Messaging Services and Distributed Systems in the Cloud

Jed Sundwall, Global Open Data Lead

What is cloud computing?

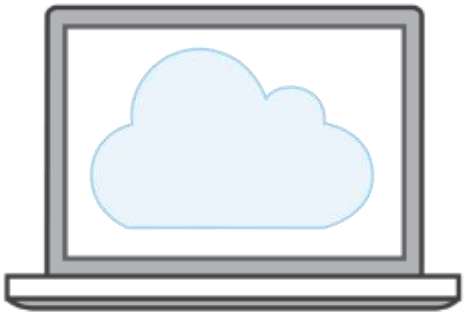

The **on-demand** delivery of IT resources over public or private networks with **zero up-front costs**, **no long-term contracts**, and **pay-as-you-go pricing**.

“Invention requires two things: the ability to try a lot of experiments, and not having to live with the collateral damage of failed experiments.”

— Andy Jassy, CEO, AWS

“Everything fails, all the time.”

— Werner Vogels, CTO, Amazon

Loosely-coupled systems

The **looser** they are coupled,

the **bigger** they will scale,

the more **fault tolerant** they will be,

the fewer **dependencies** they will have,

the **faster** you will innovate.

CAP* Theorem

Consistency - A read is guaranteed to return the most recent write for a given client.

Availability - A non-failing node will return a reasonable response within a reasonable amount of time (no error or timeout).

Partition Tolerance - The system will continue to function when network partitions occur.

* A different CAP than what we've been talking about!

Messaging enables decoupling

Message Queueing

- Asynchronous
- Point-to-point

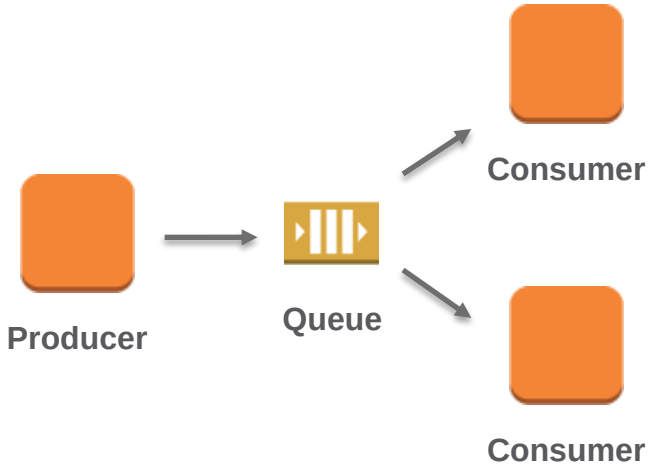

Publish-subscribe (pub-sub)

- Broadcast
- Point-to-multipoint

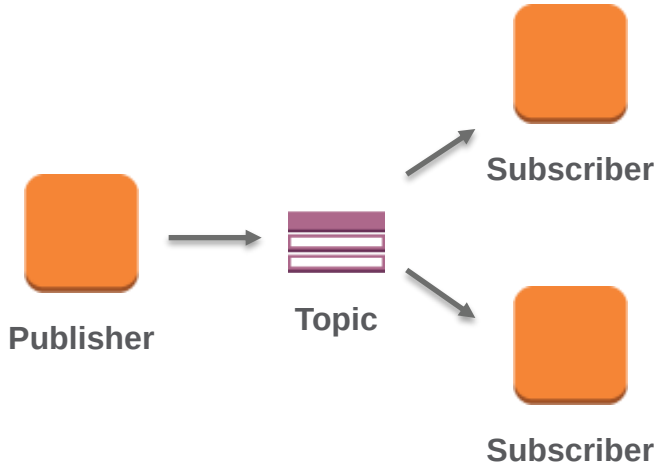

When to use messaging

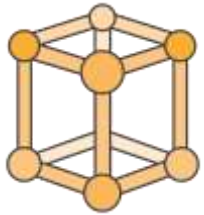

Separate parts of an application

- Web tier instances create work, worker instances complete it
- Scale and manage tiers separately

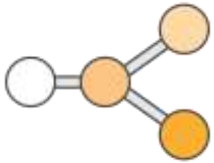

Perform tasks asynchronously

- Long-running tasks (e.g. transcoding, transactions)
- Don’t need to wait for a response (e.g. JS web apps)
- Independent and fault-tolerant

When to use messaging

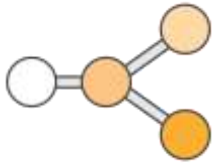

Connect multiple components

- Send individual messages or fan-out to many recipients
- Provide instant or delayed notification

Batch and burst processing

- Be resilient to spikes in traffic
- Perform work only as fast as necessary to lower costs
- Don’t lose data

Amazon Simple Notification Service (Amazon SNS)

- Fast, reliable, scalable fully managed pub-sub service
- Message notifications pushed to subscribers
- Use topics to fan out messages to:
  - Amazon SQS queues
  - HTTP endpoints (web servers)
  - AWS Lambda functions
  - Mobile push, SMS, and email

Amazon SNS: key features

- Proven reliability with messages are stored across multiple AZs

- Flexible message delivery over multiple transport protocols

- Nearly unlimited throughput

Amazon SNS: key features

- Instantaneous or delayed, push-based delivery

- Simple APIs and easy integration

- Amazon CloudWatch metrics and alerts

- Message payloads up to 256 KB

SNS is Simple to Use

1. Create Topic
2. Subscribe
3. Publish

See full API list:
http://docs.aws.amazon.com/sns/latest/api/Welcome.html

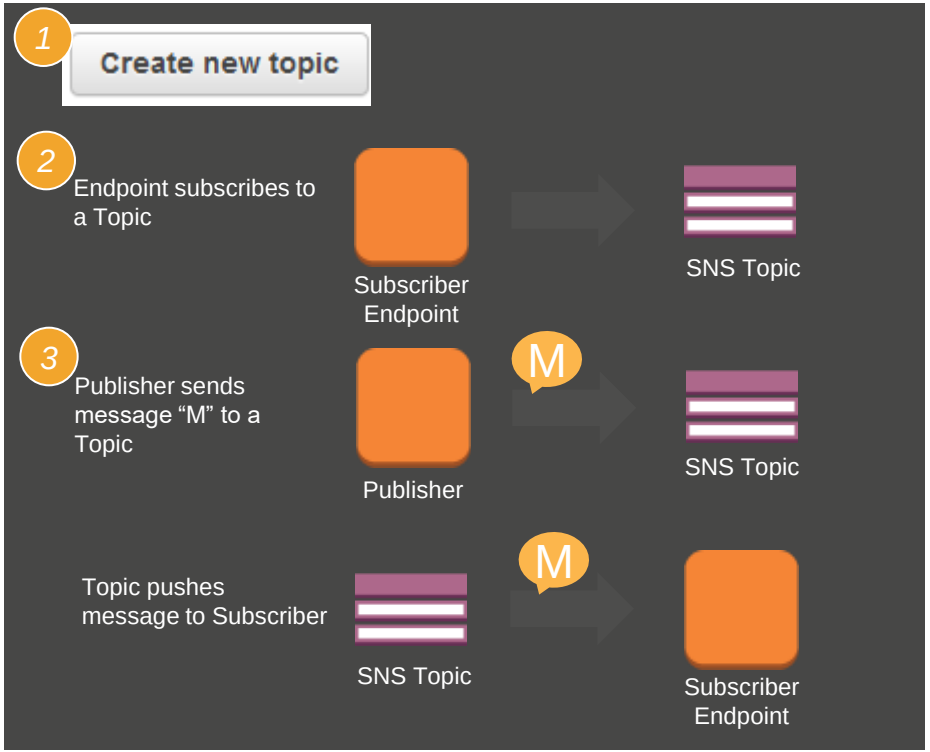

Earthquake Response Prototype

Goal:

To provide meaningful, actionable information as quickly as possible to respond to disaster situations arising from earthquakes.

Inputs/Outputs

- earthquake.usgs.gov
- USGS Landsat satellite imagery
- ESA Sentinel-1 satellite imagery
- ESA Sentinel-2 satellite imagery
- OpenStreetMap data
- Bonus: bring your own high resolution imagery

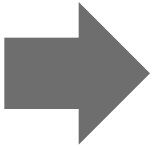

- A publicly accessible S3 bucket with:
  - pre and post event optical and SAR imagery
  - a list of medical facilities within a 1,000 km radius
  - available within seconds of reporting

Infrastructure Diagram, SNS feed

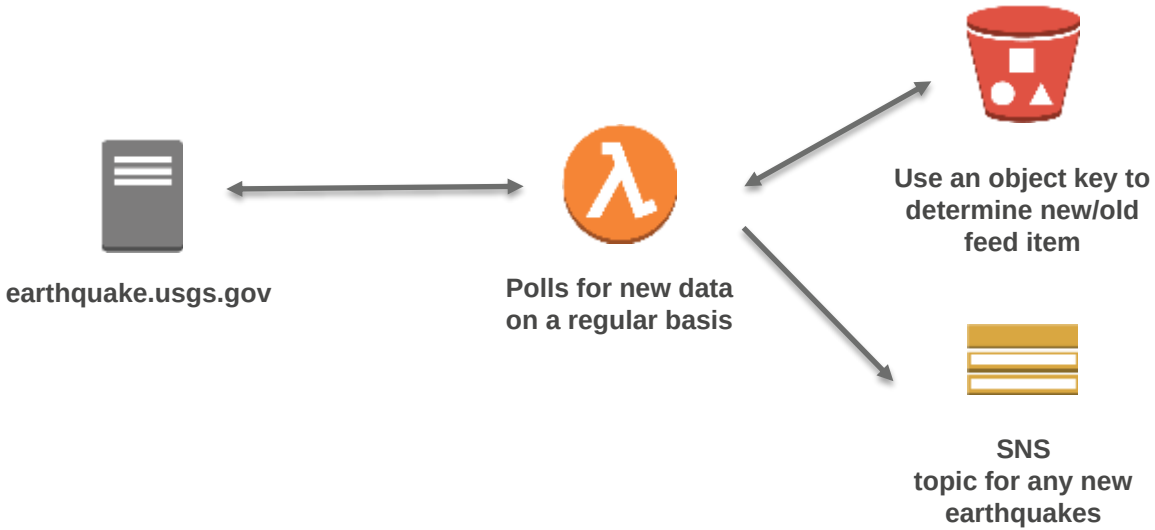

Infrastructure Diagram, data handler

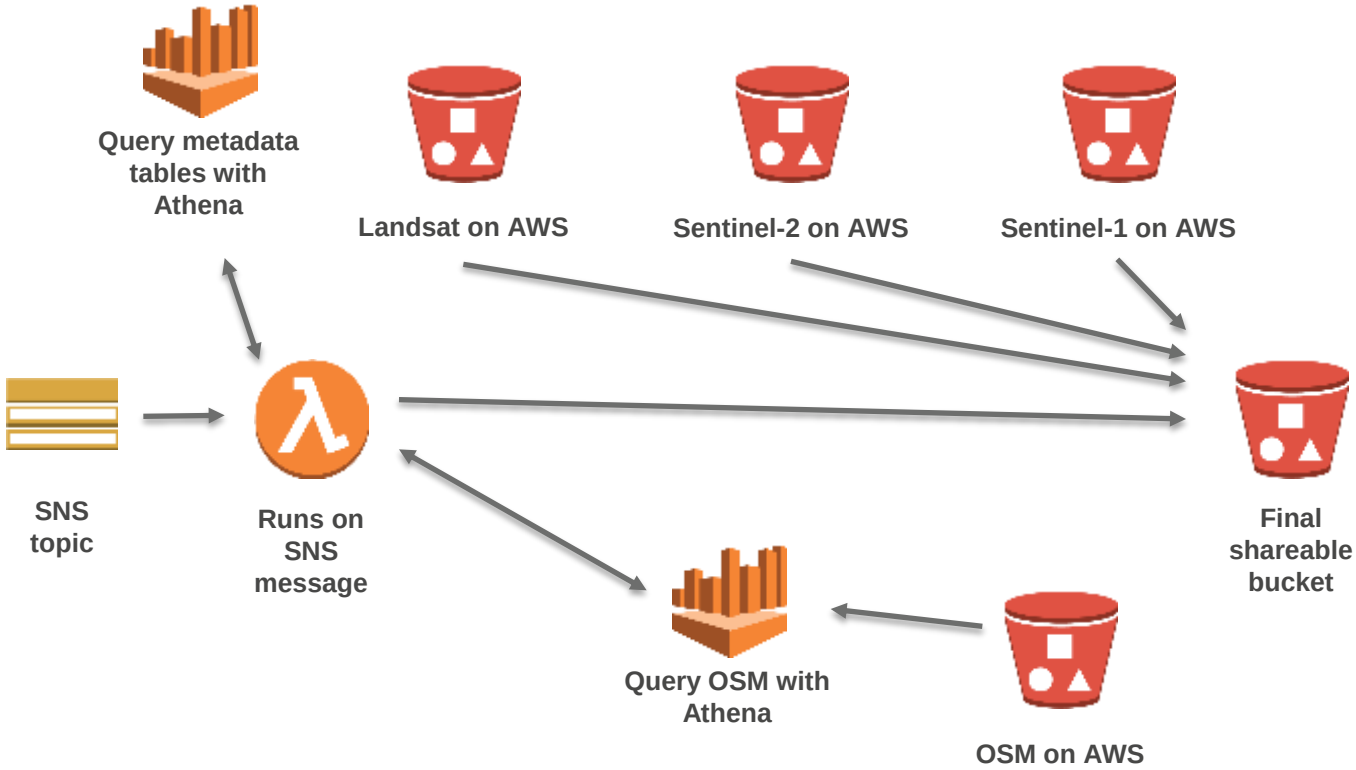

See it live!

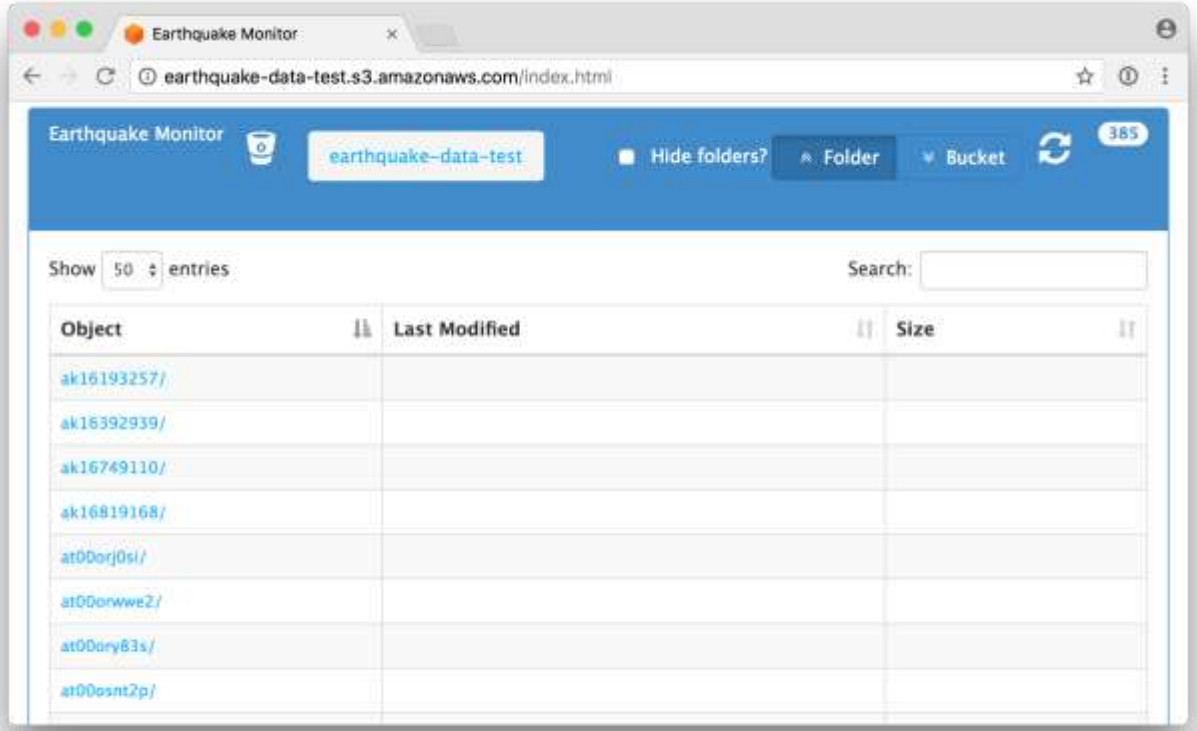

http://earthquake-data-test.s3.amazonaws.com/index.html

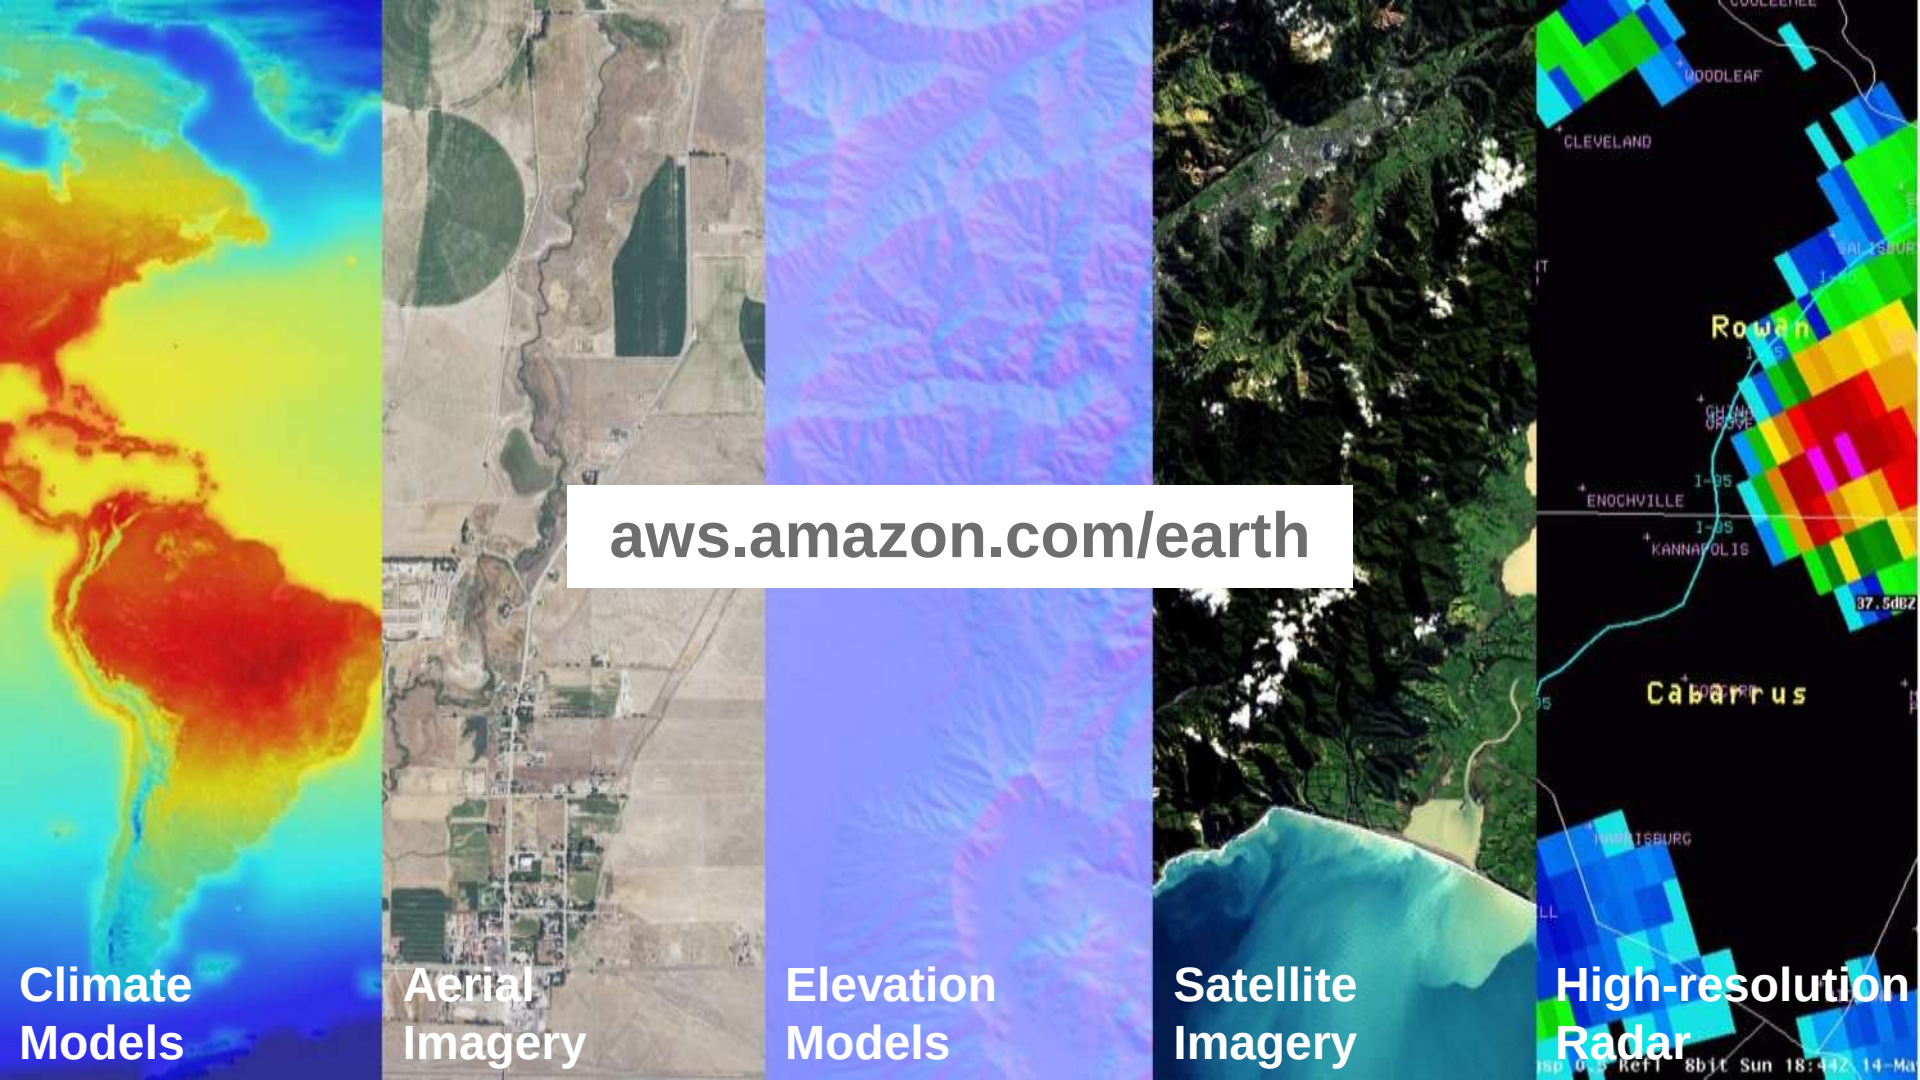
aws.amazon.com/earth
Climate Models
Aerial Imagery
Elevation Models
Satellite Imagery
High-resolution Radar
WOODLEAF
CLEVELAND
Rowan
ENOCHVILLE
KANNAPOLIS
Cabarrus
37.5dBZ

AWS Cloud Credits for Research provide promotional AWS cloud credits for **anyone** to conduct research using Earth Observation data.

aws.amazon.com/earth/research-credits