

AWS CERTIFIED SOLUTION ARCHITECT: ASSOCIATE

AWS CERTIFIED SOLUTION ARCHITECT

ASSOCIATE



AMAZON WEB SERVICES

CLOUDFRONT

- Understanding Cloudfront
- Understanding Edge Locations
- Understanding Regional Edge Locations
- Cloudfront Lab

Cloudfront - Content Delivery Network (CDN)

Cloudfront is a service that securely delivers data, videos, applications, and APIs to your viewers with low latency and high transfer speeds.

CloudFront is integrated with AWS – including physical locations that are directly connected to the AWS global infrastructure, as well as software that works seamlessly with services including AWS Shield for DDoS mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications.

Its a distributed servers and is not the part of Regions/AZ



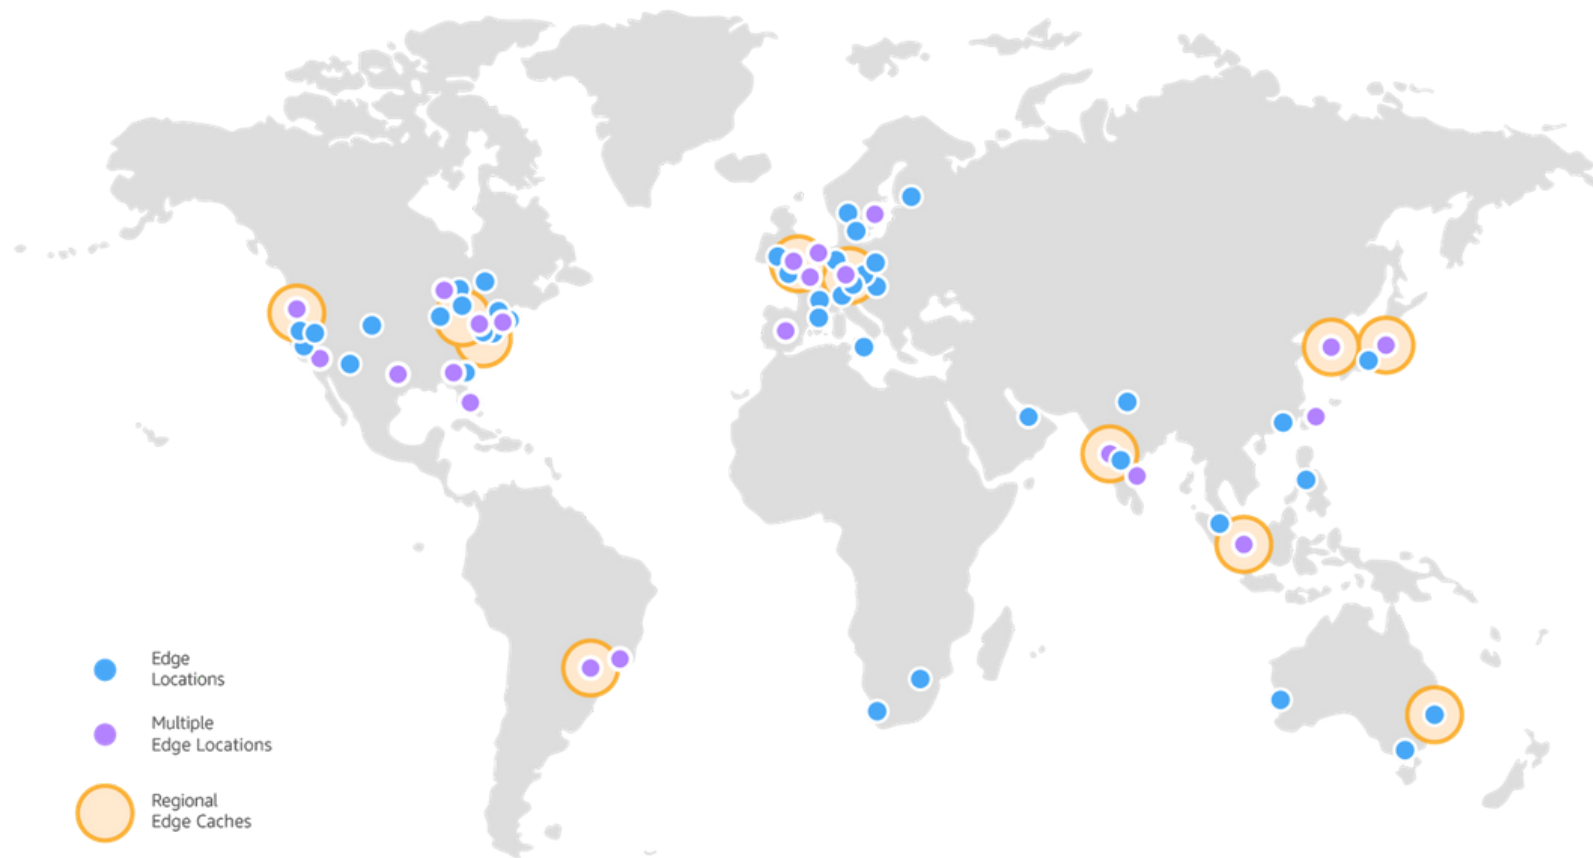
Cloudfront - Content Delivery Network (CDN)

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users.


CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.



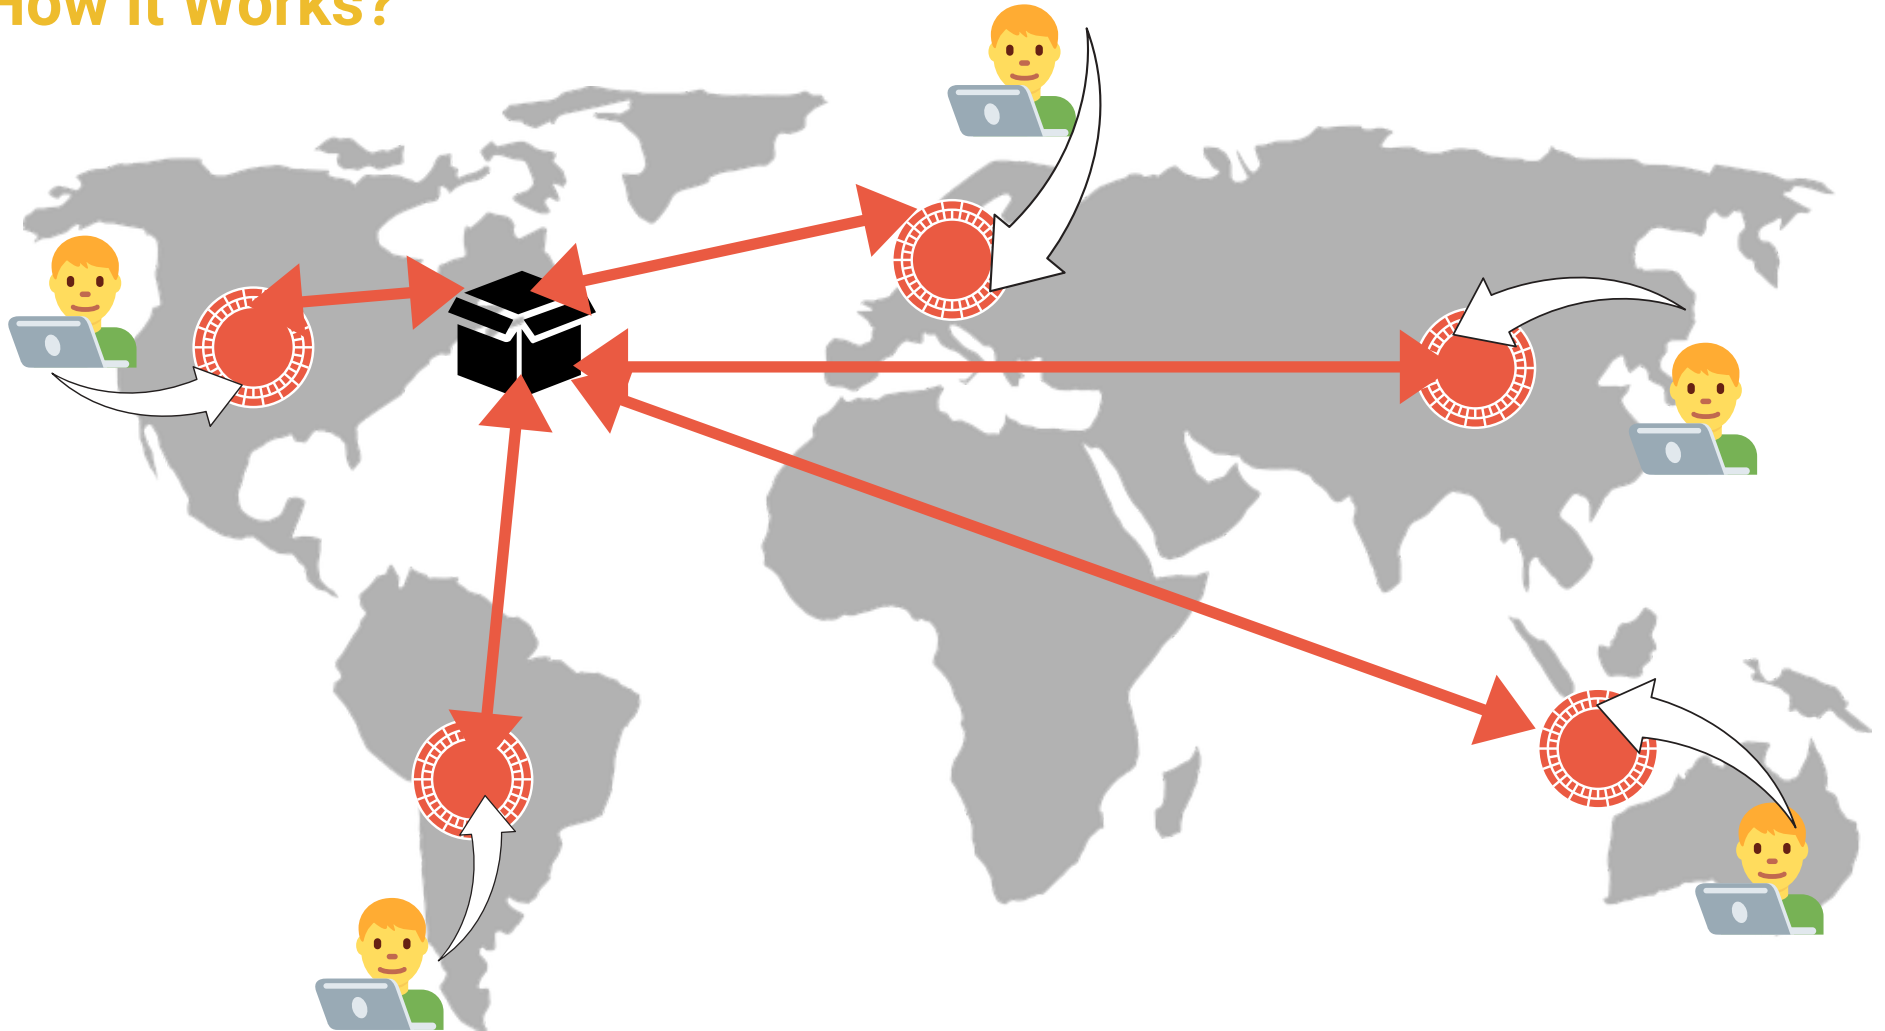
Amazon CloudFront uses a global network of 136 Points of Presence (125 Edge Locations and 11 Regional Edge Caches) in 62 cities across 29 countries. Amazon CloudFront Edge locations are located in: (as on September, 2018)




How it Works?



How it Works?





Step #1. You specify origin servers, like an Amazon S3 bucket or your own HTTP server, from which CloudFront gets your files which will then be distributed from CloudFront edge locations all over the world.

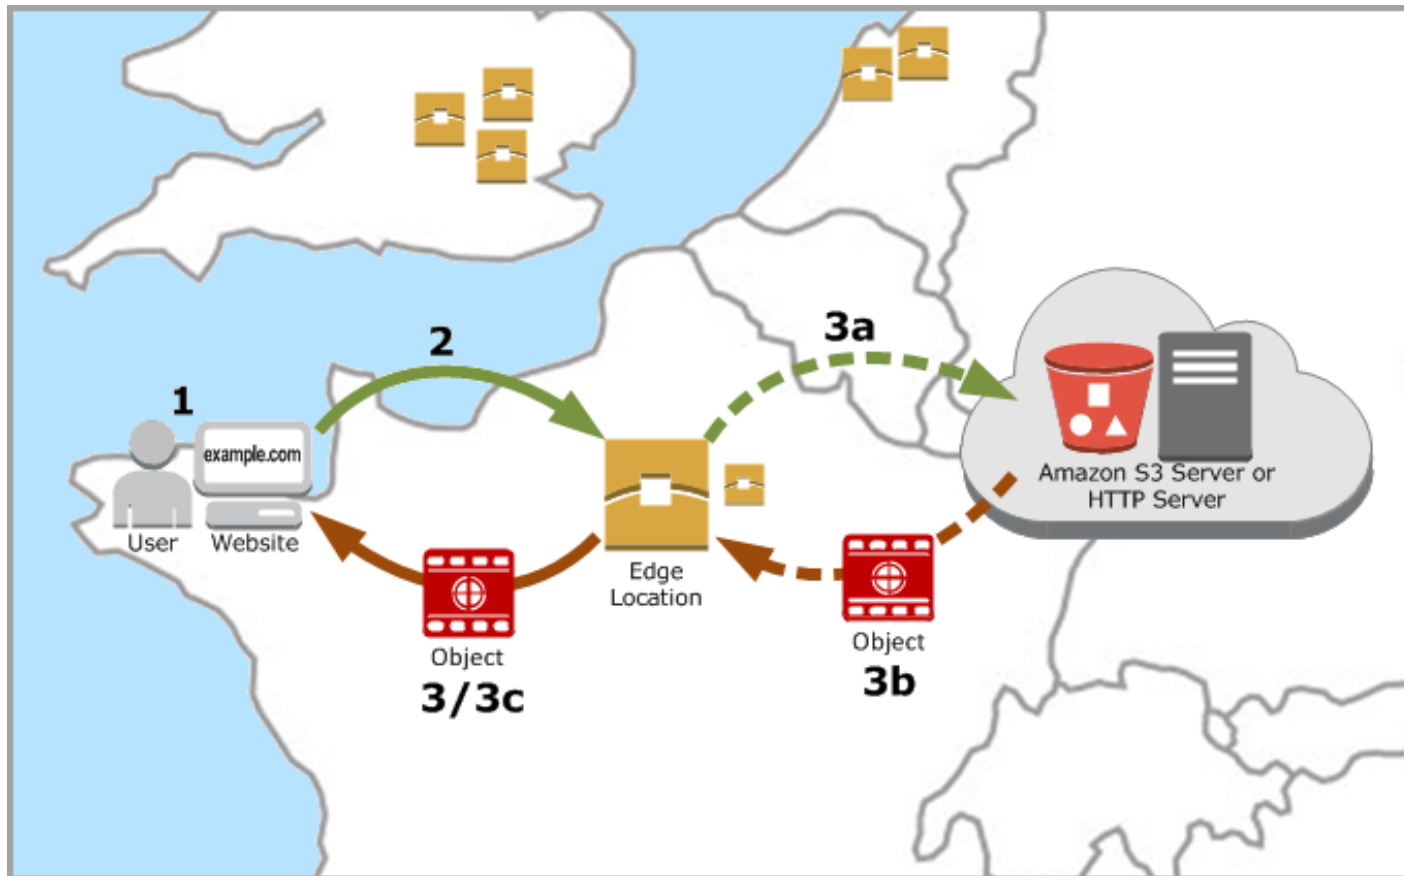
Step #2. You upload your files to your origin servers. Your files, also known as objects, typically include web pages, images, and media files, but can be anything that can be served over HTTP or a supported version of Adobe RTMP, the protocol used by Adobe Flash Media Server.

Step #3. You create a CloudFront distribution, which tells CloudFront which origin servers to get your files from when users request the files through your web site or application. At the same time, you specify details such as whether you want CloudFront to log all requests and whether you want the distribution to be enabled as soon as it's created.

Step #4. CloudFront assigns a domain name to your new distribution that you can see in the CloudFront console, or that is returned in the response to a programmatic request, for example, an API request.

Step #5. CloudFront sends your distribution's configuration (but not your content) to all of its edge locations—collections of servers in geographically dispersed data centers where CloudFront caches copies of your objects.

How Cloudfront Delivers



Step #1. A user accesses your website or application and requests one or more objects, such as an image file and an HTML file.

Step #2. DNS routes the request to the CloudFront edge location that can best serve the request—typically the nearest CloudFront edge location in terms of latency—and routes the request to that edge location.

Step #3. In the edge location, CloudFront checks its cache for the requested files. If the files are in the cache, CloudFront returns them to the user. If the files are not in the cache, it does the following:

- a. CloudFront compares the request with the specifications in your distribution and forwards the request for the files to the applicable origin server for the corresponding file type—for example, to your Amazon S3 bucket for image files and to your HTTP server for the HTML files.
- b. The origin servers send the files back to the CloudFront edge location.
- c. As soon as the first byte arrives from the origin, CloudFront begins to forward the files to the user. CloudFront also adds the files to the cache in the edge location for the next time someone requests those files.

Cloudfront - Edge Locations & Regional Edge Caches

CloudFront edge locations make sure that popular content can be served quickly to your viewers. CloudFront also has regional edge caches that bring more of your content closer to your viewers, even when the content is not popular enough to stay at a CloudFront edge location, to help improve performance for that content.

Regional edge caches help with all types of content, particularly content that tends to become less popular over time.

They're located between your origin server and the global edge locations that serve content directly to viewers.

As objects become less popular, individual edge locations may remove those objects to make room for more popular content. Regional edge caches have a larger cache than an individual edge location, so objects remain in the cache longer at the nearest regional edge cache location.

This helps keep more of your content closer to your viewers, reducing the need for CloudFront to go back to your origin server, and improving overall performance for viewers.

When a viewer makes a request on your website or through your application, DNS routes the request to the CloudFront edge location that can best serve the user's request. This location is typically the nearest CloudFront edge location in terms of latency. In the edge location, CloudFront checks its cache for the requested files. If the files are in the cache, CloudFront returns them to the user. If the files are not in the cache, the edge servers go to the nearest regional edge cache to fetch the object.

In the regional edge cache location, CloudFront again checks its cache for the requested files. If the files are in the cache, CloudFront forwards the files to the requested edge location. As soon as the first byte arrives from regional edge cache location, CloudFront begins to forward the files to the user. CloudFront also adds the files to the cache in the requested edge location for the next time someone requests those files.

For files not cached at both the edge location and the regional edge cache location, CloudFront compares the request with the specifications in your distributions and forwards the request for your files to the origin server.

After your origin server sends the files back to the regional edge cache location, they are forwarded to the requested edge location, and CloudFront forwards the files to the user.

In this case, CloudFront also adds the files to the cache in the regional edge cache location in addition to the edge location for the next time a viewer requests those files. This makes sure that all of the edge locations in a region share a local cache, eliminating multiple requests to origin servers. CloudFront also keeps persistent connections with origin servers so files are fetched from the origins as quickly as possible.