



HOW TO:

LAUNCH A JETNEXUS ALB-X WITHIN AMAZON WEB SERVICES (AWS)

Using a jetNEXUS ALB-X within Amazon Web Services (AWS), you can either;

- (i) Use jetNEXUS ALB-X to accelerate your Applications or Websites running within AWS and make substantial savings on your AWS infrastructure and bandwidth costs.
- (ii) Easily trial jetNEXUS in front of your existing, in-house infrastructure

Creating a jetNEXUS ALB-X on AWS couldn't be simpler and getting up and running should take less than 5 minutes.


Please feel free to contact us at any time via support@jetnexus.com should you require any assistance in deploying or configuring your new ALB-X. We want to ensure you have as successful an experience as possible using our technology and are more than happy to assist you in getting everything up and running, so please don't hesitate to get in touch.

Please note, you'll need to have an AWS account in order to deploy an ALB-X within that environment. If you don't already have an AWS account, setting one up is a very quick and simple process – please see <http://aws.amazon.com/ec2/> for more details

The process of setting up a jetNEXUS ALB-X within AWS is extremely simple and can be carried out by following the instructions below:

1. Log into your Amazon Web Services Management Console
2. Select the AWS Elastic Compute Cloud (EC2)
3. Select the Region you wish to set the Appliance up in
4. Select Images → AMIs
5. Search under 'All Images' for the AMI code appropriate to your Region of choice. Please see the table below for the jetNEXUS AMI Codes

AWS REGION	JetNEXUS ALB-X AMI ID
US EAST	ami-b30dd8da
US WEST (CALIFORNIA)	ami-65356b20
US OREGON	ami-c0a12cf0
EUROPE	ami-ef744b9b
SINGAPORE	ami-4481c416
TOKYO	ami-1ed6601f
SOUTH AMERICA	ami-802cf39d



6. When you've found the jetNEXUS AMI, right click on it and select 'Launch Instance'
7. Next, follow the standard AWS process for setting up a new Instance. On screen instructions from AWS will guide you through this process if you're not

already familiar with it, but we've flagged below some of the key items to note in relation to the jetNEXUS appliance specifically.

INSTANCE DETAILS

The size of AWS Instance you should chose to deploy your jetNEXUS ALB-X onto depends on the amount of traffic it will be required to handle and we would advise testing to determine the best Instance size to meet your specific needs.

However, please note that a **Micro Instance is not adequate for testing an ALB-X** as it does not have sufficient resources to run an ALB-X and is likely to be forcibly rebooted by AWS when put under even the slightest load.

Once you've decided which Instance size you wish to deploy on, you can then select whichever Availability Zone you deem to be most appropriate. As a rule of thumb, we would suggest that you select the Availability Zone in which the majority of your web servers reside but the ALB-X is able to load balance and accelerate traffic across different Availability Zones without any difficulties.

Request Instances Wizard Cancel

CHOOSE AN AMI **INSTANCE DETAILS** CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Provide the details for your instance(s). You may also decide whether you want to launch your instances as "on-demand" or "spot" instances.

Number of Instances: Instance Type:

Launch Instances

EC2 Instances let you pay for compute capacity by the hour with no long term commitments. This transforms what are commonly large fixed costs into much smaller variable costs.

Launch into:

EC2

Availability Zone:

Request Spot Instances

[Back](#)

ADVANCED INSTANCE OPTIONS

None of the Advanced Instance Options are necessary to set up the jetNEXUS ALB-X, but you may wish to configure some of the options shown below, such as Termination

Protection, for your own reasons. As a rule of thumb, we would suggest that you set the 'Shutdown Behaviour' to Stop.

Request Instances Wizard Cancel X

CHOOSE AN AMI | **INSTANCE DETAILS** | CREATE KEY PAIR | CONFIGURE FIREWALL | REVIEW

Number of Instances: 1
Availability Zone: eu-west-1a

Advanced Instance Options

Here you can choose a specific kernel or RAM disk to use with your instances. You can also choose to enable CloudWatch Detailed Monitoring or enter data that will be available from your instances once they launch.

Kernel ID: Use Default **RAM Disk ID:** Use Default

Monitoring: Enable CloudWatch detailed monitoring for this instance (additional charges will apply)

User Data:

as text

as file base64 encoded

Termination Protection: Prevention against accidental termination.

Shutdown Behavior: Stop Choose the behavior when the instance is shutdown from within the instance.

[< Back](#) Continue

CREATE KEY PAIR

You can either choose to use an Existing Key Pair or Create a New Key Pair, but do not proceed without a Key Pair. For the purposes of the example shown below, we'll use our existing key pair 'RTSKeyPair1'

Request Instances Wizard Cancel X

CHOOSE AN AMI | INSTANCE DETAILS | **CREATE KEY PAIR** | CONFIGURE FIREWALL | REVIEW

Public/private key pairs allow you to securely connect to your instance after it launches. To create a key pair, enter a name and click **Create & Download your Key Pair**. You will then be prompted to save the private key to your computer. Note, you only need to generate a key pair once - not each time you want to deploy an Amazon EC2 instance.

Choose from your existing Key Pairs

Your existing Key Pairs*: RTSKeyPair1

Create a new Key Pair

Proceed without a Key Pair

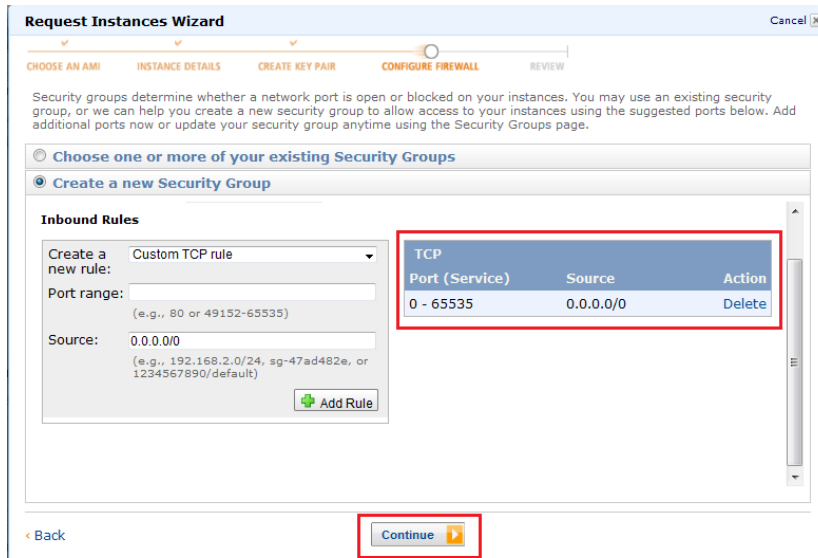
CONFIGURE FIREWALL

You can either chose to use an Existing Security Group or Create a New Security Group. The key thing is to ensure that the requisite ports are open, otherwise it won't

be possible for you to connect to the Instance via the console. So, for the purposes of this example, we'll Create a New Security Group.

For the sake of ease and simplicity, in this example we'll;

- Create a new Custom TCP Rule
- Use Port Range 0 – 65535
- Use a Source of 0.0.0.0/0
- Give a Group Name & Description of 'jetnexus group'



We can now [Launch the Instance](#)

Of course, it is likely you'll want to create more stringent security rules than this for anything other than the most basic of testing, but this is sufficient for this example. For reference, the ALB-X's web interface connects on Port 27376

An Instance will typically take a couple of minutes to launch. When it has fully launched and has a Status of 'Running' you should be able to connect to the ALB-X but it's worth noting that the AWS Console can show an Instance itself as 'Running' before all of the jetNEXUS components have been fully configured and are themselves running.

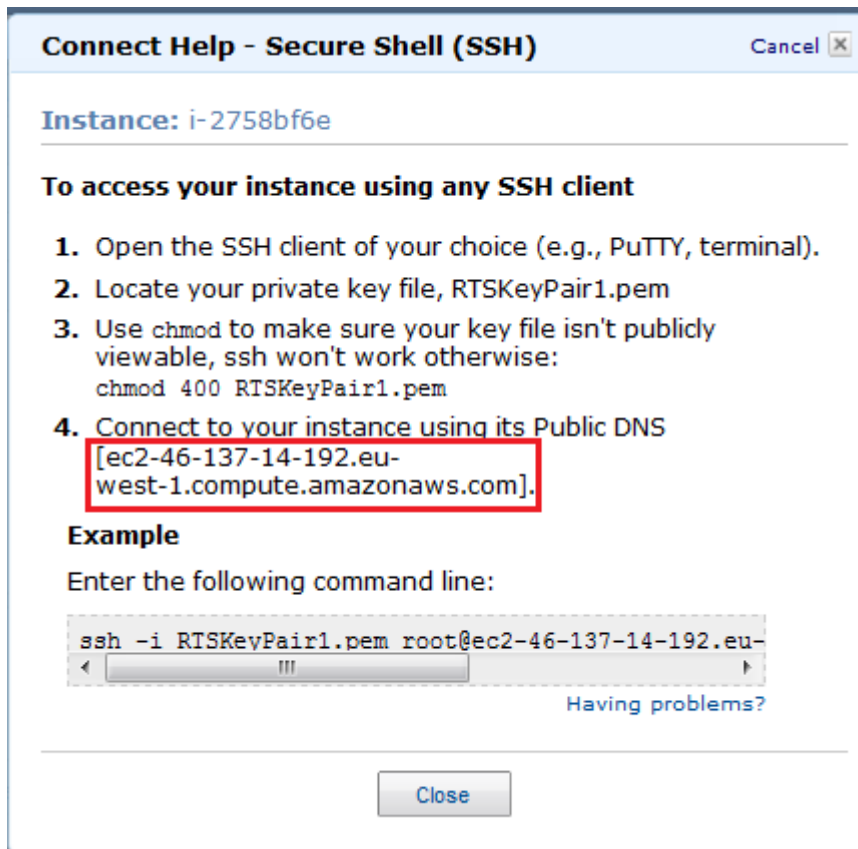
You can easily check if your new ALB-X is running by copy and pasting the Instance's DNS into your browser window – providing the appliance is running successfully, you'll be automatically directed to a test page that your appliance has been automatically configured with.

CONNECTING TO THE JETNEXUS APPLIANCE

Simply right click on the Instance and select **Connect**

This will provide details of how to Connect to your Instance using its Public DNS

Simply copy the Public DNS as per the screenshot below



In this example the Public DNS is ec2-46-137-14-192.eu-west-1.compute.amazonaws.com

Therefore, we need to browse to;

<https://ec2-46-137-14-192.eu-west-1.compute.amazonaws.com:27376>

Please note that the Instance-specific Public DNS should be prefixed by **https://** and suffixed by **:27376** as per the example above.

Alternately, you can assign an Elastic IP to the jetNEXUS ALB-X Instance and access it via that IP. To do this, simply;

- Go to Elastic IPs in your AWS Management Console
- Allocate a New Address (assuming you don't already have one spare)
- Right click on the IP address in question and then associate it with the newly created Instance
- You'll then be able to connect to the instance via the browser using the IP address – for example, <https://46.137.171.99:27376>

You'll now be able to connect to the jetNEXUS appliance directly from your web browser.

We recommend using Internet Explorer or Firefox to access your jetNEXUS ALB-X as we cannot guarantee all functionality within alternate browsers.

When you connect, your browser will throw up an 'Untrusted Connection' warning. This is because the ALB-X issues a self-signed certificate. You can therefore ignore this error and proceed by adding a Security Exception

Use the default username and password below to log into the Appliance

User: admin

Password: jetnexus

Wait for a few moments for the Appliance to load and then you'll be able to configure the Appliance itself.

For security reasons, we would suggest that you immediately change the password, which can be done under 'Configure' → 'Security'

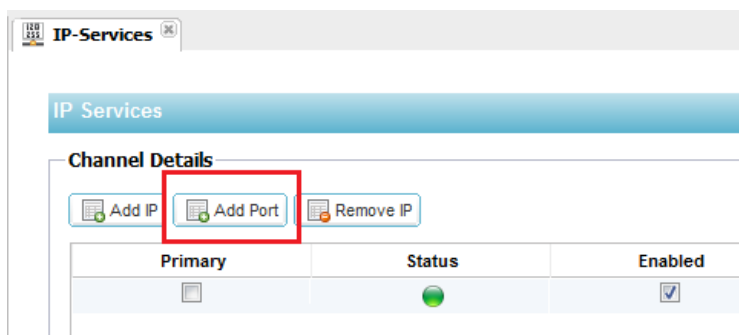
Configuration

A number of settings on the ALB-X are pre-configured for you upon start-up to make the appliance as simple and easy to use as possible.

However, there are certain networking restrictions within the AWS environment and this means that some options within the ALB-X are restricted

Please note that the Channel IP Details and the Default Gateway have been automatically configured for you as part of the installation and should not be changed or amended under any circumstances. Changing the Channel IP Details and/or Default Gateway will cause the Appliance to fail.

You can however add additional Ports to the Channel Details as shown in the screenshot below;



You should also not attempt to modify the Adapter settings (which can be found under Set-Up → Appliance) or the Hardware settings available under Set-Up → Hardware. This is because of networking restrictions within the AWS environment and altering the pre-configured settings will cause the appliance to fail.

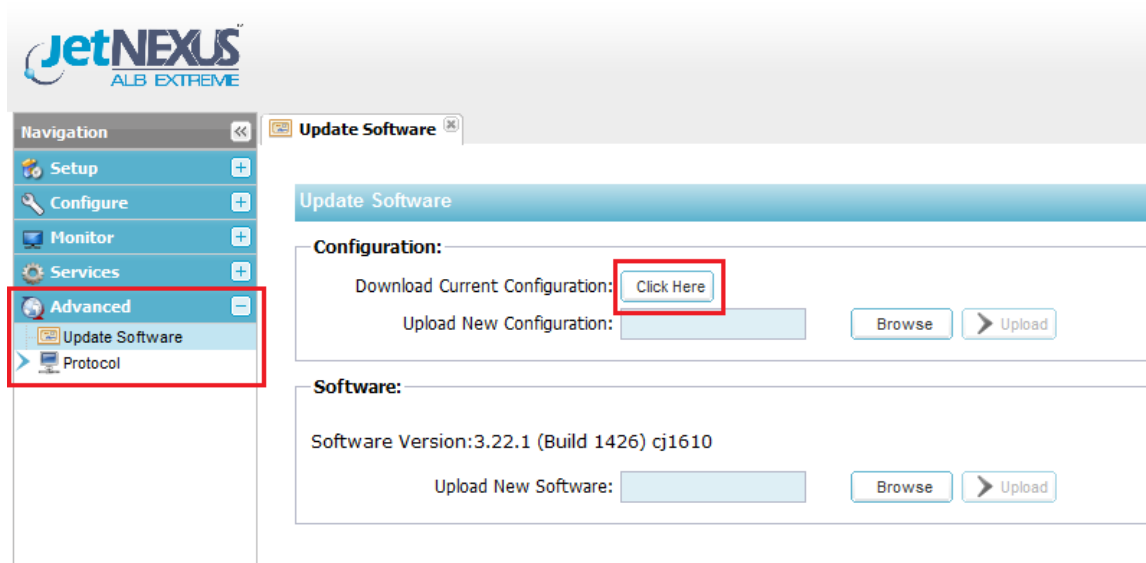
Similarly, these restrictions mean that the Bonding options are not relevant in the AWS environment and should therefore not be configured.

Uploading a Configuration File

If you wish to upload an existing configuration file from one ALB to another – for example, to replicate settings across two Highly Available jetNEXUS ALBs – this can be done in the normal fashion, albeit with one minor amendment.

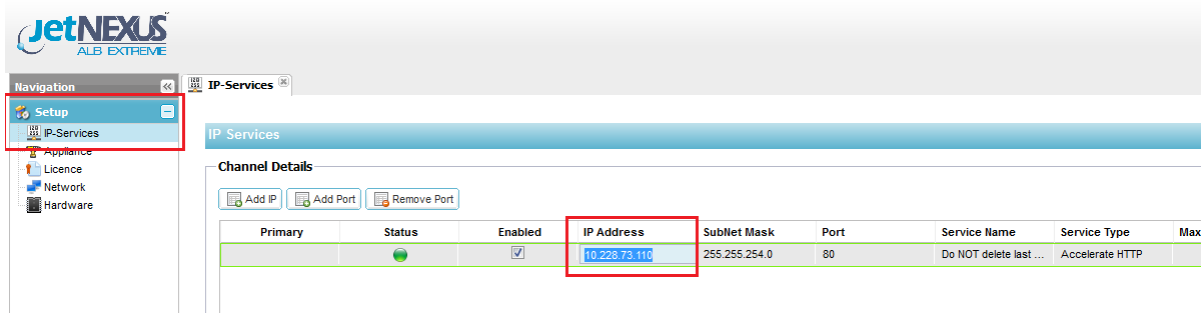
The process you should follow is to;

1. Download the Configuration File from the first ALB under Advanced → Update Software → Configuration → Download Current Configuration



2. Upload this Config file to the new ALB under Advanced → Update Software → Configuration → Upload New Configuration
3. Once the configuration has been uploaded, you'll need to change the Channel IP address to the correct one, as the ALB will now have inherited the Channel IP address from the Configuration File which will not be correct for this new Instance. You should instead replace the Channel IP address with the Internal IP address of the Amazon Instance in question.

This Internal IP address for the underlying Instance can be easily found using the AWS Web Console.



DEPLOYING ALB-X ON A HIGHLY AVAILABLE BASIS IN THE AMAZON CLOUD

Due to networking restrictions within the AWS Cloud it is not possible to directly configure multiple ALB-Xs in High Availability mode.

For customers wanting to deploy ALB-Xs in an HA pair, there are two separate methods available to achieve High Availability;

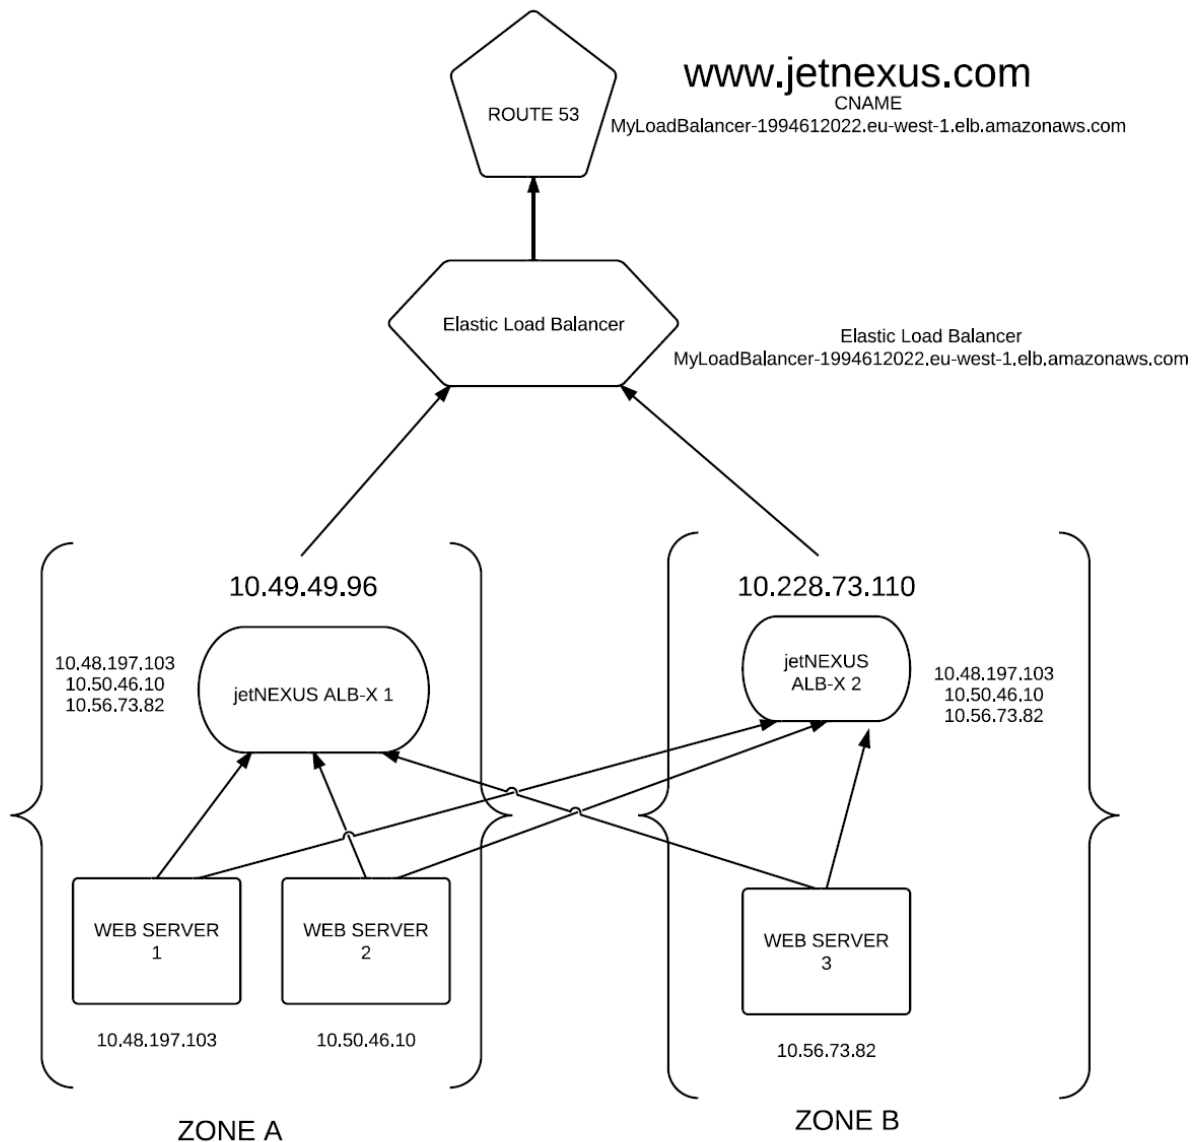
Using Elastic Load Balancing (ELB)

Two or more jetNEXUS ALBs can be configured behind an AWS Elastic Load Balancing Group in an Active-Active configuration, in the same manner that you might use ELB to balance two web servers in AWS. Using ELB with jetNEXUS ALB ensures that you'll still get all the benefits of caching, compression, connection pooling, flightPATH Traffic Management and many other rich features that ELB alone can't provide.

The diagram shows a simple example configuration of two HA ALBs using Elastic Load Balancing. As you'll note, a configuration such as this can sustain the loss of any individual Instances or a whole AWS Zone and still be able to serve accelerated content uninterrupted.

The basic process for doing this would be to;

- * Create and configure ALB 1
- * Create and configure ALB 2 with identical settings
- * Create an Elastic Load Balancer
- * Add the two ALB instances to the newly created Elastic Load Balancing Group
- * Use Elastic Load Balancing to provide the principal interface for accessing your website or application, such as via the creation of a CNAME record



Using DNS

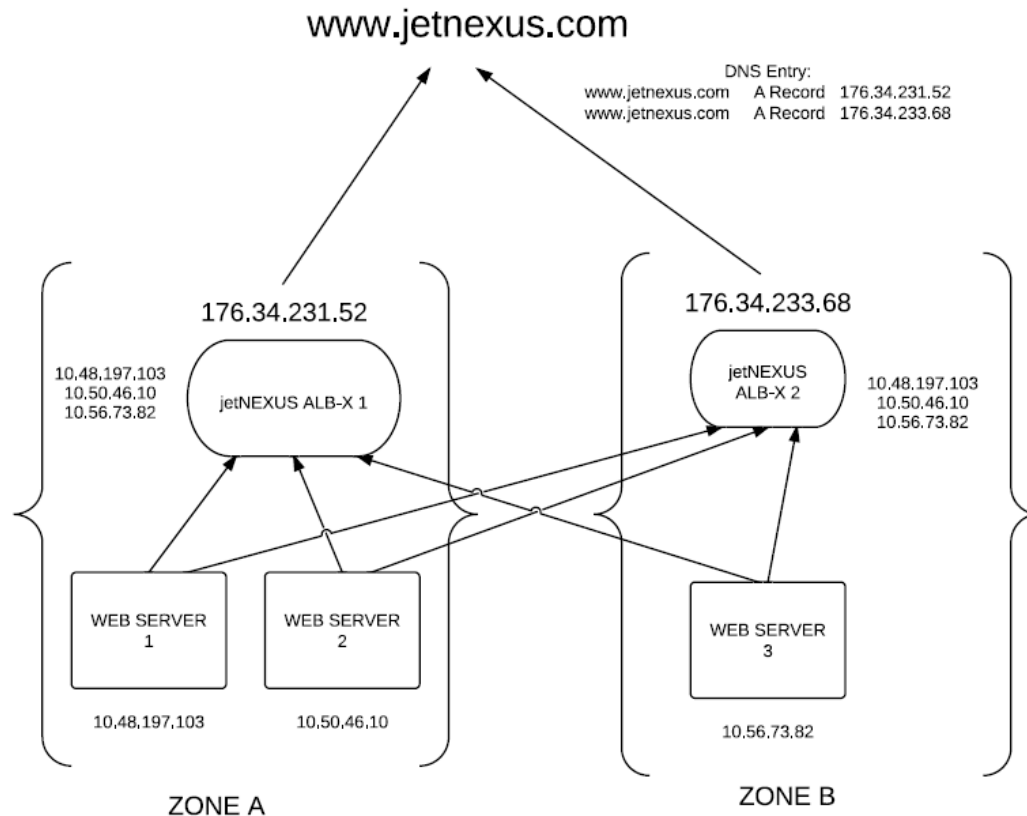
Alternately, you can use DNS-based load balancing to help achieve High Availability for ALB-Xs in an Active-Active pair. This will also ensure continuity of service even in the event that the underlying Instance running one of the ALB-Xs should fail.

The diagram below shows a simple example configuration of two HA ALBs using DNS Load Balancing. As you'll note, a configuration such as this can sustain the loss of any individual Instances or a whole AWS Zone and still be able to serve accelerated content uninterrupted.

The basic process for doing this would be to;

- * Create and configure ALB 1

- * Create and configure ALB 2 with identical settings
- * Assign Public (aka Elastic IPs) to each of the ALBs
- * Add two A Record entries to the DNS for the domain in question, one for each of the Public IP addresses assigned to the ALB



Should you require any assistance in configuring two or more ALBs within the AWS environment, please don't hesitate to contact us via support@jetnexus.com

DEPLOYING ALB-X IN CONJUNCTION WITH AWS AUTO-SCALING

The current version of jetNEXUS ALB-X is not directly integrated with AWS Auto-Scaling but can be configured to work with your existing auto-scaling policies if required. Please contact our engineering team for more details on how to achieve this, available via support@jetnexus.com

For further details on configuring the ALB-X itself, please consult the User Manual for further details, available from within the ALB-X GUI under Help → Reference.

Please note that any guidance given within this deployment guide should take precedence over any conflicting information in the ALB-X manual available via the GUI.

If you need any additional assistance, please don't hesitate to contact us via support@jetnexus.com

Author(s):	RTS
Contributors:	OG
Document Ref:	ALB/AWS Deployment Guide
Version:	1.5
Status:	Live
Created On:	17 August 2011
Modified On:	23 December 2011