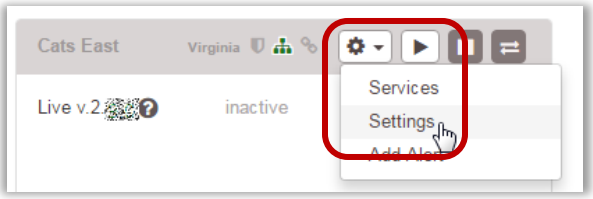

   You are taken to the Settings tab of the deployment.
3. In the bottom right corner of the page, click the **Delete Deployment** button.

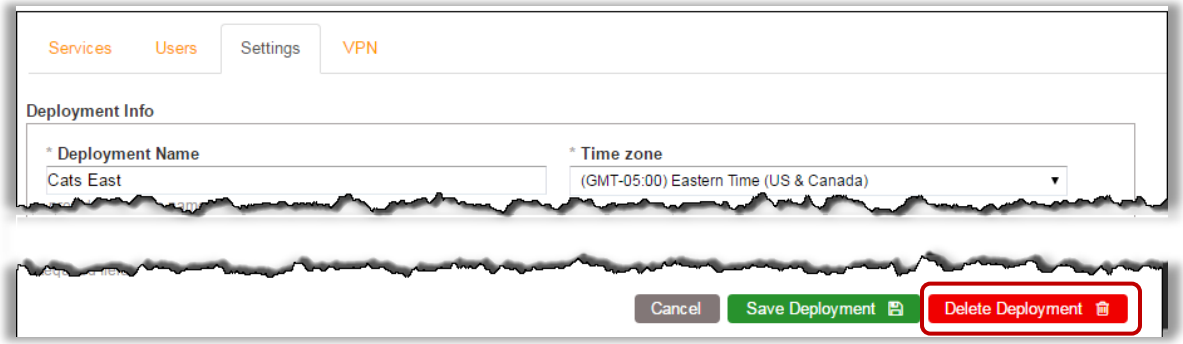

4. Confirm the deletion.

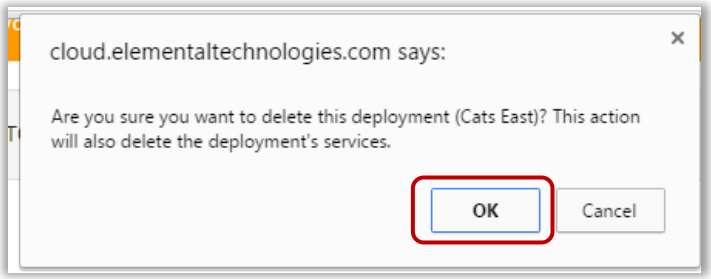

   The deployment is removed.

5.5. Change Services Settings

To enable/disable services, or change settings on services (such as node limits and additional options):

1. Select **Services** from the **configuration** drop-down in the deployment header.

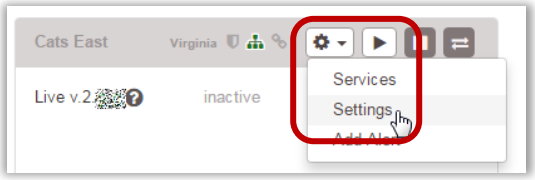

   The deployment's Services tab is displayed.
2. Make changes as needed.
3. Click the **Save and Exit** button when all required adjustments have been made.

5.6. Upgrade Service Versions

When new service versions are available, you receive notification in the System Messages in the AWS Elemental Cloud menu bar, or under **Notifications** > **System Messages**.

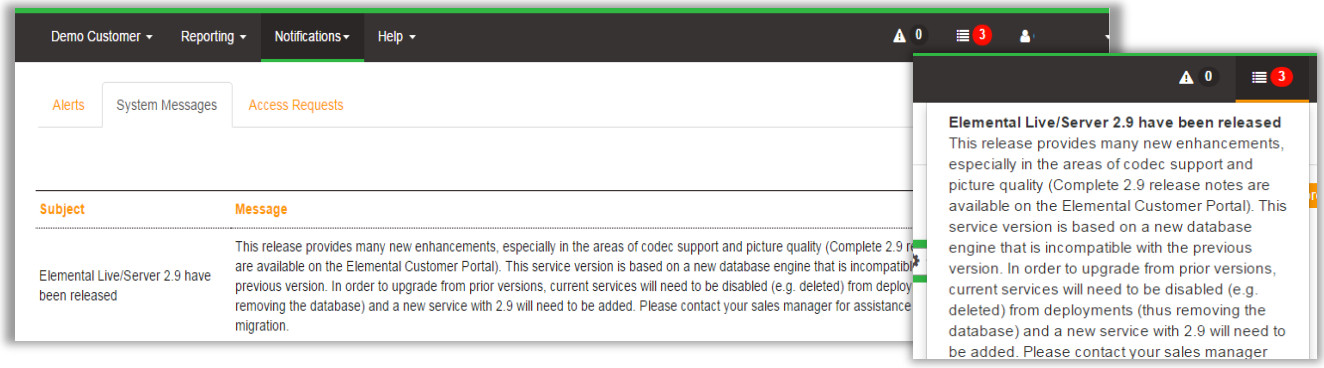

Note that:

- You can upgrade up to two "dot" versions. This means that if you're on 2.6.x, you can upgrade to 2.8.x but not to 2.9.x. Conversely, if you're on 2.6.1, you *can* upgrade to 2.6.8 since "2.6" is the dot version information.
- You cannot upgrade a service when there is an active access request associated with the deployment.
- You cannot downgrade an enabled service. To return to an earlier version, see To Use an Earlier Service Version on page 28.
- Upgrading service versions does not affect your databases unless specifically stated in the release notes. However, backing up pertinent information (such as profiles and presets) provides additional security.
- The URLs for your services will remain the same, but internal deployment addressing will change. If you have reserved Live IPs, these will remain the same.
- Since all services have to be stopped and restarted, it may take up to 30 minutes for the upgrade to high availability to complete.
- You can test the new version by creating a test deployment prior to upgrading. See the next section for more information.

Test New Version

1. Download the XML for the profiles, presets, events, and jobs you will be using for testing.
2. Create a new deployment using the instructions from Creating a Deployment on page 7.
3. Enable the services with the new versions and configure services to duplicate your existing deployment (upload the XML files you downloaded in step 1).
4. To test the upgrade, use the Launcher as described on page 21.
5. When you have verified that the new versions work in your workflow, upgrade your existing deployment as described in Upgrade Service Versions.

Upgrade Service Versions

1. Stop the services to be upgraded, as described on page 14.
2. Select **Services** from the **configuration** drop-down in the deployment header.

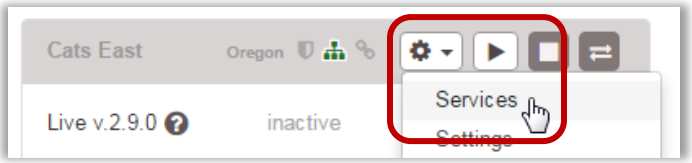

3. In the **Version** drop-down for applicable services on the Services tab, select the version you are upgrading to.

   Note that you cannot downgrade to an earlier version, nor can you select an End of Life version.

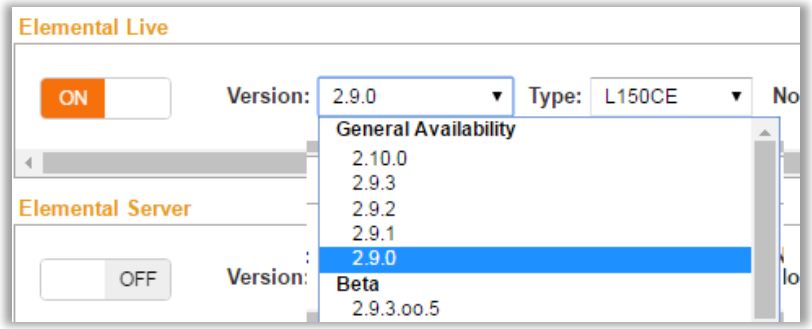

4. When you have selected the new version on all applicable services, click the **Save and Exit** button.

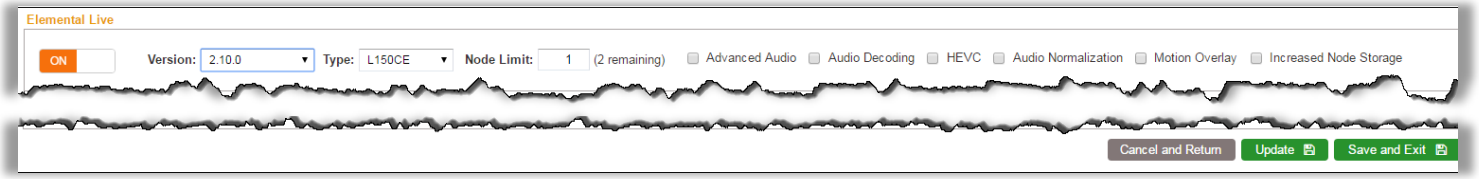

When the services are started again, they will be recreated on the upgraded version.

To Use an Earlier Service Version

You cannot downgrade an existing service. However, if you need to return to an earlier service version, you can create a new service with the selected version.

Note that disabling the service purges all data for the service, including database backups. Download the XML for the service to make setup on the new service version more efficient.

1. On the existing service, download the XML for the profiles, presets, events, and jobs you will be using for testing.
2. In the AWS Elemental Cloud web interface, stop the service if it is running (see page 14).
3. Access the Services tab in the deployment as described on page 14 and click to toggle the service from ON to OFF and click the **Update** button.
4. When the page refreshes, toggle the service back on. Be sure to include the correct node limits and additional options as needed. Click the **Save and Exit** button.

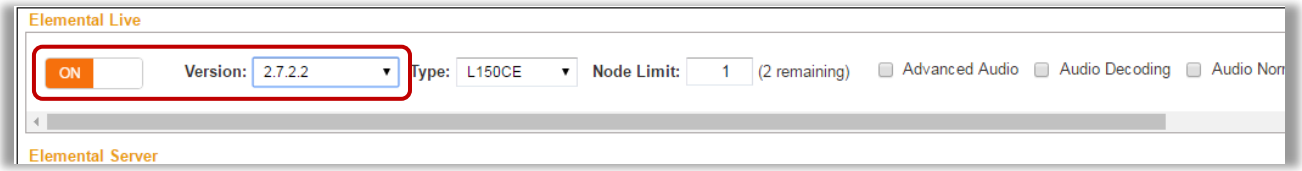

5. Create the service as described in Create Individual Services on page 13.
6. When the service is running, configure it using the XML you downloaded in step 1.

5.7. Upgrade to High Availability

You can upgrade to high availability if you determine after creating the deployment that you need service resiliency and/or regional failover.

Note that:

- Upgrading to high availability removes your existing databases. This means that all filters, profiles, settings, logs, and presets will be lost. Back-up all pertinent information prior to upgrading to high availability.
- The URLs for your services will remain the same, but internal deployment addressing will change. If you have reserved Live IPs, these will remain the same.
- If you have any services running when you initiate the upgrade to high availability, it may take up to 45 minutes for the process to complete. This is because running services have to be stopped and restarted.

To upgrade to high availability, click the **upgrade to HA** button in the deployment header.

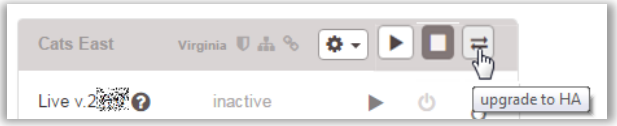

The upgrade process begins: the services are duplicated and restarted (if they were running prior to the upgrade).

To test failover, see page 22.

5.8. Delete Service Backups

You can delete your service backups in the following ways:

- On active accounts:
  - Contact AWS Elemental Support through your company's Private Space in AWS Elemental User Community (https://community.elemental.com) to request the deletion of service backups.
  - To delete users, see the help page titled **Subscription** on the AWS Elemental Cloud web interface.
  - Manage and delete content from your asset repository outside of AWS Elemental Cloud.

    Content delivered to Amazon S3 buckets might contain customer content.
    For more information about removing personally identifiable information, see
    S3 documentation here: https://docs.aws.amazon.com/AmazonS3/latest/dev/delete-or-empty-bucket.html
- For archived accounts: all account data is automatically removed within 90 days of archive. If you need data deleted sooner, contact AWS Elemental Support through your company's Private Space in AWS Elemental User Community (https://community.elemental.com).

5.9. Delete Inactive Service

You can use the REST API to delete an inactive service from a deployment.

When the service is in an inactive state, send a POST request using the `delete` command to delete it.

```
POST https://<subscription_name>/service_managers/<service_ID>/command/delete
```

For more information about using the REST API, see the REST API guide in the Help pages of your AWS Elemental Cloud subscription.