



Adobe® Experience Manager Managed Services

Overview of cloud deployment models

Outsource management of your Adobe Experience Manager solution to the same folks that wrote the application. Take advantage of Adobe's experience with designing, operating, and maintaining our solutions so that you can focus on launching your websites, mobile apps, landing pages, and digital marketing campaigns.



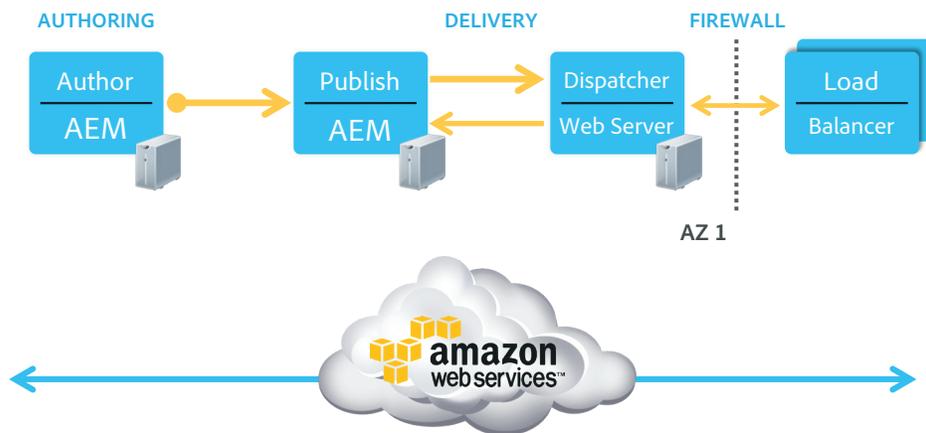
Adobe Experience Manager Managed Services can be deployed in various ways to suit your size and needs:

- Non-high availability (99.5%) deployment topology
- High availability (99.9%) deployment topology
- Geo and high availability (99.99%) deployment topology

Adobe Experience Manager Sites Deployment models

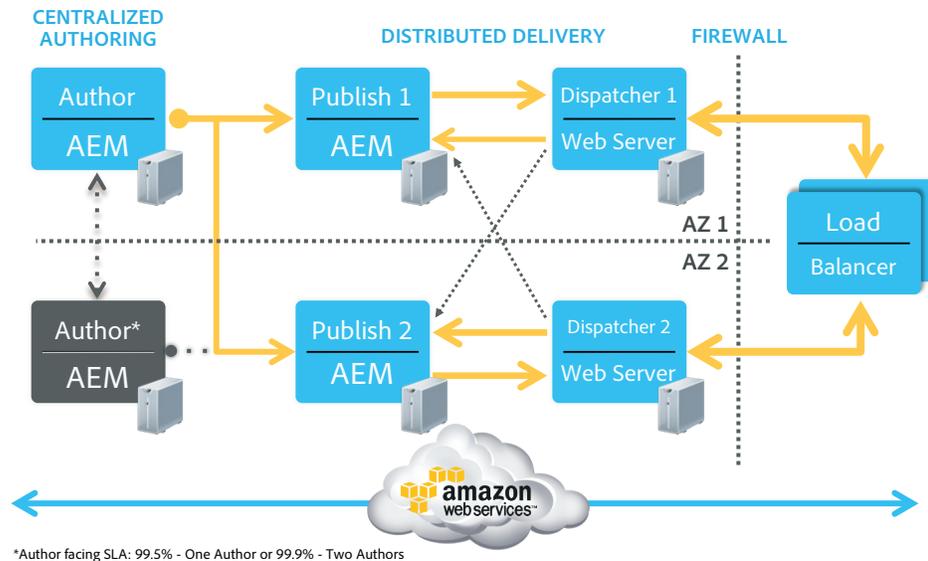
Adobe Experience Manager is hosted on the Amazon Elastic Compute Cloud (EC2) in a safe, dedicated environment while managed and fully supported by Adobe. Adobe offers several deployment models based on your needs for high availability. The options are listed in order of increasing availability.

Non-high availability (99.5%) deployment topology (single availability zone, single region)



This single-leg topology consisting of one availability zone (AZ) is optimal for smaller Experience Manager deployments because it provides optimized cost performance for smaller brands and offers the potential for future expansion. All Experience Manager Managed Services systems can grow both vertically—you can add more CPUs to any of the three servers—and horizontally by adding more publish or dispatcher servers.

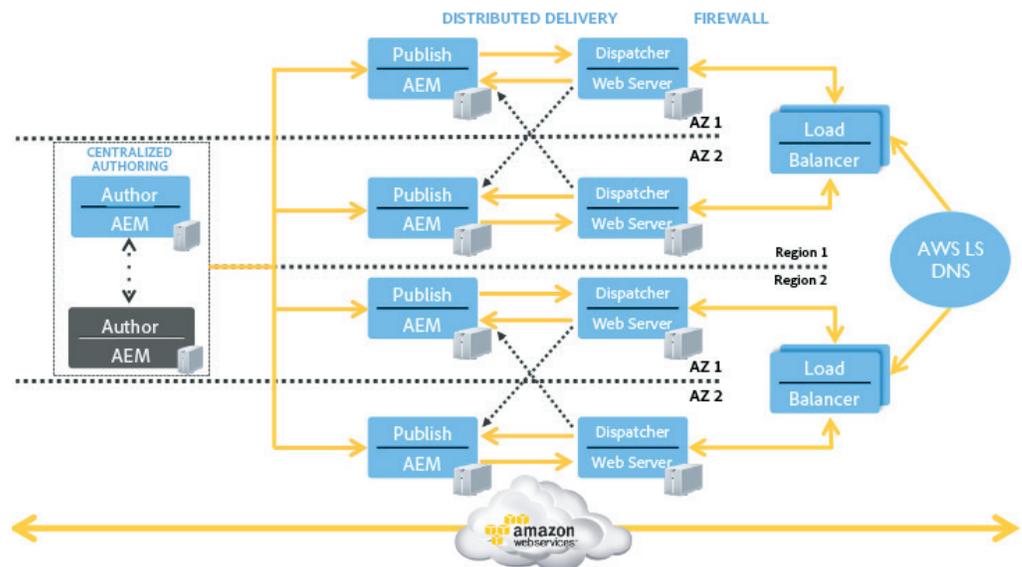
High availability (99.9%) deployment topology (dual availability zone, single region)



Amazon EC2 provides the ability to place instances in multiple geographical locations that are composed of regions and availability zones. Regions are geographically dispersed and consist of one or more availability zones. An availability zone is a distinct location that is engineered to be insulated from failures in other availability zones and to provide inexpensive, low-latency network connectivity to other availability zones in the same region. By launching instances in separate availability zones, you can protect your applications from failure of a single location. This topology distributes the application among two availability zones located inside a single region to provide for high availability system.

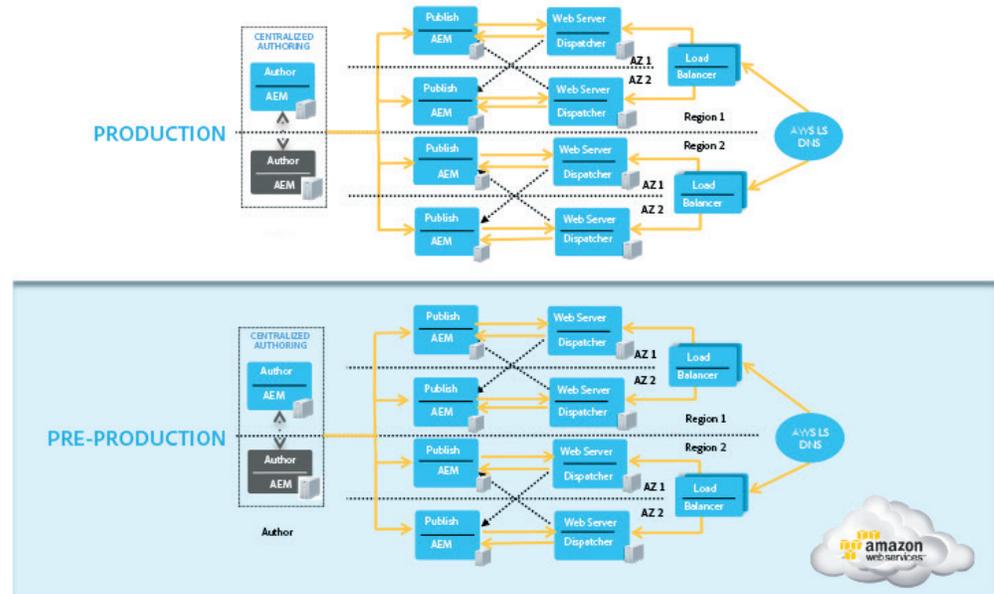
In the event of the total destruction of a data center and the consequent failure of a whole Availability Zone, the load balancer in this high availability topology will direct all incoming traffic to the operational Availability Zone. This way no service outage is visible to the end customer while the failed Availability Zone is restored inside an operational Availability Zone within the same region.

Geo and high availability (99.95%) deployment topology (dual availability zone, multiple regions)



This dual region, dual availability zone topology utilizes a high-availability design in each of the regions. Latency sensitive DNS directs customers to the lowest latency location as well as to an operational location in case part of the system fails. This deployment model is ideal for customers with a global audience looking to provide an uninterrupted brand experience.

Geo and high availability (99.99%) deployment topology (dual availability zone, multiple regions)

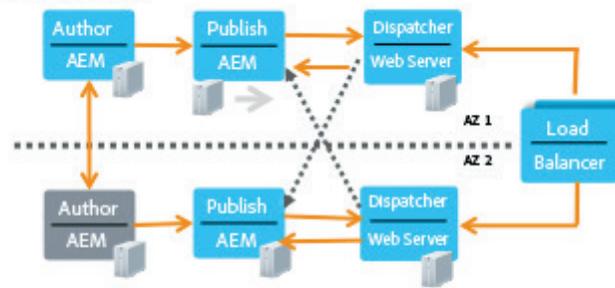


If you need additional system-failure protection, you can deploy your system in up to six regions to have the instances hosted in separate geographic areas or countries. This “four-leg” architecture—two legs in each region—reduces latency and upgrade downtime. Latency-sensitive Domain Name System (DNS) servers direct clients to the lowest latency location to get the content in the fastest way possible. If the system in one region fails, DNS redirects the traffic to the working region.

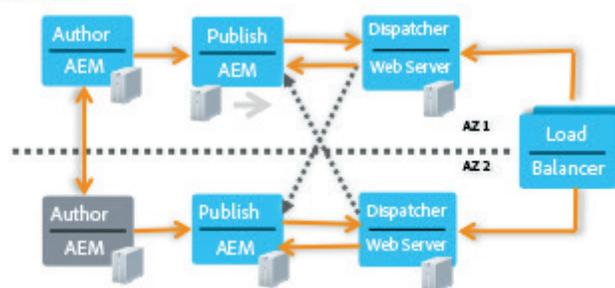
Ideal deployment topology (99.9%, dual availability, single region)

This dual availability zone, dual region topology comes with additional content and upgrade mechanisms to provide a staging system with customer access and testing capability. In this case, preproduction and the production system are under Adobe control and consist of the 99.9% dual availability topology.

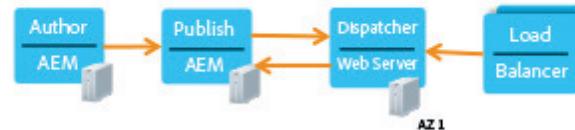
PRODUCTION



STAGING



DEV/QA



For example, the demand and customer marketing team at a well-established retailer wants to drive a deeper connection with existing and new users through its online channel. As a result, they have planned a three-year demand and social engagement campaign to build awareness of its products and drive more purchases at retail stores in a specific geographical region. Yet, it's unclear how much traffic they will generate over the next three years. Unfortunately, the company's internal IT does not have the resources or time to manage this project, and they do not want to add staff for a finite project.

Understanding that premium customer experience is related to failover and high availability, their ideal deployment should span across North America but inside two availability zones in order to reduce latency to the customers in different parts of the region.

An ideal high-availability deployment would consist of:

- Production Environment
 - AEM Sites - 99.9% SLA
 - 1 author, 2 publish, 2 dispatcher
 - AEM Additional Instance
 - 1 author
 - AEM Add-Ons
 - Multi-Site Manager, Mobile, Social
- Stage Environment
 - 2 author, 2 publish, 2 dispatcher
- Dev Environment
 - 1 author, 1 publish, 1 dispatcher

This deployment is ideal because staging is a replication of the production. Customers can vary the integration based on the instance sizes.

Special considerations

To provide the most stable environment for an Experience Manager Managed Services deployment, Adobe requires that customers who purchase the Social Communities add-on also purchase either Cloud Extension (see below) or add an additional instance to their Production system.

What's in the box

Every cloud deployment of Experience Manager comes with the following features:

- Single tenant clusters
- Engineering-tested installations
- Support for customization and additional executables
- Fast start engineering support
- 24/7 monitoring of the entire stack
- Upgrades and updates (entire stack) handled by Adobe
- Uptime guarantees of 99.5% to 99.99%**
- Specialized Platinum Support for Managed Services
- Availability of staging and development servers (temporary or permanent)*
- ISO27001, HIPAA, and FEDRAMP security

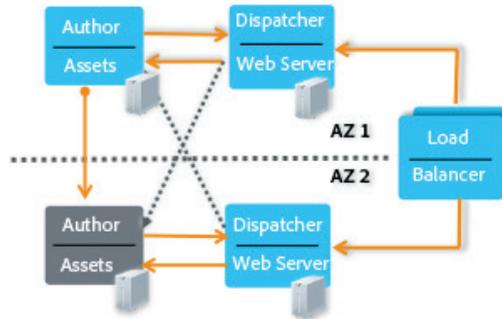
Summary

Service-level agreement	Availability zones	Regions	Advantages
99.5%	1	1	<ul style="list-style-type: none">• Ideal for smaller Experience Manager deployments• Option of scaling the system vertically and horizontally inside the AZ
99.9%	Multiple	1	<ul style="list-style-type: none">• Deployed in geographically dispersed AZs to protect from application failure• Option of scaling the system within the region as well as AZs
99.95%	Multiple	Multiple	<ul style="list-style-type: none">• Deployed in multiple AWS Regions to maximize failover• Minimized latency for globally distributed audiences
99.99%	Multiple	Multiple	<ul style="list-style-type: none">• Ideal for minimizing latency and upgrade downtime• Can be deployed in up to six regions to provide maximum uptime

Adobe Experience Manager Assets Deployment Models

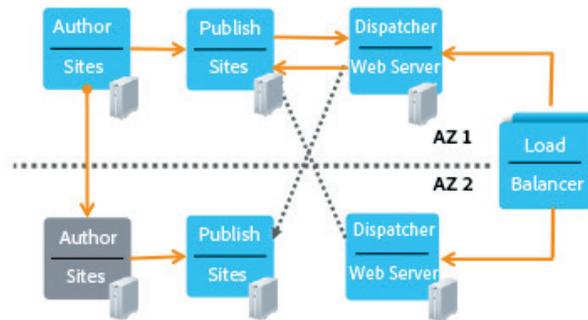
Customers purchasing Adobe Experience Manager Assets can deploy their Digital Asset Management application under the 99.9% SLA in two variations: with or without Media Portal.

High Availability (99.9%) Deployment Topology for AEM Assets without Media Portal (Dual Availability Zone, Single Region)



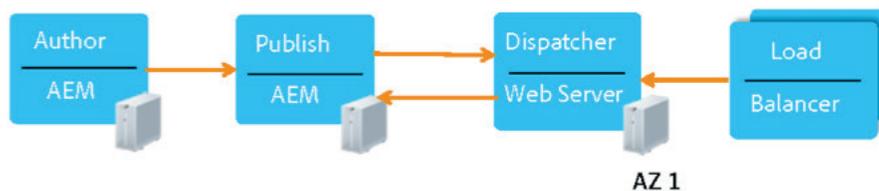
Customers not requiring Media Portal with their Adobe Experience Manager Assets application can deploy their system across two AWS Availability Zones, each containing one Author and one Dispatcher instance in order to maximize availability.

High Availability (99.9%) Deployment Topology for AEM Assets with Media Portal (Dual Availability Zone, Single Region)



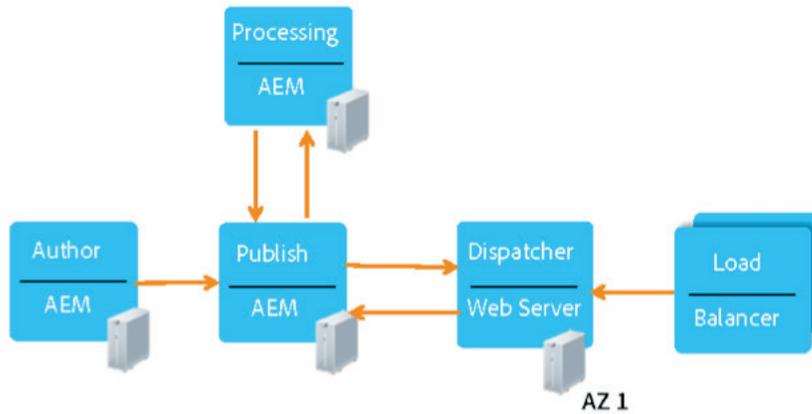
When customers expect a high level of collaboration with external parties (e.g. Creative agencies), it is recommended they purchase the Media Portal which will streamline collaboration across different organizations. Experience Manager Assets with Media Portal are deployed across two AWS Availability Zones with each containing an Author, a Publish and a Dispatcher instance.

Adobe Experience Manager Apps Deployment (Non High Availability, 99.5% SLA)



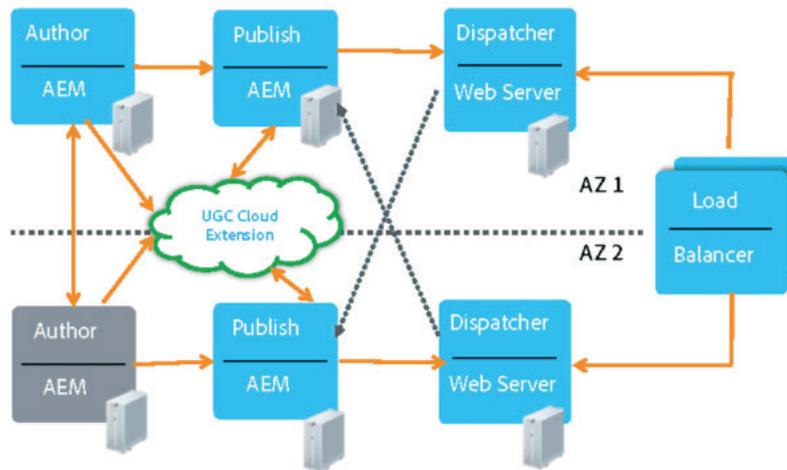
Adobe Experience Manager Apps are deployed standalone inside a single AWS Availability Zone with a guarantee of at least 99.5% uptime. In the case of Adobe Experience Manager Apps, Staging and Dev environments will mirror the Production environment.

Adobe Experience Manager Forms Deployment (Non High Availability 99.5% SLA)



Adobe Experience Manager Forms are deployed standalone inside a single AWS Availability Zone with a guarantee of at least 99.5% uptime. In addition to the standard 99.5% deployment topology, Adobe Experience Manager Forms will include an additional instance for forms processing. Staging and Dev environments will mirror the Production environment.

Adobe Experience Manager Social Communities: Cloud Extension Deployment



In the case of a customer requiring Adobe Experience Manager Social Communities expecting high volumes of User Generated Content (UGC), the recommended option is to store UGC inside the Adobe Marketing Cloud storage. This option, called "Cloud Extension", has been introduced to both improve customer experience when they are predicting high levels of the UGC, as well as for performance tuning and cost reduction for such deployments. By hosting UGC inside the Adobe Marketing Cloud, customer can also future proof their investment through easy out-of-the-box expansion into other Adobe services such as Adobe Social.

For more information
www.adobe.com/solutions/web-experience-management.html

