

A graphic featuring a stylized world map with glowing blue lines representing global connectivity. Two server racks are positioned on the map, one in North America and one in Europe. The text "Powering innovation, globally" is overlaid in a large, bold, sans-serif font, with "Powering" in white and "innovation, globally" in a gradient of blue and purple.

Powering innovation, globally

Contact Us

Explore AWS



Amazon.com
AWS Global Infrastructure

AWS Global Infrastructure: The Foundation of Cloud Computing

AWS's global infrastructure forms the backbone of their cloud services, designed with unparalleled flexibility, reliability, scalability, and security to support millions of customers across industries worldwide. This infrastructure enables businesses to deploy applications anywhere with consistent performance and redundancy.

AWS Regions: Strategic Geographic Distribution

What Are AWS Regions?

AWS Regions are distinct geographic areas containing multiple Availability Zones. Each Region is completely isolated from other Regions to ensure the highest fault tolerance and stability.

With over 30 Regions globally (and more planned), AWS provides customers the flexibility to deploy applications close to their users while meeting compliance and regulatory requirements for data sovereignty.



Regions are physically isolated from one another to mitigate risks of region-wide failures, ensuring business continuity even during major outages.

Availability Zones: The Building Blocks of Reliability



Physical Infrastructure

Each Availability Zone consists of one or more discrete data centers with redundant power, networking, and connectivity, housed in separate facilities.



Independent Resources

AZs operate with independent power sources, cooling systems, and networking infrastructure to isolate failures and prevent cascading outages.

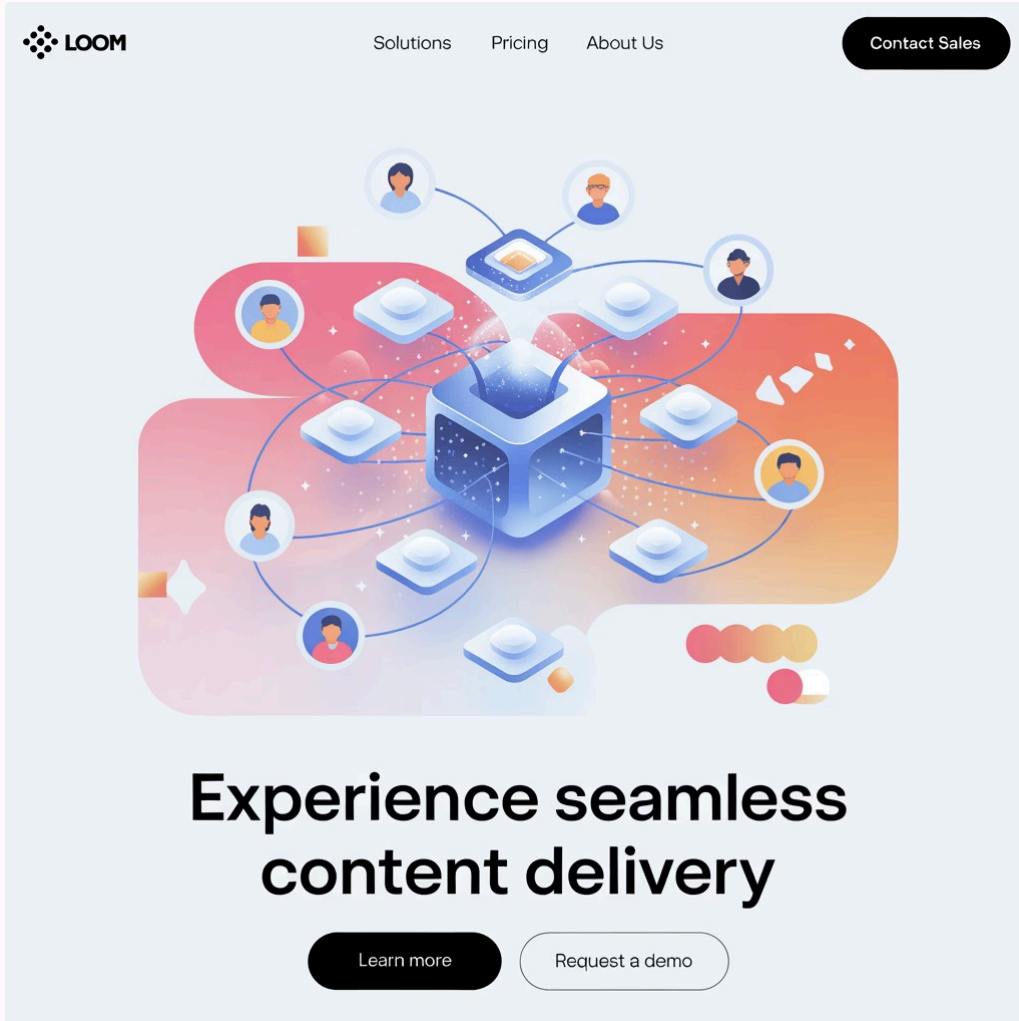


High Availability Design

Connected by low-latency links, AZs enable applications to be deployed across multiple locations for fault tolerance and high availability.

AWS recommends deploying critical applications across multiple AZs to achieve 99.99% availability and protect against data center-level failures.

Edge Locations: Bringing Content Closer to Users



The banner features the Loom logo in the top left corner, followed by navigation links for 'Solutions', 'Pricing', and 'About Us', and a 'Contact Sales' button. The central graphic is a 3D illustration of a blue cube with a glowing blue sphere on top, surrounded by several smaller blue cubes and connected by lines to circular icons of people. The background is a gradient of pink and orange with sparkling effects. At the bottom, the text 'Experience seamless content delivery' is displayed in a large, bold, black font, with 'Learn more' and 'Request a demo' buttons below it.

Experience seamless content delivery

Global Content Delivery Network

Edge Locations form the foundation of Amazon CloudFront, AWS's content delivery network, with over 400 points of presence worldwide.

- Cache static and dynamic content close to end users
- Reduce latency for global applications
- Absorb distributed denial-of-service (DDoS) attacks
- Accelerate both static assets and dynamic API responses

By serving content from locations closest to users, Edge Locations significantly improve load times and user experience.

Points of Presence: The Extended Network



Edge Locations

Serve as endpoints for CloudFront to deliver cached content directly to users with minimal latency. They handle requests for dynamic, static, and streaming content.



Regional Edge Caches

Larger caches positioned between origin servers and edge locations. They store content that's not accessed frequently enough to remain at edge locations.



User Experience

This two-tiered cache system reduces origin load by up to 90% and decreases latency by serving content from locations closer to users, improving overall application performance.

Together, these Points of Presence ensure that content delivery is optimized across AWS's global network, regardless of user location or traffic patterns.

Local Zones and Wavelength Zones: Specialized Infrastructure

AWS Local Zones

Infrastructure deployments that place compute, storage, and database services closer to large population and industry centers not served by an AWS Region.

- Ideal for latency-sensitive applications like media production
- Extends existing VPC to the Local Zone
- Currently available in major metropolitan areas

AWS Wavelength Zones

Infrastructure deployments embedded within telecommunications providers' 5G networks.

- Enables ultra-low latency applications (under 10ms)
- Perfect for IoT, AR/VR, and autonomous vehicles
- Eliminates network hops between devices and applications



**Powering tomorrow's
connections**




AWS Outposts: The Hybrid Cloud Solution



Fully Managed Infrastructure

AWS-designed hardware racks installed in your data center, fully managed by AWS, providing the same hardware used in AWS cloud regions.



Consistent Hybrid Experience

Uses the same APIs, tools, and management console as the AWS cloud, enabling a seamless hybrid experience with consistent operations.



Key Use Cases

Ideal for applications requiring low latency to on-premises systems, local data processing, or data residency requirements that prevent cloud migration.

Benefits of AWS Global Infrastructure

1 Unmatched Reliability and Resilience

The distributed nature of AWS infrastructure enables redundancy at multiple levels, from multiple power sources within a data center to geographic distribution across regions, ensuring business continuity even during major disruptions.

2 Compliance and Data Sovereignty

The ability to choose specific regions allows businesses to meet regulatory requirements regarding where data is stored and processed, addressing industry-specific compliance needs and data privacy laws like GDPR.

3 Global Reach with Local Performance

By deploying across multiple regions and leveraging edge locations, businesses can serve customers worldwide with low-latency, high-performance applications regardless of geographic location.

4 Scalability and Elasticity

AWS's massive global footprint provides virtually unlimited capacity to scale applications on demand, whether serving a local market or expanding globally.

Practice Exam

Online Training Course