



AWS Webinar

Exploiting Multi-Region Data Locality with Lambda@Edge

Vijay Potharla, Senior Product Manager, AWS

Seth Fitzsimmons, Technical Fellow, Radiant Earth Foundation

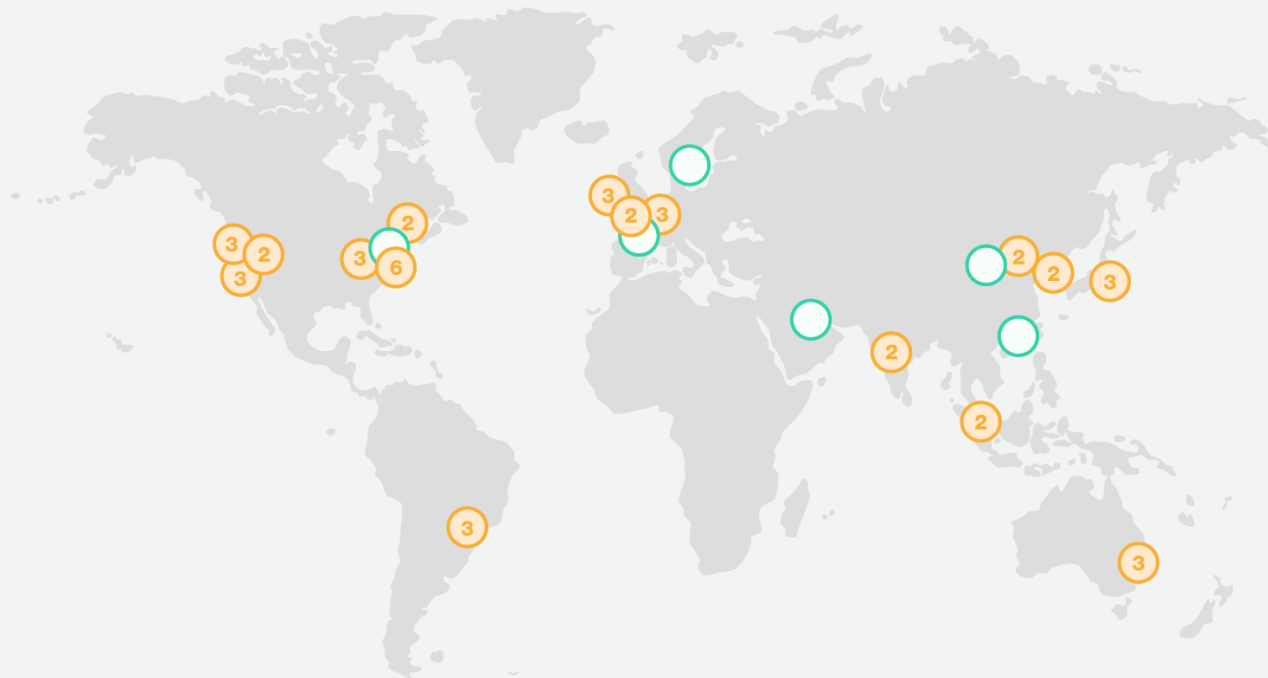
17 June 2019



What is covered in this session

- Lambda@Edge overview
 - What is Global Serverless Computing?
 - How Lambda@Edge works?
 - What can you do with Lambda@Edge?
- How Radiant Earth implemented multi-region data locality with Lambda@Edge ?

AWS Locations: Where is Your Customer?



180 Points
of Presence

69 cities

30 countries

21 AWS
Regions



AWS Lambda: Why Serverless?

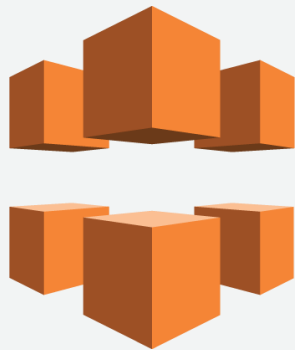
Build and run applications without managing servers

- No servers to manage
- Run at scale
- Respond quickly to events
- Only pay for compute time that you use
- Developer productivity



Amazon CloudFront: Global Content Delivery

- Global, growing content delivery network
- High performance—speed up distribution of your static and dynamic web content to your users across the world.
- DDoS resiliency, scale, and cost effective
- Flexible and highly “programmable”, with AWS Lambda@Edge



Global Serverless: Run Lambda Functions Across AWS Locations



Benefits of AWS Lambda@Edge



No
servers
to
manage



Continuous
scaling

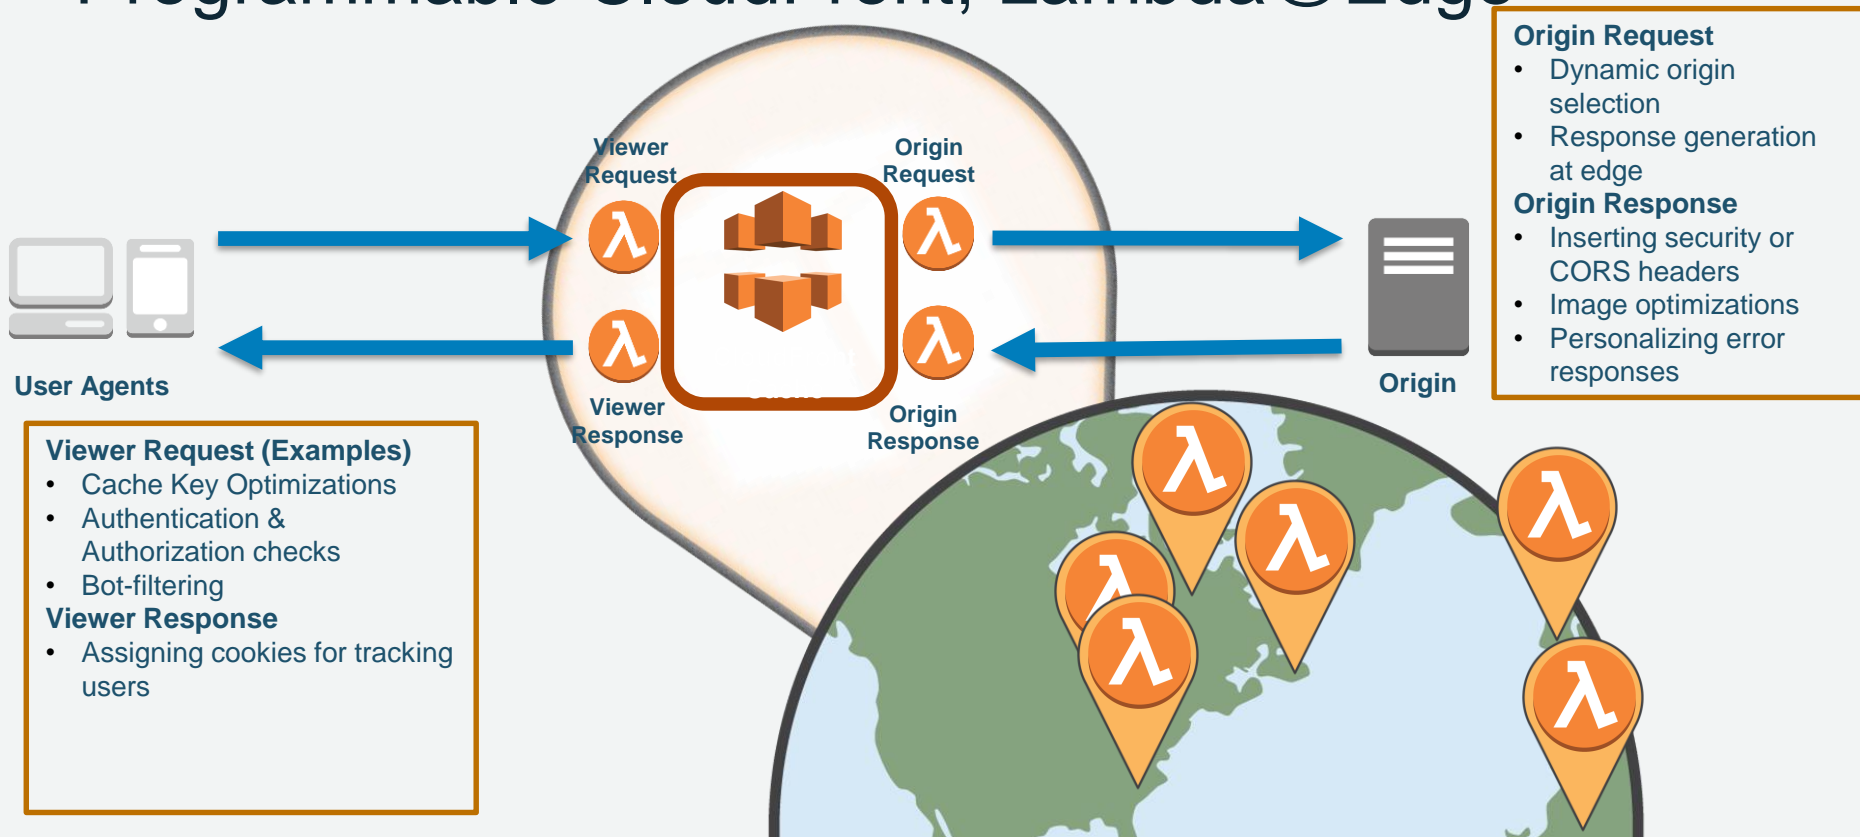


Never pay
for idle



Global
presence

Programmable CloudFront; Lambda@Edge



Exploiting Multi-Region Data Locality with Lambda@Edge

Seth Fitzsimmons for Radiant Earth Foundation

17 June 2019





Seth Fitzsimmons

Seth is an independent consultant who architects software solutions for raster imagery, data pipelines, and humanitarian needs.

Amongst other things, he is a Technical Fellow at Radiant Earth, focusing on tooling and evangelization of standards and specifications like COG (cogeo.org) and STAC (stacspec.org).



seth@pacatlas.com

  @mojodna



Radiant Earth Foundation

EARTH IMAGERY FOR IMPACT



About

- Non-profit, founded in 2016
- Funded by Omidyar Network, Gates Foundation, Schmidt Futures & McGovern Foundation

Vision

- Machine Learning and EO for Positive Impact

Expected Impact

- Democratize ML data & models
- High quality data for better SDG monitoring & decision-making
- Increase use of ML on EO

Services



Audience

**ML
Experts**

Library: Training Datasets and Models

**Standards &
Geospatial
Community**

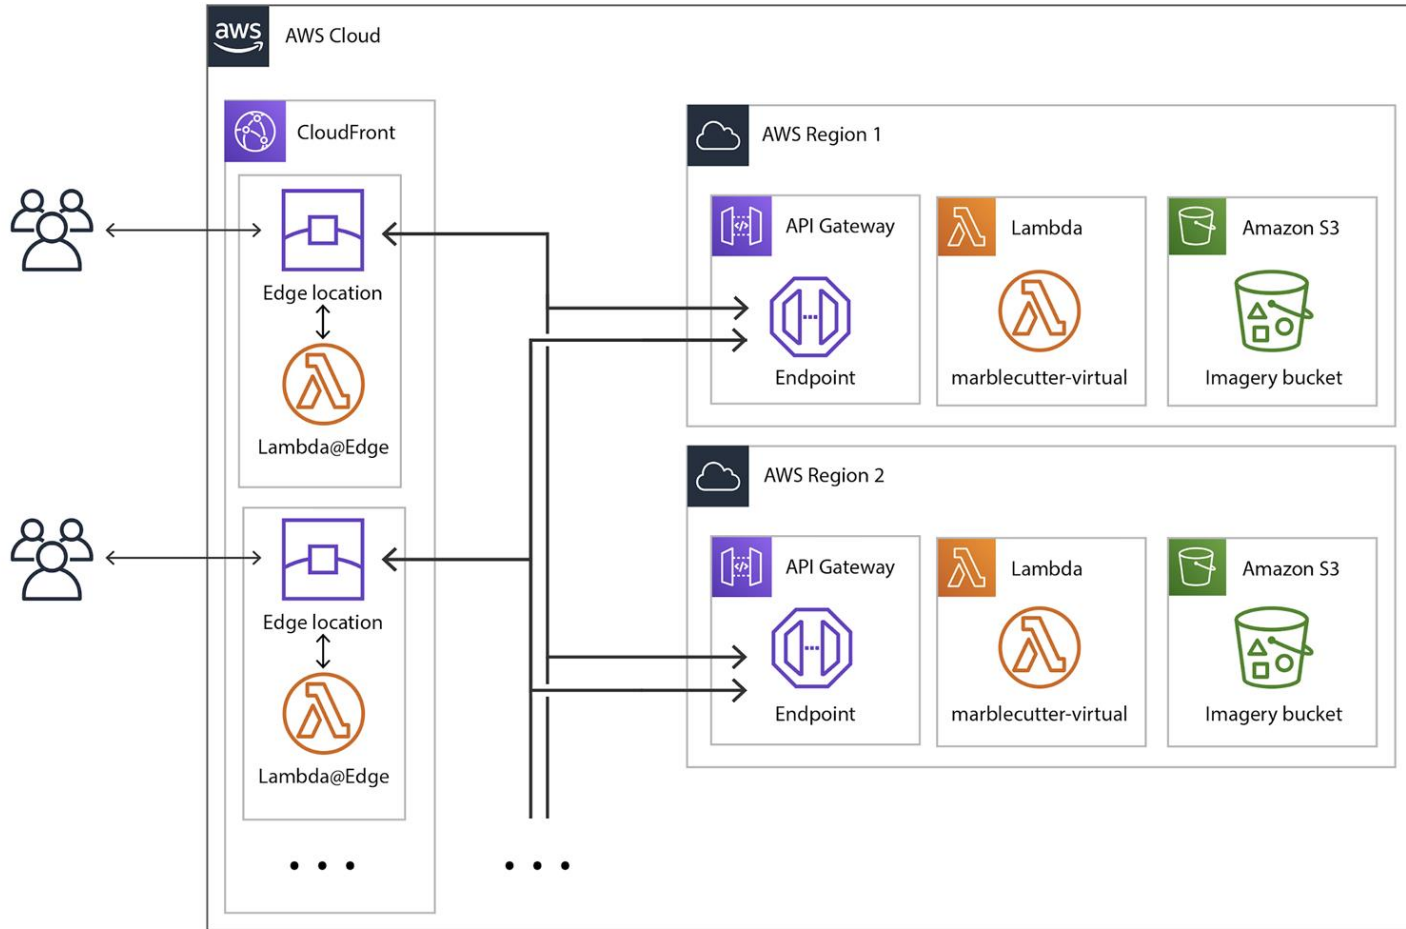
Community: Best Practices, Standards
(STAC), Fellowships, Events, Publications

**Public with interest in ML / EO
& Sustainable Development
Goals (SDGs)**

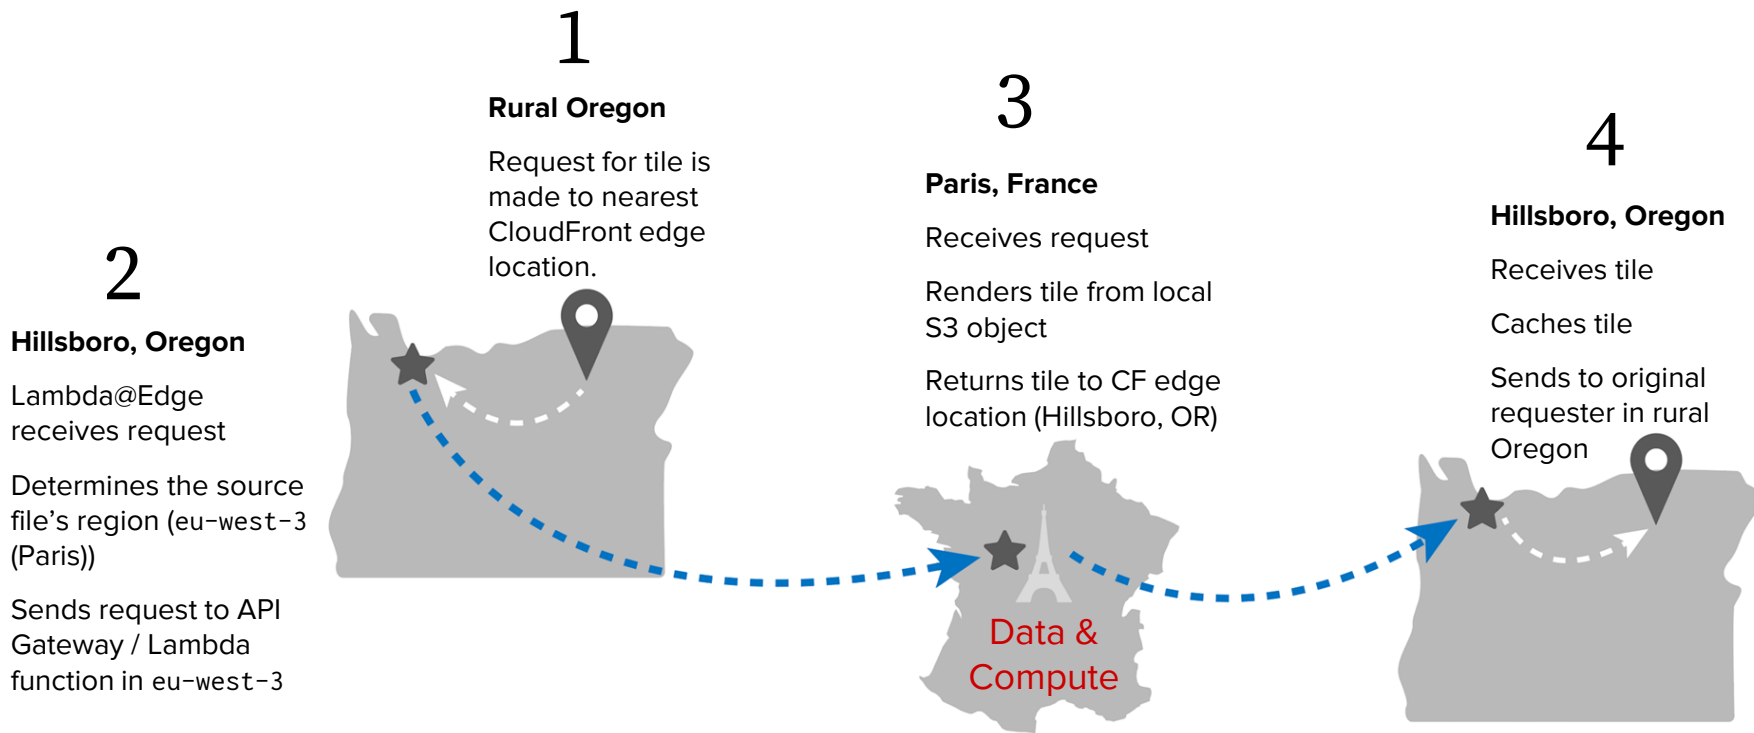
Raise awareness on Data, Policy, Ethics &
EO Market

Why we're here today:

<https://tiles.rdnt.io>



How it works

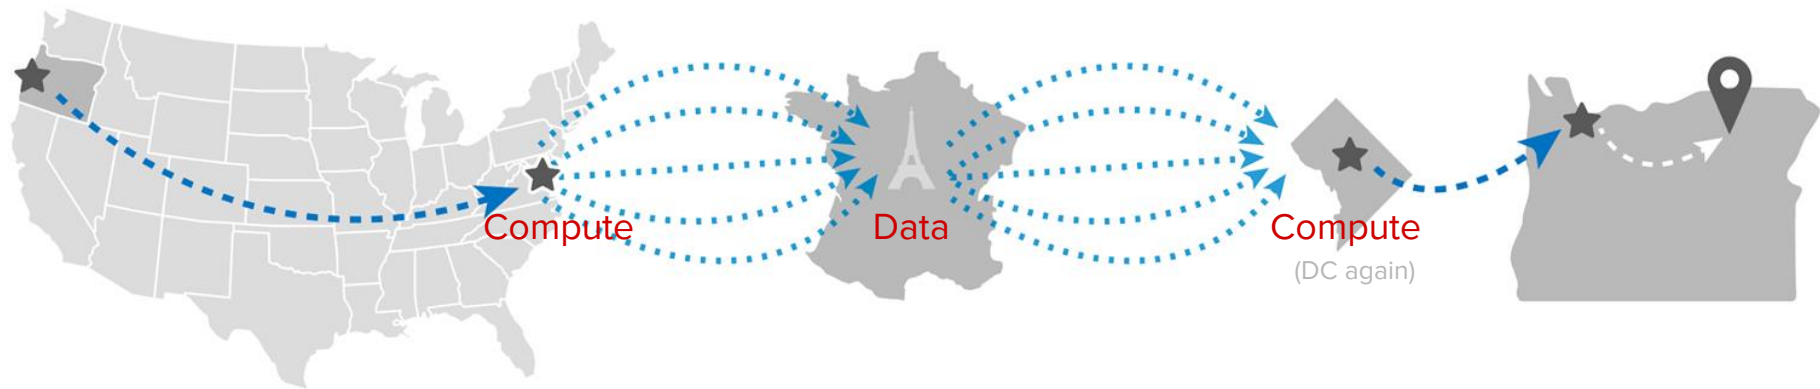


Example of how it works



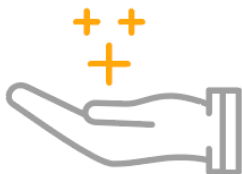
- Initially computed further from **users**, but closer to **data**
- Will be cached **close to users**

Before Lambda@Edge



- Data and compute functions are in different regions
- Multiple requests must be made between regions. This incurs latency and data transfer costs.

Benefits of Lambda@Edge + tiles.rdnt.io



- Tangible benefit for adopting a standardized format [Cloud-Optimized GeoTIFF (COG)]
- Able to browse multi-petabyte archives (HUGE) easily, interactively
- Affordable: free or cheap to access
Affordable: costs nothing when not in use
Affordable: minimize data transfer costs
- Access to imagery and data from AWS Public Dataset program


AWS Public Dataset Program

+

← → ↺

https://aws.amazon.com/opendata/public-datasets/

☆ ⚙ 🔴 🖋 👤 ⋮



Contact Sales Support ▾ English ▾ My Account ▾ [Sign In to the Console](#)

Products Solutions Pricing Documentation Learn Partner Network AWS Marketplace Explore More 🔍

AWS Public Dataset Program

The AWS Public Dataset Program covers the cost of storage for publicly available high-value cloud-optimized datasets. We work with data providers who seek to:

1. Democratize access to data by making it available for analysis on AWS.
2. Develop new cloud-native techniques, formats, and tools that lower the cost of working with data.
3. Encourage the development of communities that benefit from access to shared datasets.

You can see examples of datasets supported by the AWS Public Dataset Program on the [Registry of Open Data on AWS](#).

Requirements

To share your dataset through the AWS Public Dataset Program, you must agree to the AWS Public Dataset Program Terms and Conditions, which are available at: <https://aws.amazon.com/public-datasets/terms/>

Application process

AWS evaluates applications to the AWS Public Dataset Program every three months. The quarterly deadlines for submitting AWS Public Dataset Program applications are: March 31, June 30, September 30, and December 31 (or the first business day after those dates). Application decisions are typically communicated within one month following the applicable quarterly deadline.


If we bring your dataset into the AWS Public Dataset Program, we will cover the costs of storage and data transfer for a period of two years, in accordance with the AWS Public Dataset Program Terms and Conditions.

To apply, simply fill out the secure form at <https://application.opendata.aws>

Registry of Open Data on AWS

← → ↻ https://registry.opendata.aws

☆ ⚙ 🔴 🖌 👤 ⋮

Registry of Open Data on AWS

About

This registry exists to help people discover and share datasets that are available via AWS resources. [Learn more about sharing data on AWS.](#)

See [all usage examples for datasets listed in this registry.](#)

Search datasets (currently 110 matching datasets)

Search datasets

Add to this registry

If you want to add a dataset or example of how to use a dataset to this registry, please follow the instructions on the [Registry of Open Data on AWS GitHub repository](#).

Unless specifically stated in the applicable dataset documentation, datasets available through the Registry of Open Data on AWS are not provided and maintained by AWS. Datasets are provided and maintained by a variety of third parties under a variety of licenses. Please check dataset licenses and related documentation to determine if a dataset may be used for your application.

Sentinel-2

earth observation satellite imagerygeospatialnatural resource sustainabilitydisaster response

The [Sentinel-2 mission](#) is a land monitoring constellation of two satellites that provide high resolution optical imagery and provide continuity for the current SPOT and Landsat missions. The mission provides a global coverage of the Earth's land surface every 5 days, making the data of great use in on-going studies. L1C data are available from June 2015 globally. L2A data are available from April 2017 over wider Europe region and globally since December 2018.

[Details →](#)

Usage examples

- [Using Vector tiles and AWS Lambda, we can build a really simple API to get Landsat and Sentinel images](#) by Remote Pixel
- [FME Landsat-8/Sentinel-2 File Selector](#) by Safe Software
- [Sentinel Playground](#) by Sinergise
- [Exploring the Chile wildfires with Landsat and Sentinel-2 imagery](#) by Timothy Whitehead
- [EO Browser](#) by Sinergise

[See 16 usage examples →](#)

Landsat 8

earth observation satellite imagerygeospatialnatural resource sustainabilitydisaster response

An ongoing collection of satellite imagery of all land on Earth produced by the Landsat 8 satellite.

[Details →](#)

Usage examples

- [EOS Land Viewer](#) by Earth Observing System
- [Development Seed Geolambda](#) by Matthew Hanson
- [Using Vector tiles and AWS Lambda, we can build a really simple API to get Landsat and Sentinel images](#) by Remote Pixel

Amazon Sustainability Data Initiative



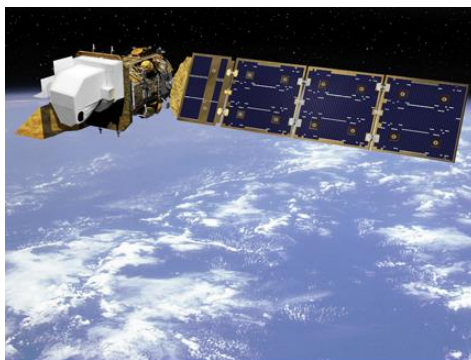
blog.aboutamazon.com/sustainability/unlocking-sustainability-insights-around-the-world

aws.amazon.com/blogs/publicsector/geo-diverse-open-training-data-as-a-global-public-good/

Sentinel-2



Landsat-8

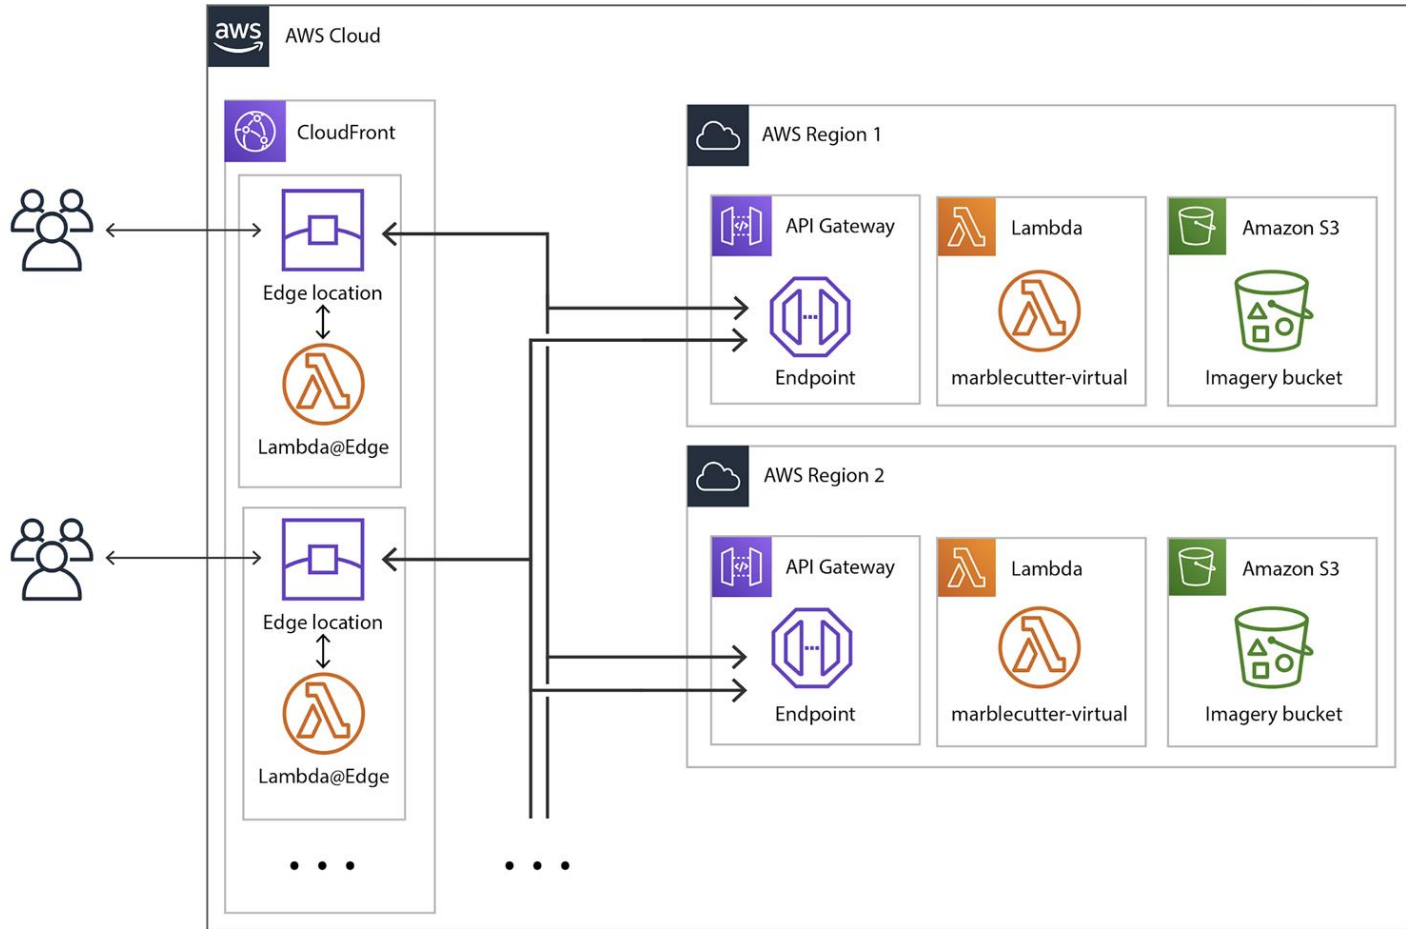


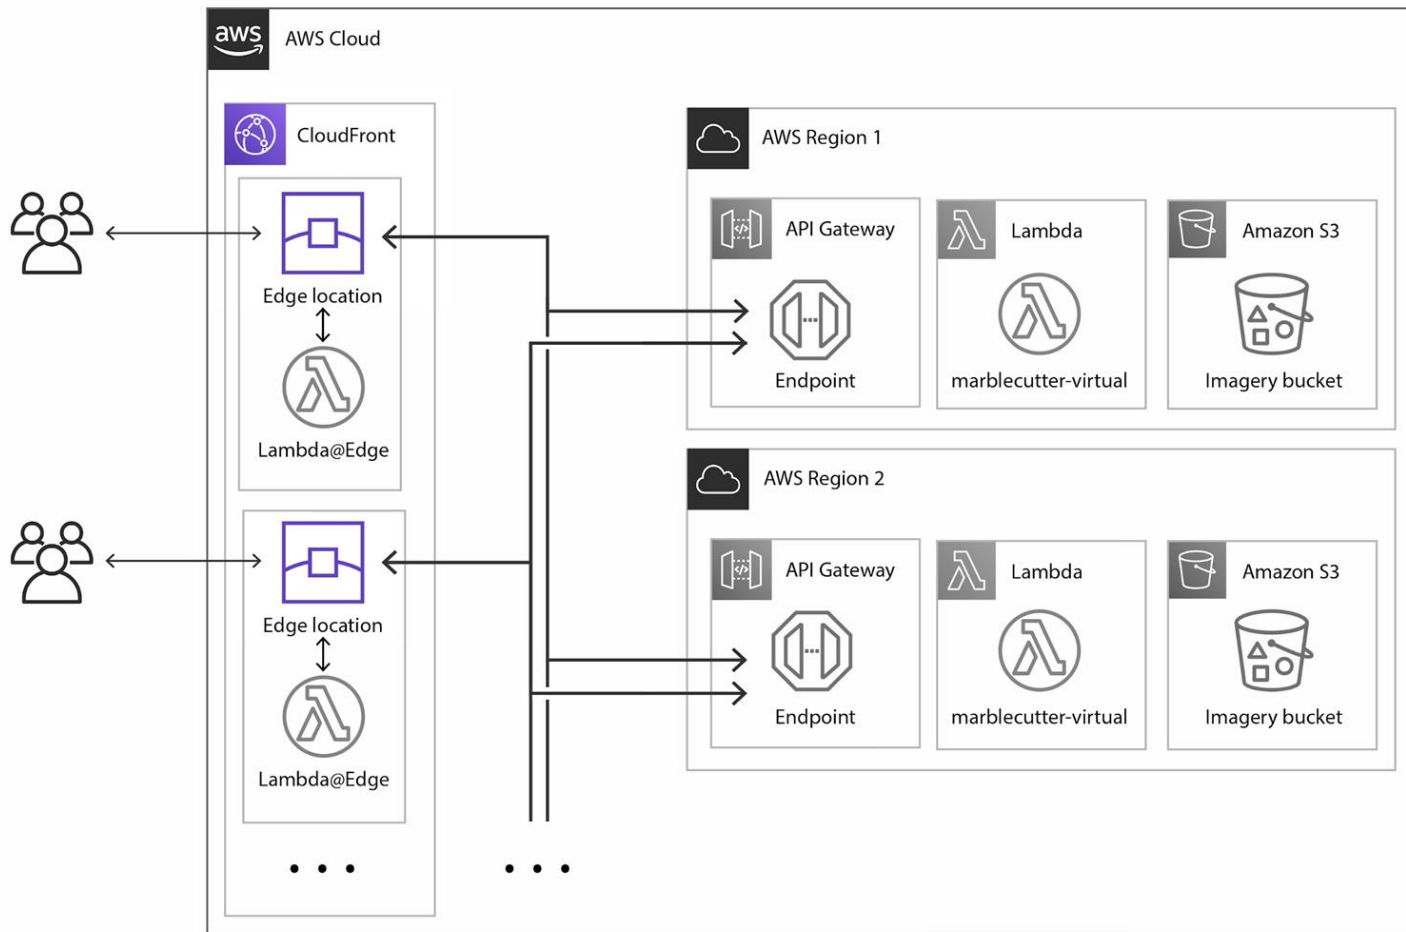
NAIP



Architecture

<https://tiles.rdnt.io>





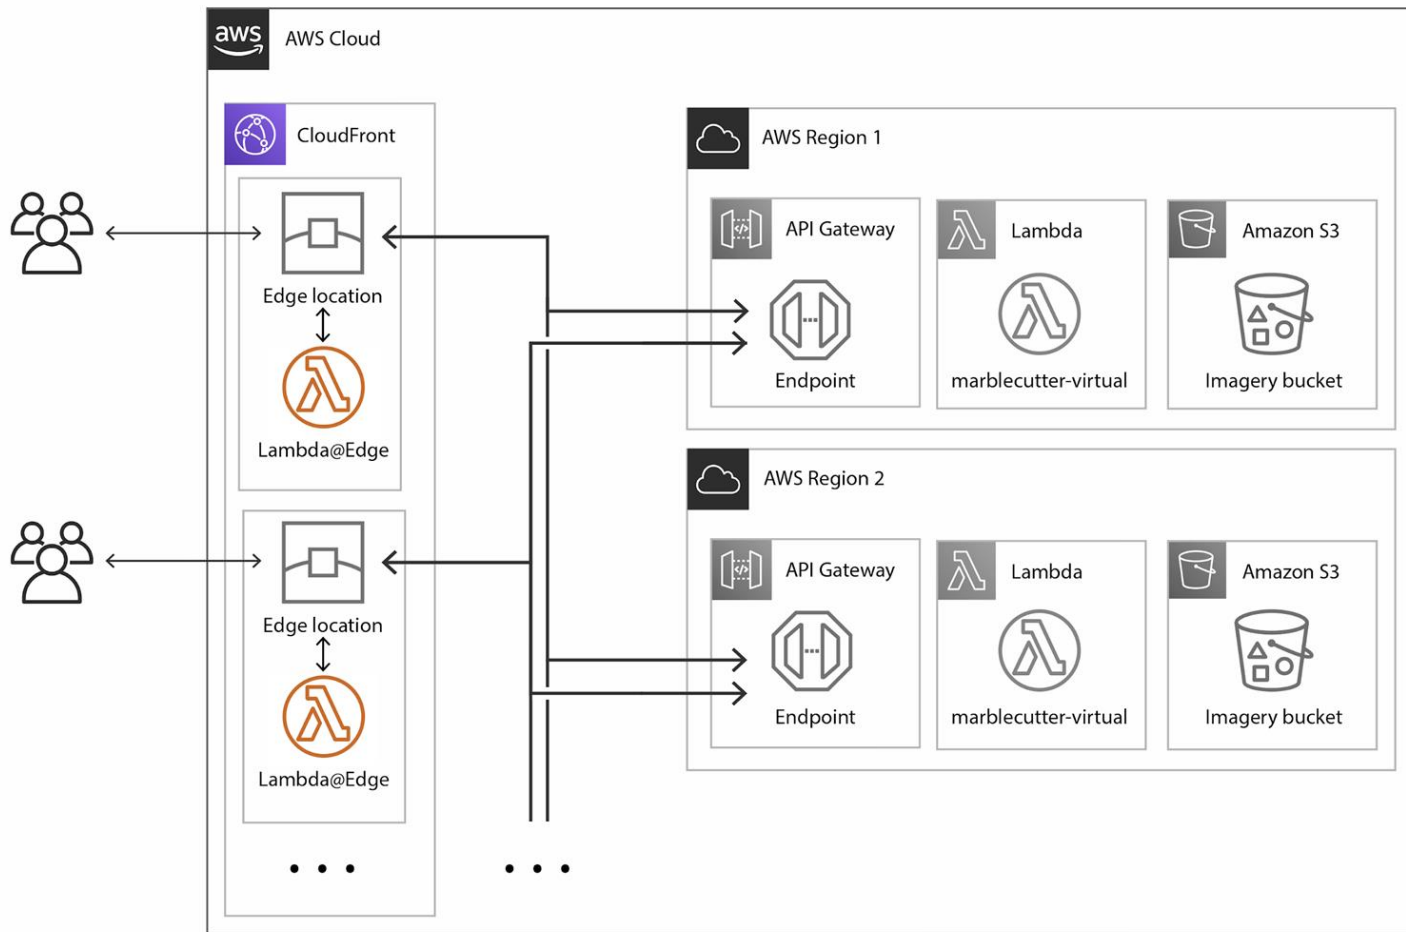
Anatomy of a request

CF ENDPOINT
ZOOM *TILE X* *TILE Y* *SCALE* *TARGET URL*

https://tiles.rdnt.io/tiles/16/10679/23825@2x?url=s3://naip-visualization/or/2016/100cm/rgb/44121/m_4412162_sw_10_1_20160621.tif

=





Lambda@Edge Function

```
exports.handler = async (event, context, callback) => {  
  const request = event.Records[0].cf.request; GET THE REQUEST OBJECT  
  
  const qs = querystring.parse(request.querystring); PARSE THE QUERY STRING  
  
  if (qs.url == null) {  
    return callback(null, request);  
  }  
  
  const uri = url.parse(qs.url); EXTRACT THE TARGET URL  
  
  const region = await getRegion(uri); FIND THE TARGET URL'S REGION  
  let origin = MAPPING[region];  
  
  if (origin == null) {  
    region = DEFAULT_REGION; LOOKUP UP THE API GATEWAY ID  
    origin = MAPPING[region];  
  }  
  
  REWRITE THE ORIGIN REQUEST  
  request.origin.custom.domainName = `${origin}.execute-api.${region}.amazonaws.com`;   
  request.origin.custom.path = "/Prod";  
  
  request.headers["host"] = [  
    { key: "host", value: `${origin}.execute-api.${region}.amazonaws.com` }  
  ];  
  
  return callback(null, request);  
};
```

1. Extract the bucket



The diagram illustrates the structure of an S3 URI: `s3://naip-visualization/or/2016/100cm/rgb/44121/m_4412162_sw_10_1_20160621.tif`. A horizontal line with vertical end-caps is positioned above the URI. The segment from the start to the first slash is labeled "BUCKET" in a handwritten style. The segment from the first slash to the end is labeled "KEY" in a handwritten style.

s3://naip-visualization/or/2016/100cm/rgb/44121/m_4412162_sw_10_1_20160621.tif

2. Determine the bucket's region

```
$ host naip-visualization.s3.amazonaws.com
```

naip-visualization.s3.amazonaws.com is an alias for s3-**us-west-2**-w.amazonaws.com.
s3-us-west-2-w.amazonaws.com has address 52.218.232.227

```

const getRegion = async ({ hostname }) => {
  try {
    if (uri.protocol === "s3:") {
      hostname += ".s3.amazonaws.com";
    }
    CONVERT TO A HOSTNAME + RESOLVE

    let matches;
    if (hostname.endsWith(".s3.amazonaws.com")) {
      const addresses = await resolveCname(hostname);

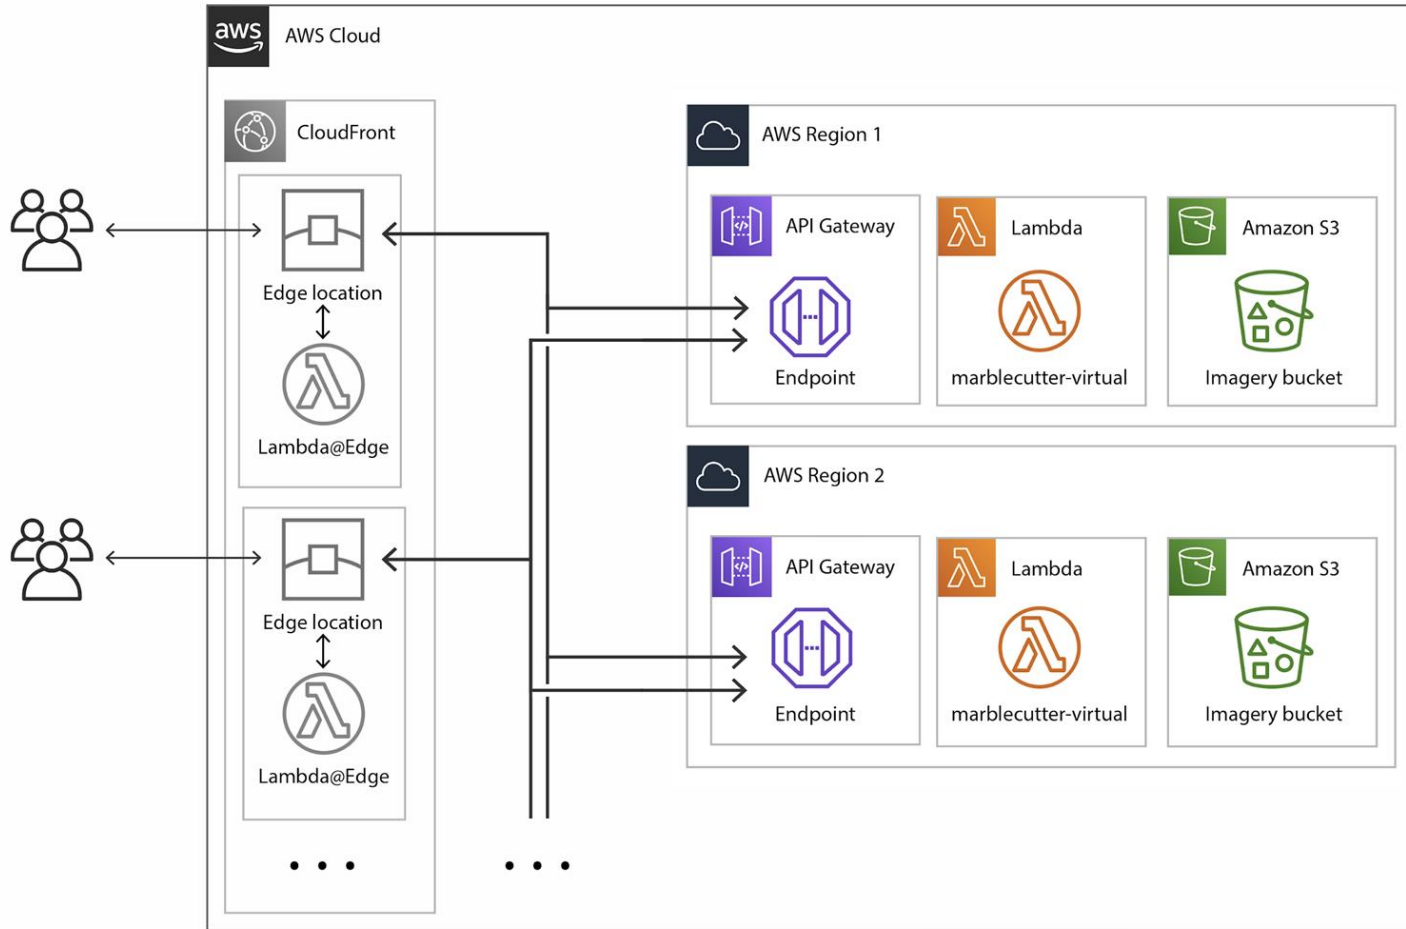
      if (addresses.length > 0) {
        if (
          (matches = addresses[0].match(
            /s3-([\w-]+)-w.amazonaws\.com/
          ))
        ) {
          EXTRACT THE REGION
          if (matches[1] === "1") {
            return "us-east-1";
          }
          SPECIAL CASES
          if (matches[1] === "3") {
            return "eu-west-1";
          }

          return matches[1];
        } else if (
          (matches = addresses[0].match(
            /s3(-w(ebsite)?)(\.dualstack)?[.-]([\w-]+)\.amazonaws\.com/
          ))
        ) {
          return matches[4];
        }
      }
    } else if (
      (matches = uri.hostname.match(
        /s3(-w(ebsite)?)(\.dualstack)?[.-]([\w-]+)\.amazonaws\.com/
      ))
    ) {
      return matches[4];
    }
  } finally {
    return DEFAULT_REGION;
  }
};

```

3. Set the origin to an API Gateway in that region

```
{  
  "us-east-1": "gkh62xc7he",  
  "us-east-2": "pmc4xepmhj",  
  "us-west-1": "9z8wrtc0y2",  
  "us-west-2": "p6q1upto2a",  
  "ap-south-1": "3ozyebs7ic",  
  "ap-northeast-1": "09ydnoc09g",  
  "ap-northeast-2": "1oj45tquq3",  
  "ap-southeast-1": "18lsnrvvvf",  
  "ap-southeast-2": "1xgmgnwxf",  
  "ca-central-1": "813ilrpqtg",  
  "eu-central-1": "1nk7vx4ye4",  
  "eu-west-1": "b6jx1o8lbi",  
  "eu-west-2": "aexc3bj0bi",  
  "eu-west-3": "bz2jkis2x3",  
  "eu-north-1": "96pywtwas9",  
  "sa-east-1": "b3bz87h9n7"  
}
```



Co-locate **interactive** compute with data



Examples

<https://tiles.rdnt.io> enables other applications

ISERV STAC browser

IPR201410161410290729N012

← → ↺ https://iserv.stac.cloud/item/UMYsgQCkrMHGTJjF34u7onu/hZ8chRfyMqL8F5i4ZhfiJBNzsmw/2389... ☆ 🔍 🌐 📄 👤 ⋮

ISERV / 2014 / 2014/10 / 2014/10/16 / IPR201410161410290729N01223W

IPR201410161410290729N01223W

<https://iserv-stac.s3.amazonaws.com/0.6.1/2014/10/16/IPR201410161410290729N01223W.json>

Preview

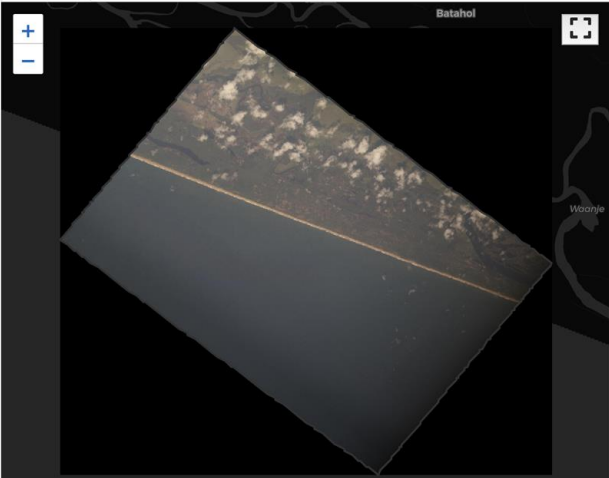
Thumbnail

Assets

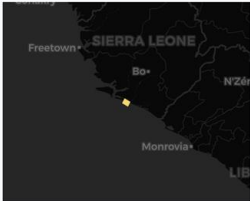
Bands

+

-

A satellite image showing a coastal area with a river and a road. The image is tilted and has a black border. The text 'Batahol' is visible in the top right corner of the image area. The text 'Woonje' is visible in the bottom right corner of the image area.

Map data © OpenStreetMap contributors, © CARTO, Imagery PDDL-1.0 SERVIR

A map of Sierra Leone with a yellow dot indicating the location of the image. The map shows the coastline and some major cities like Freetown, Bo, and Monrovia. The text 'SIERRA LEONE' is visible at the top of the map.

License

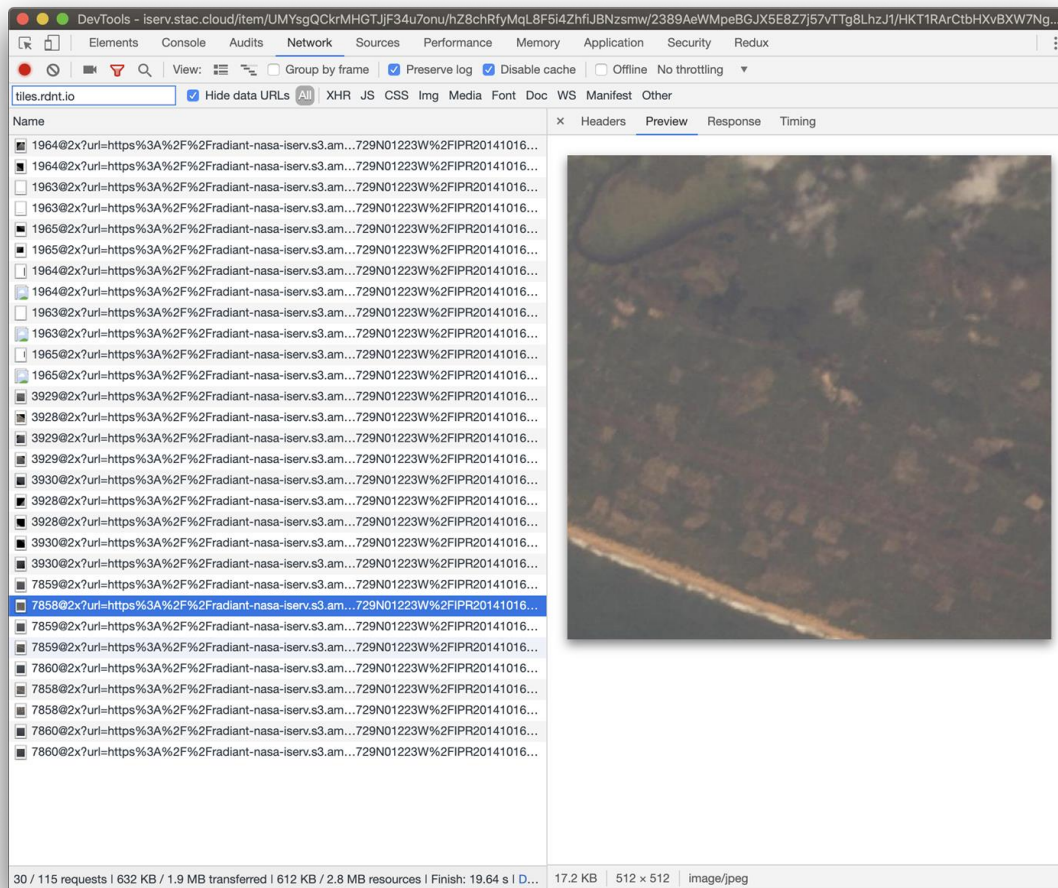
PDDL-1.0 by SERVIR

Acquired

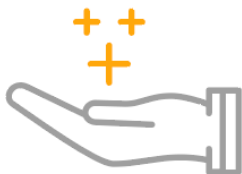
10/16/2014, 2:10:29 PM UTC

Powered by STAC Browser

ISERV STAC browser



Customer benefits



- Region-agnostic
- Lower bandwidth costs
- Lower S3 costs
 - No need for multi-region replication
 - No cross-region S3 reads
- Lower latency

Cloud-Optimized Files / Formats

Principles of Cloud-Optimized Files / Formats



- Predictable filenames / S3 keys
- A few large objects are better than many small objects
- Specify well-known offsets to avoid random reads to derived locations
- HTTP Range Request-friendly
- Related data is contiguous / local
- Block-level compression
- Block-level summaries
- Multi-resolution

Benefits of Cloud-Optimized Formats



- Doesn't read everything
- Lower costs
- Customers don't need copies of data
- Reduced data transfer when accessed interactively
- Easier object management
- No need for external indices

Examples of Cloud-Optimized Formats



- Web video
- ORC
- Parquet
- Cloud-Optimized GeoTIFF (COG)
- ZIP

Questions?

