

AWS Region

a) Region

When configuring the Amazon environment for a client, one of the first things to do is to identify which Amazon region will be used to run services required by the client.

A primary characteristic of an AWS region is that different services required by a particular client - e.g. EC2, RDS, S3, CloudFront - should all be provisioned under the same region. In other words, it's important to choose the AWS region prior to installing the services.

There are a few considerations when selecting the region:

- The region should be geographically close to the client's customers and primary places of operation.
- Due to privacy and other regulations, the client may require that hosting infrastructure be located in a specific physical region.
- Some AWS infrastructure is shared between multiple clients, as the services present a fixed cost that can be shared across multiple clients. This characteristic requires that Kensium limit the number of regions that it can support, at least initially.

Given the above, the AWS regions that Kensium currently supports are:

- us-east-1 - deployments to customers with operations mainly in the eastern US or eastern Canada.
- us-west-2 - deployments to customers with operations in the western US, western Canada, or Mexico.

Kensium may support additional regions in the future.

b) Availability Zone (AZ)

Enter the region's availability zone selected for the client.

When selecting other AWS services such as EC2 or RDS, ensure that the availability zones for each service match the zone entered here.



There are two reasons for using the same availability zone:

- Performance
- Reduce a client's exposure to a faults in a single availability zone (i.e. a client depending on multiple zones increases the chance that a fault in a single zone will disrupt operations).