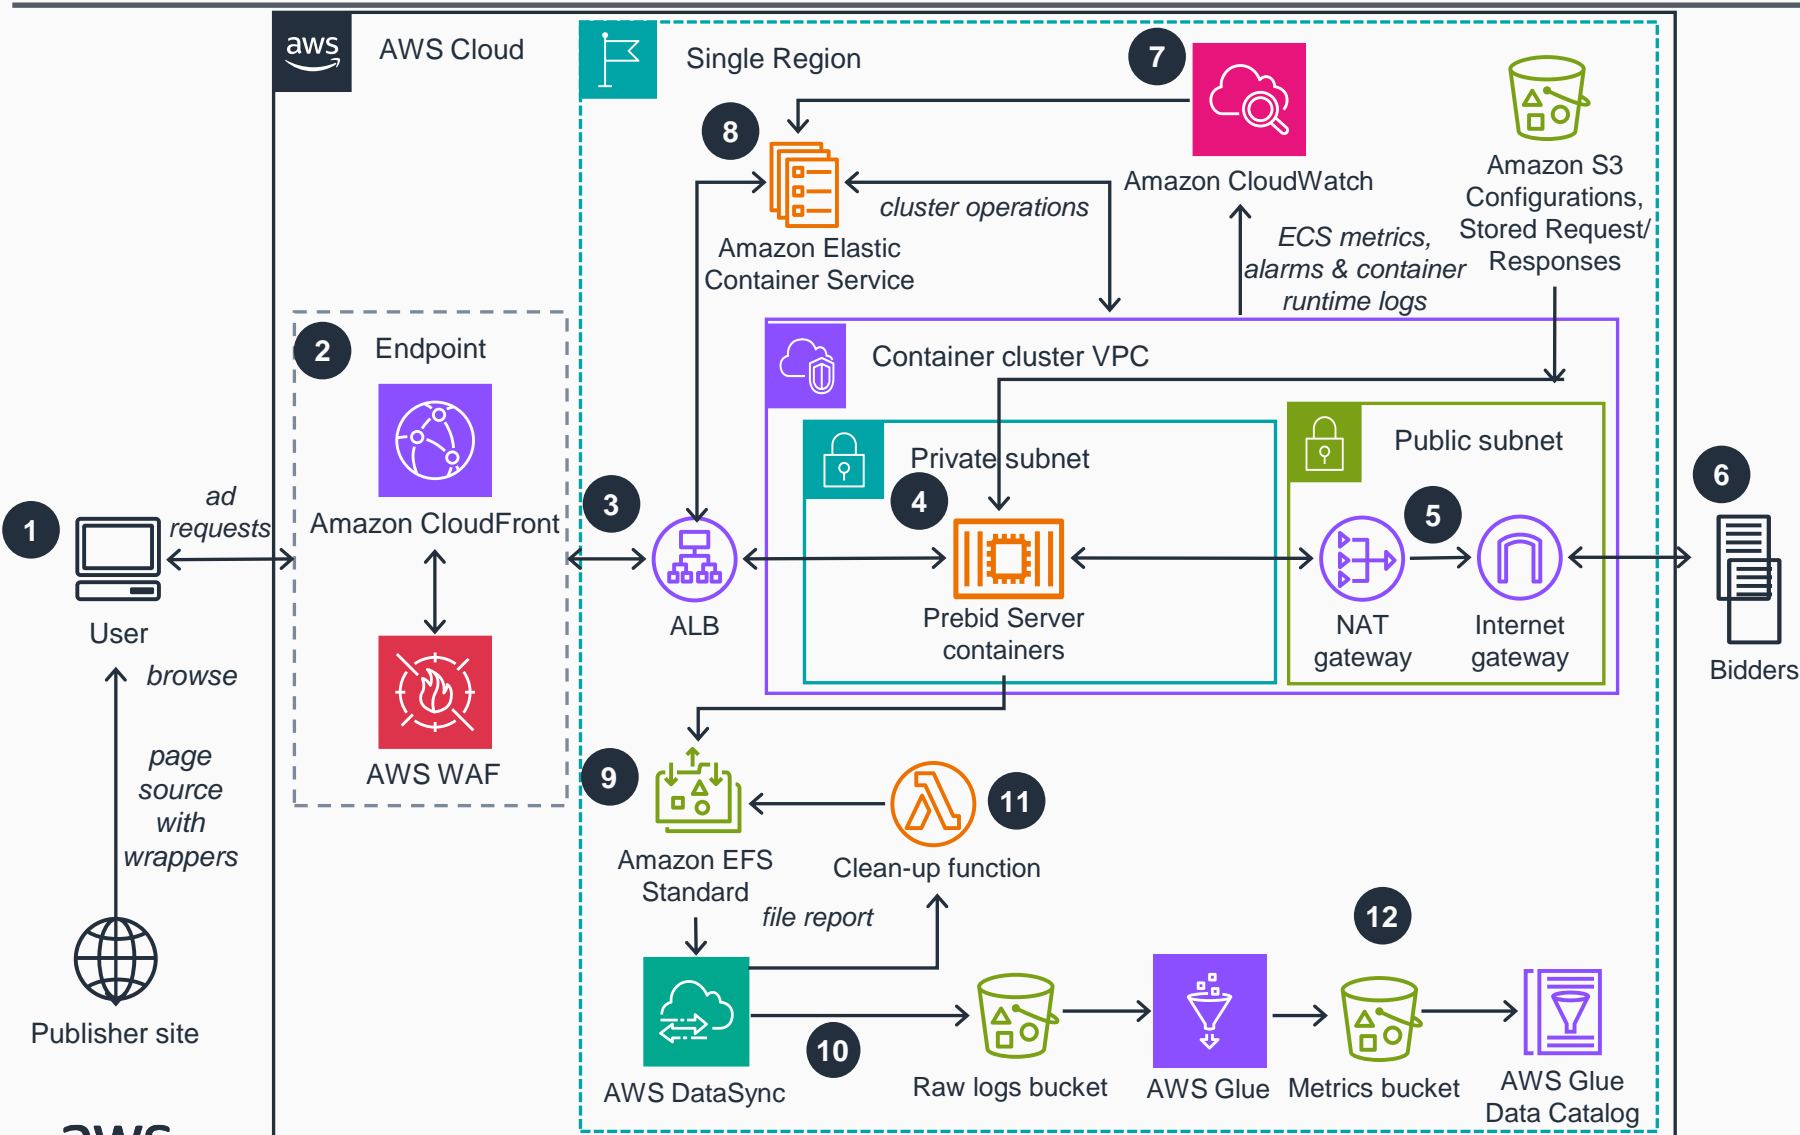


# Guidance for Deploying a Prebid Server on AWS

This architecture diagram illustrates how to effectively support Prebid server deployment on AWS. It shows the key components and their interactions, providing an overview of the architecture's structure and functionality.



- 1 A user browses to a page on a website that hosts ads. The publisher site returns the page source to the browser with resources, and one or more script modules (also called wrappers) that enable real-time bid requests and responses for ads of given dimensions, types, topics, and other criteria.
- 2 The bid requests are received from the browser at the **Amazon CloudFront** endpoint integrated with **AWS Web Application Firewall (AWS WAF)** for entry into the solution. This step helps validate legitimate traffic from malicious requests, such as penetration or denial-of-service attempts. Traffic can be received here as HTTP or HTTPS.
- 3 The request is forwarded to **Application Load Balancer (ALB)**. ALB determines which container running Prebid Server in the cluster is at a utilization level that can accept more requests. ALB has a network interface on the public internet and one in each private subnet where containers are hosted within **Amazon Virtual Private Cloud (Amazon VPC)**.
- 4 The request arrives at an **Amazon Elastic Container Service (Amazon ECS)** container, is parsed and validated, and requests to different bidding services are sent concurrently to the internet through the default internet gateway.
- 5 The **NAT gateway** and **Internet Gateway** allow containers to initiate outbound requests to the internet and receive responses. These resources are primarily used for Prebid Server containers to request and gather bids for ad auctions.
- 6 Bidders receive one or more bid requests over the internet from a Prebid Server container. Bidders respond with zero or more bids for the various requests. The response, including the body of the winning creative(s), is sent back to the browser.



This architecture diagram illustrates how to effectively support Prebid server deployment on AWS. It shows the key components and their interactions, providing an overview of the architecture's structure and functionality.

