Guidance for Deploying a Serverless IPFS Cluster on AWS

Easy to deploy and operate InterPlanetary File System (IPFS) Cluster for Web2 and Web3 applications

This diagram shows how an IPFS Cluster can run and operate in AWS Cloud. Adopt this Guidance to fit the storage needs of most use cases. Easy to operate, this cluster uses AWS managed services and has a very low operational cost.

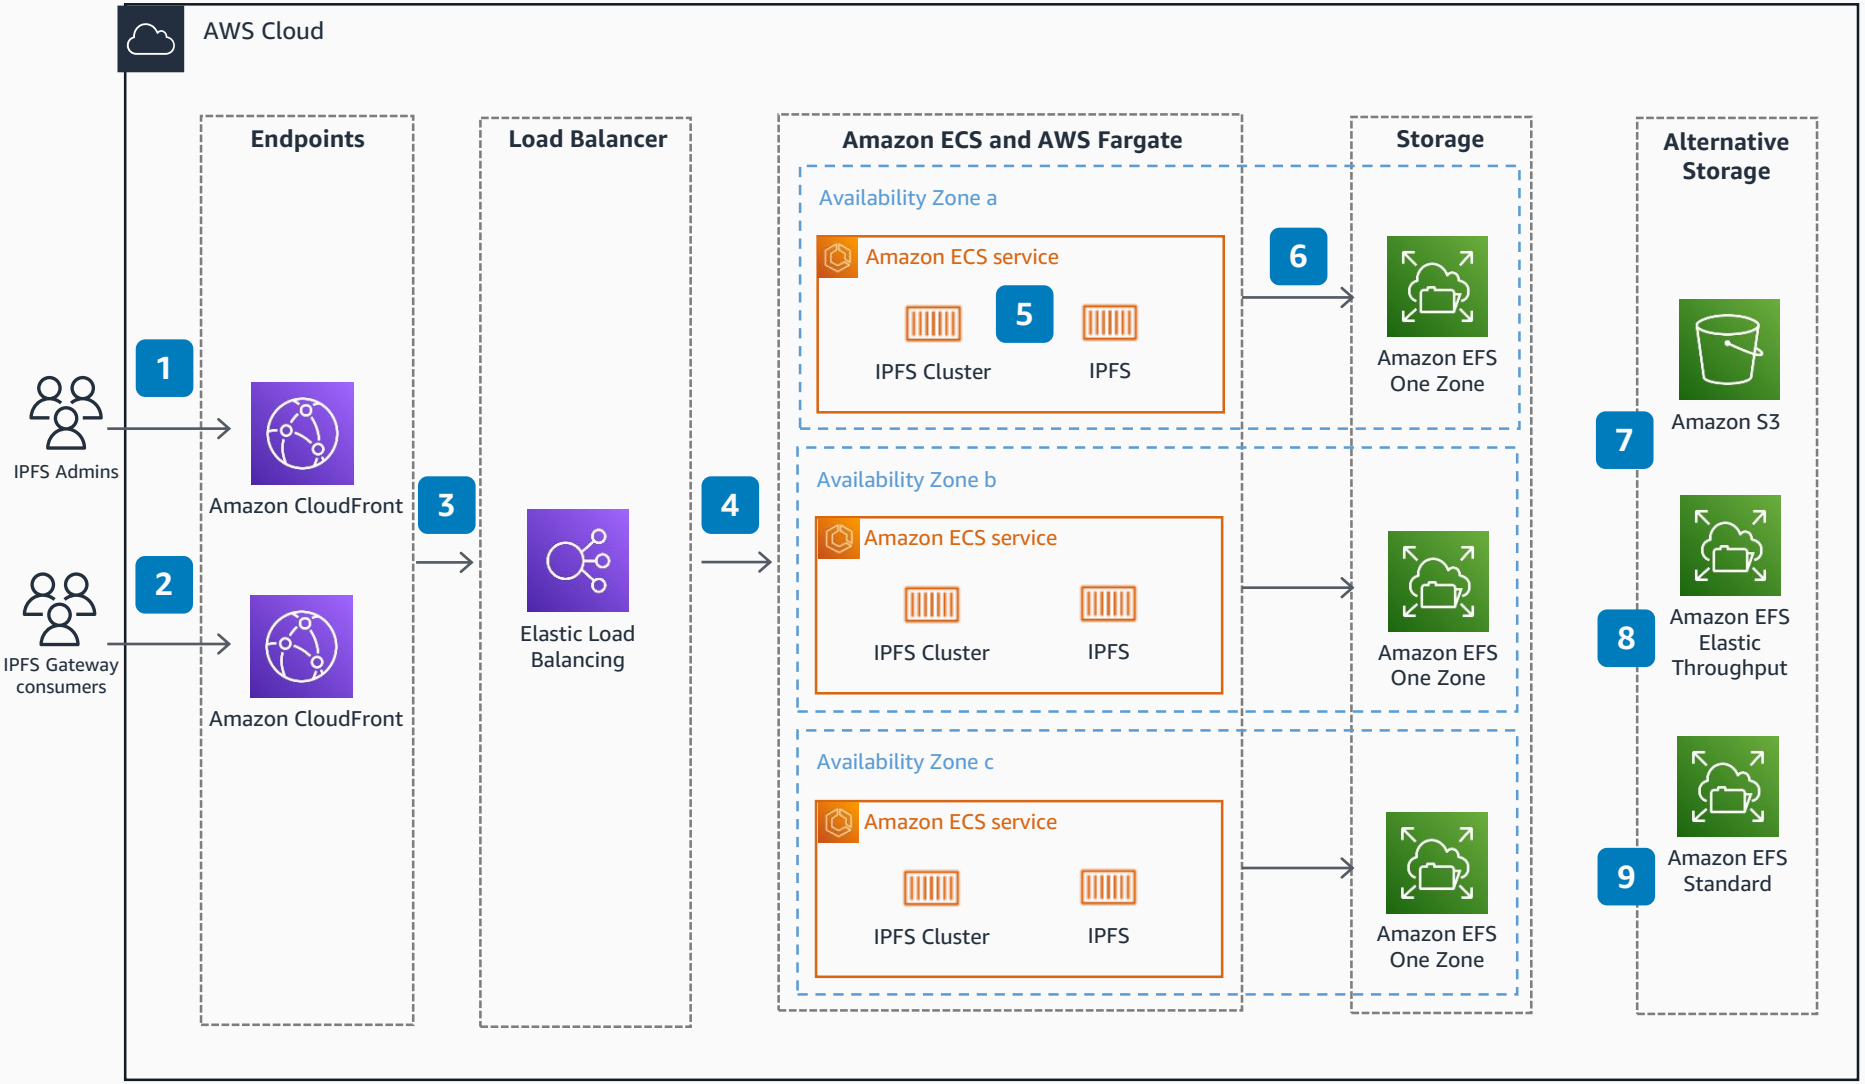

1. InterPlanetary File System (IPFS) Admins operate the IPFS Cluster over HTTPS through the REST API served by **Amazon CloudFront** on port 443.
2. Consumers access IPFS content over HTTPS through IPFS Gateway served by **CloudFront** on port 443.
3. Traffic coming from both **CloudFront** distributions goes through an Application Load Balancer (ALB).
4. The ALB distributes traffic from both **CloudFront** distributions to three **Amazon Elastic Container Service** (Amazon ECS) containers that manage **AWS Fargate** tasks.
5. **Amazon ECS** runs one **Fargate** task per Availability Zone (AZ), each running two containers. One container runs the IPFS daemon. The other runs the IPFS Cluster sidecar which provides the REST API.
6. The **Fargate** tasks use **Amazon Elastic File System** (Amazon EFS) as permanent storage. **Amazon EFS One Zone** replicates content in one AZ. It uses the 'bursting' mode by default.
7. **Amazon Simple Storage Service** (Amazon S3) can be used as a storage layer for the IPFS blockstore. **Amazon S3** can scale and is cost efficient to store lots of files. The number of GET requests to **Amazon S3** will increase your **Amazon S3** costs.
8. **Amazon EFS** Elastic Throughput is another mode for **Amazon EFS**. The performances are much better than the 'bursting' mode but it is also much more costly. Provisioned Throughput mode is better suited to set a baseline.
9. **Amazon EFS Standard** storage class replicates data over multiple AZs. Only one drive is needed but it costs more.

Reviewed for technical accuracy February 20, 2023

AWS Reference Architecture