# Orion: Under the Hood

Jharrod LaFon
VP, Cloud Development

OpenEye SCIENTIFIC

# Goals

- A better understanding of
  - What a **service** is
  - How Orion is composed of **services**
  - What Amazon **services** Orion uses

- How Orion is designed to respond to **service** disruptions

# Background

- Teams
  - User interface (front-end)
  - Orion platform & APIs (back-end)
  - Services team

Designed & implemented Orion’s original architecture

# What's a service?

- Software accessed over a network (often the Internet)
- Access is facilitated by sending and receiving messages over the network
  - Client & Server
  - Request & Response
- A single service is typically targeted at a specific domain

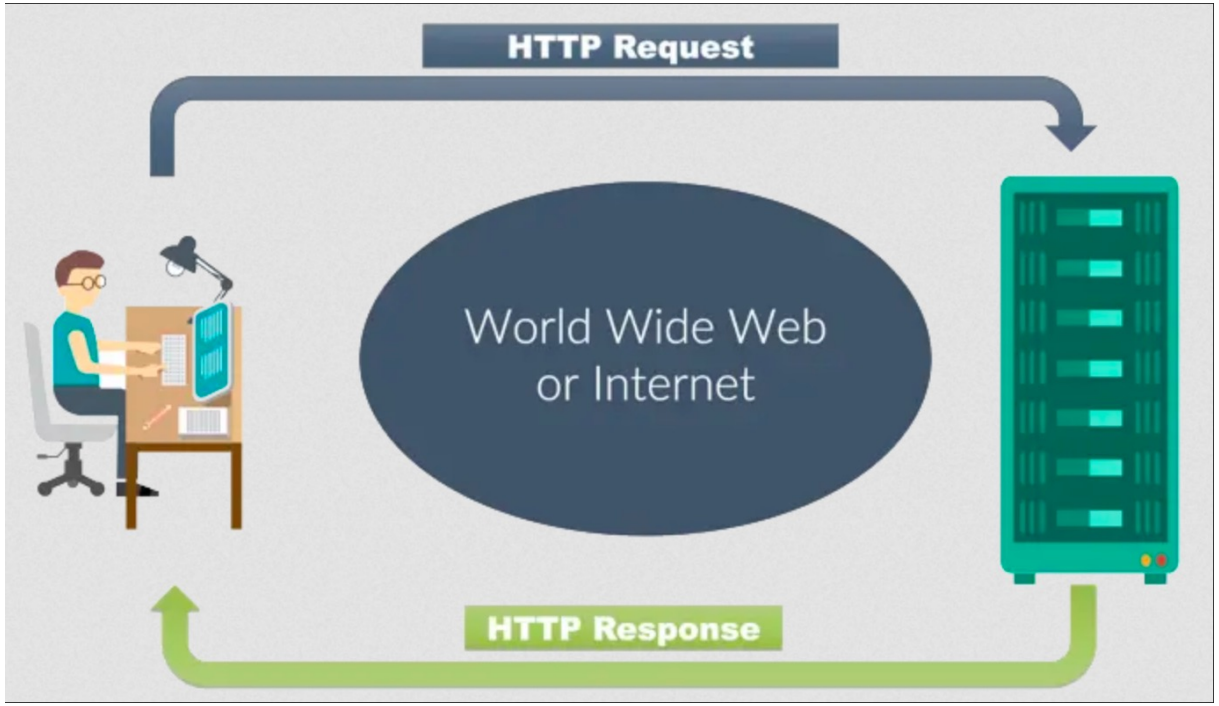

# Service Examples

- Orion
- Salesforce.com
- Office365
- Box
- Google Apps
- Docusign

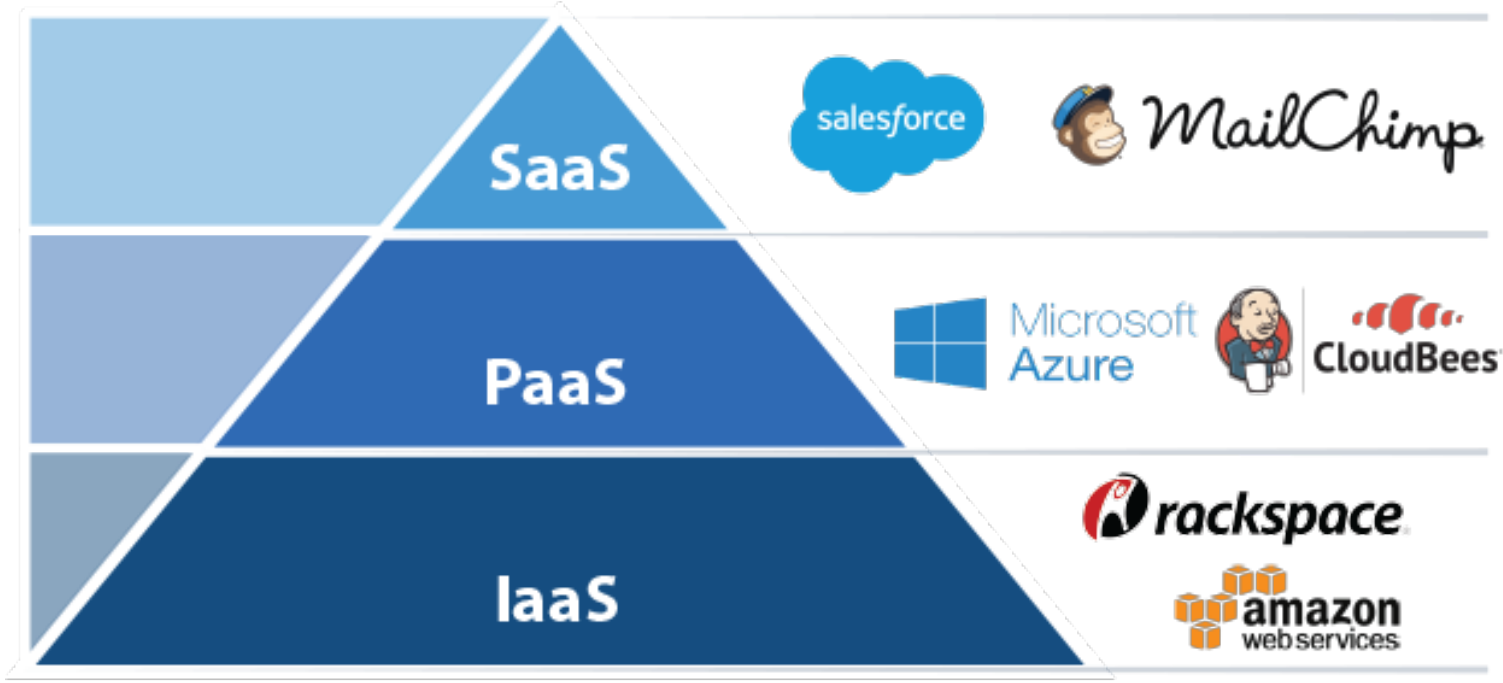

# Using a service: like using a switchboard

1. Ask the operator to connect you to your bank => (TCP, DNS) connect to the service

2. Ask the bank for your current balance => send a **request** to the service

3. The bank tells you your balance => a **response** is received

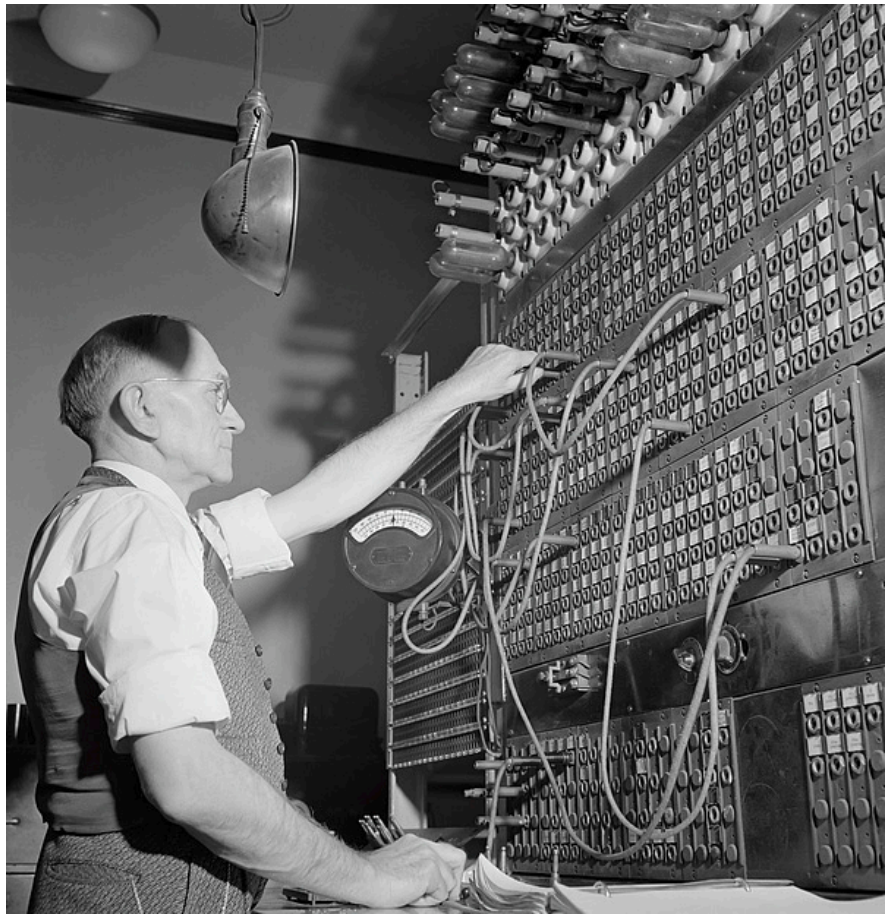

# Orion: Client & Server Model

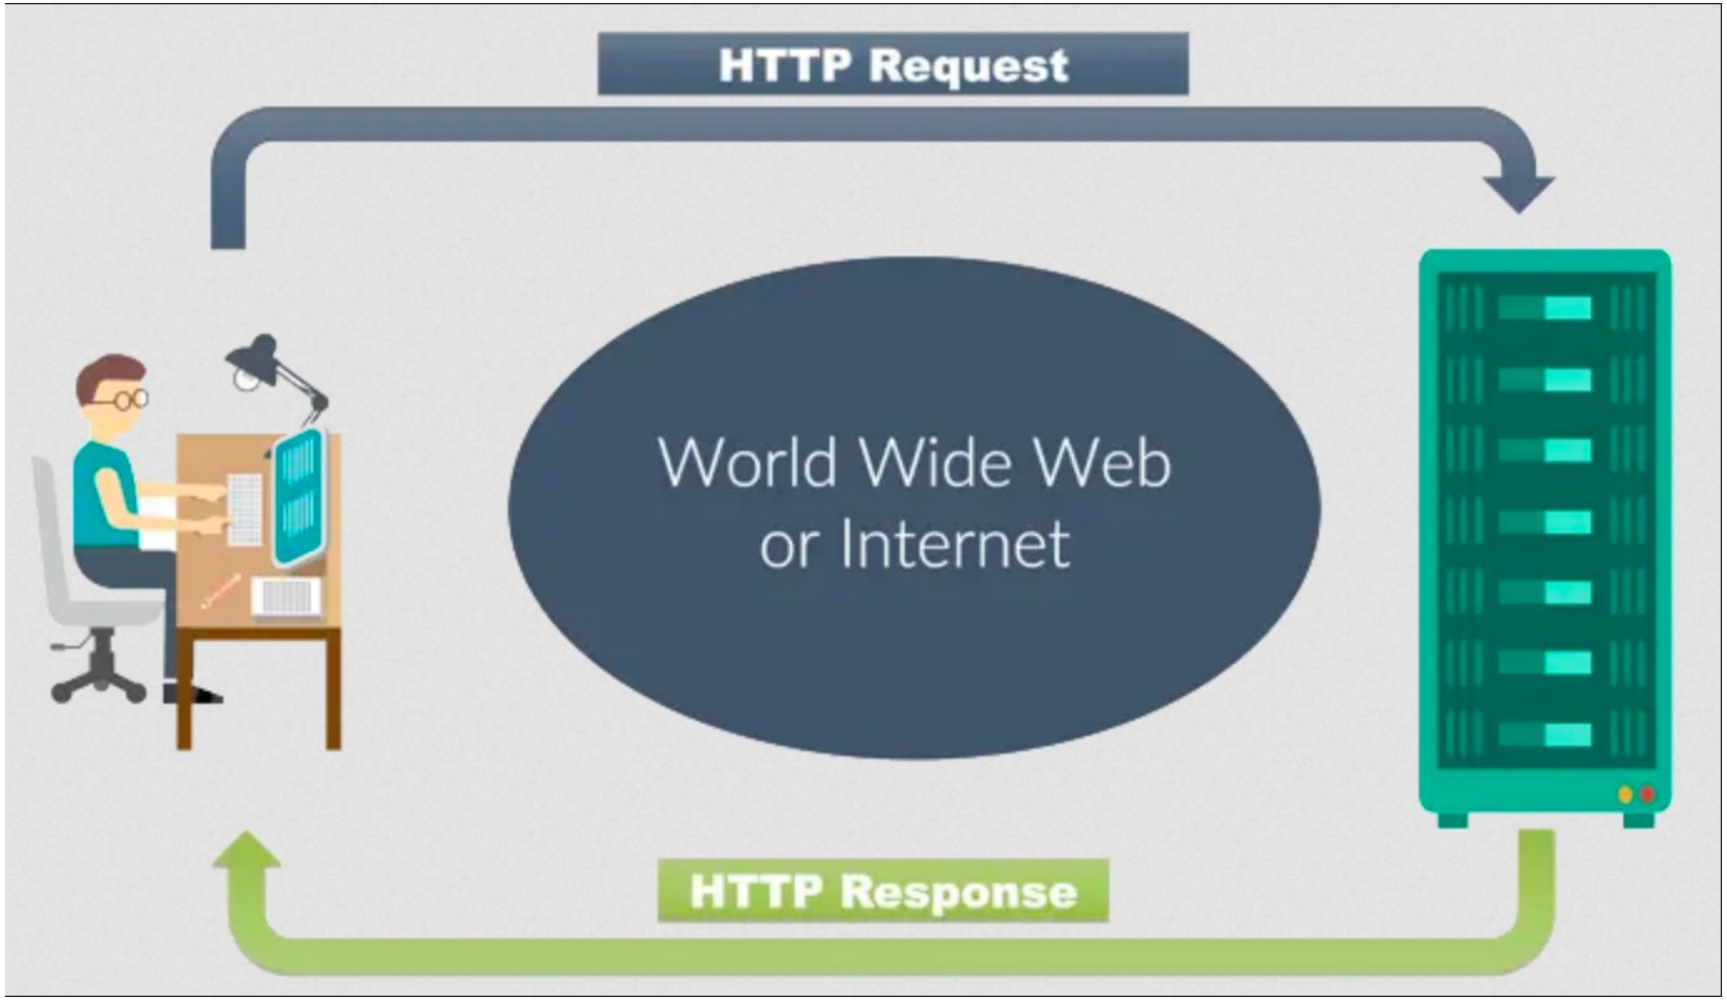

# Orion: Client & Server Model

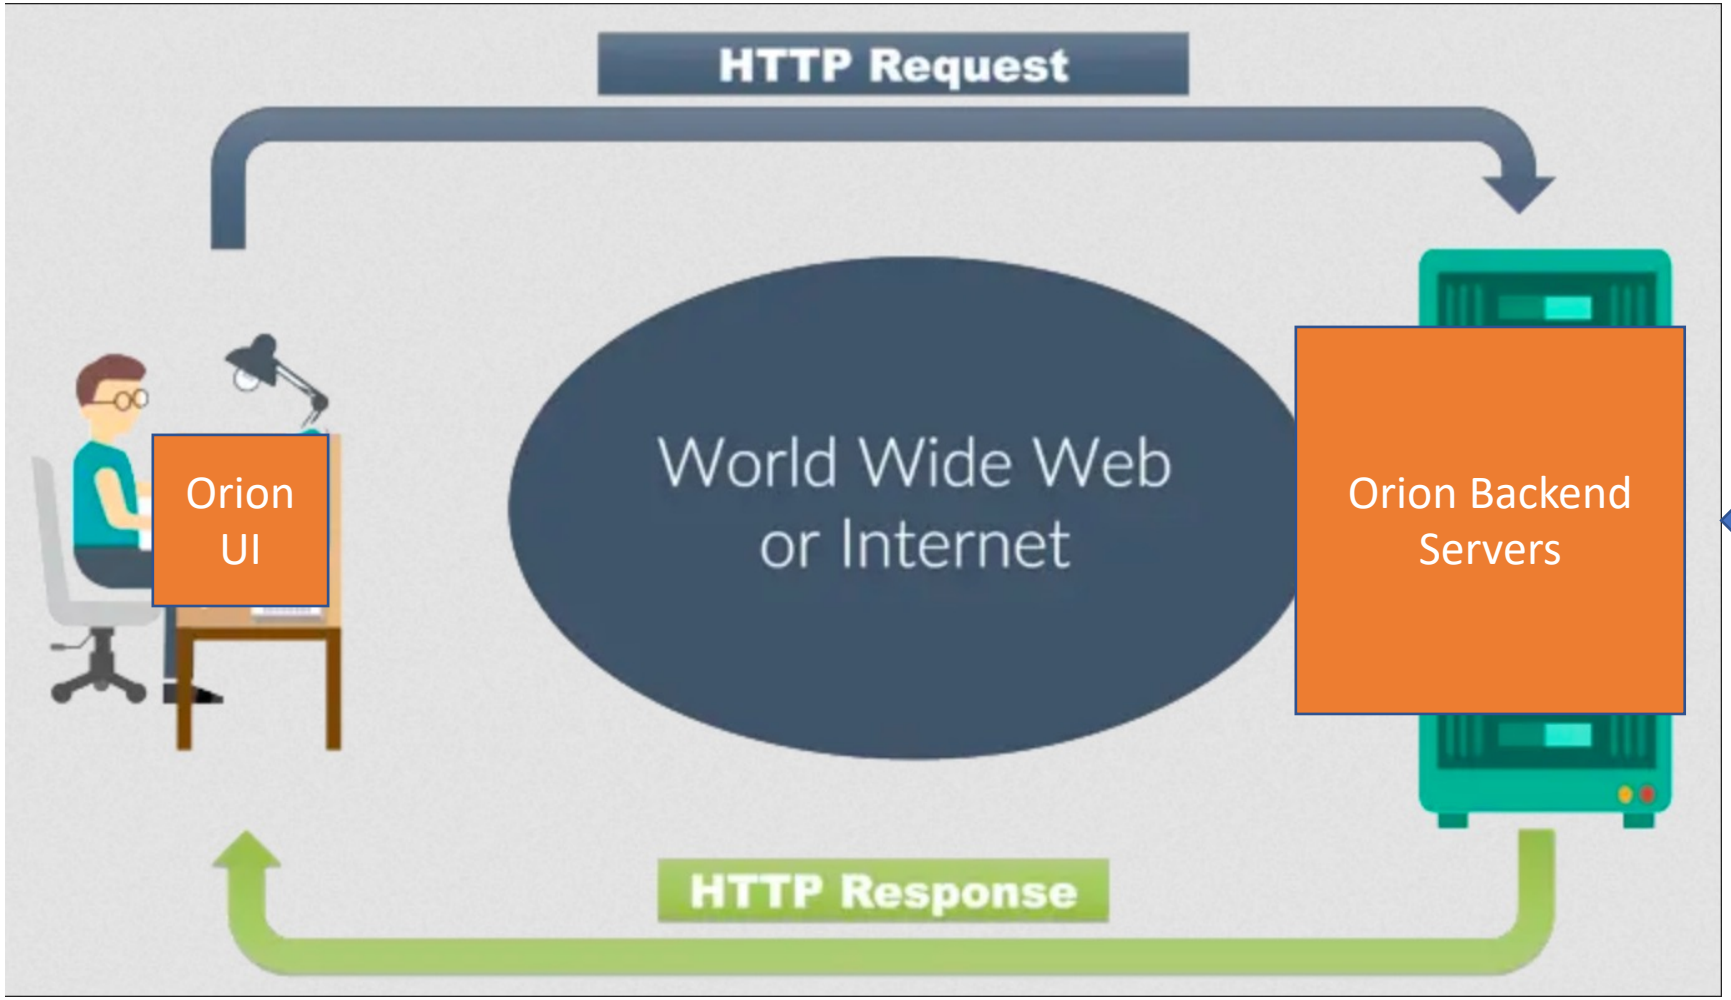


Amazon Web
Services

# Service Design

- Services typically encapsulate a set of operations on a related set of concepts
- It is “black box” for consumers as they are not aware of the service’s inner workings.
- APIs: used to describe a related set of operations
  - Ex: Orion Dataset APIs

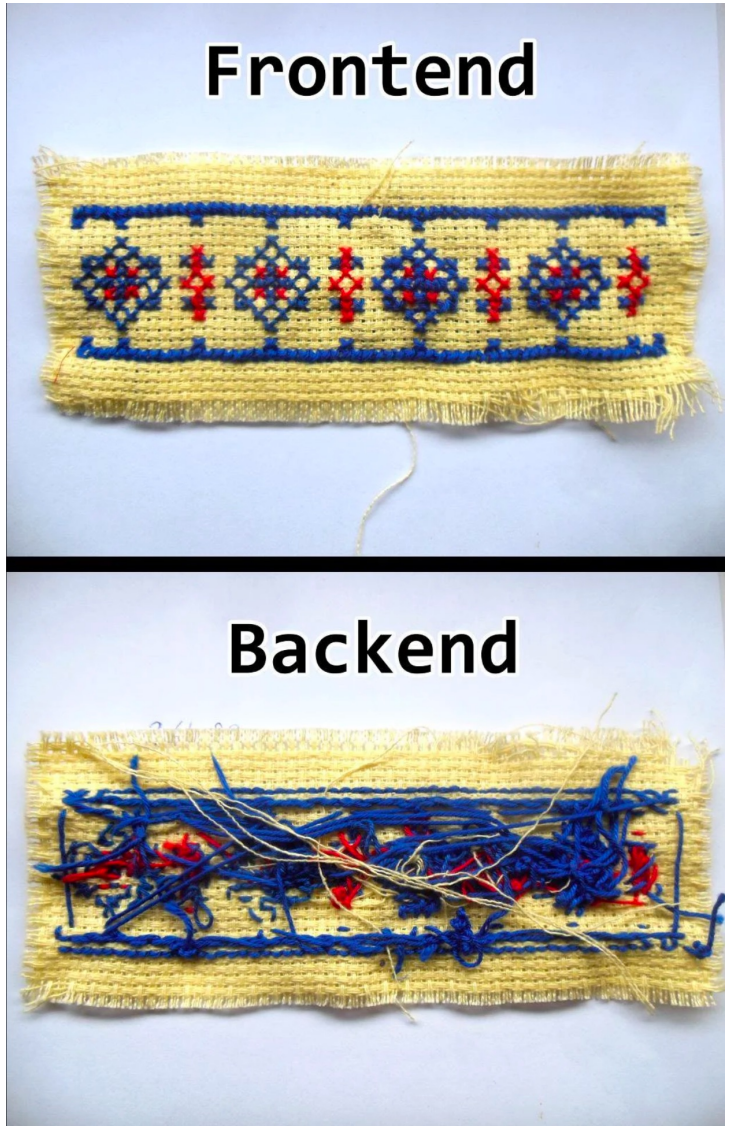

# Founding of AWS

*”All teams will henceforth expose their data and functionality through **service interfaces**…*

*…the team must plan and design to be able to **expose** the interface to developers in the outside world. No exceptions.*

*Anyone who doesn’t do this will be fired.*

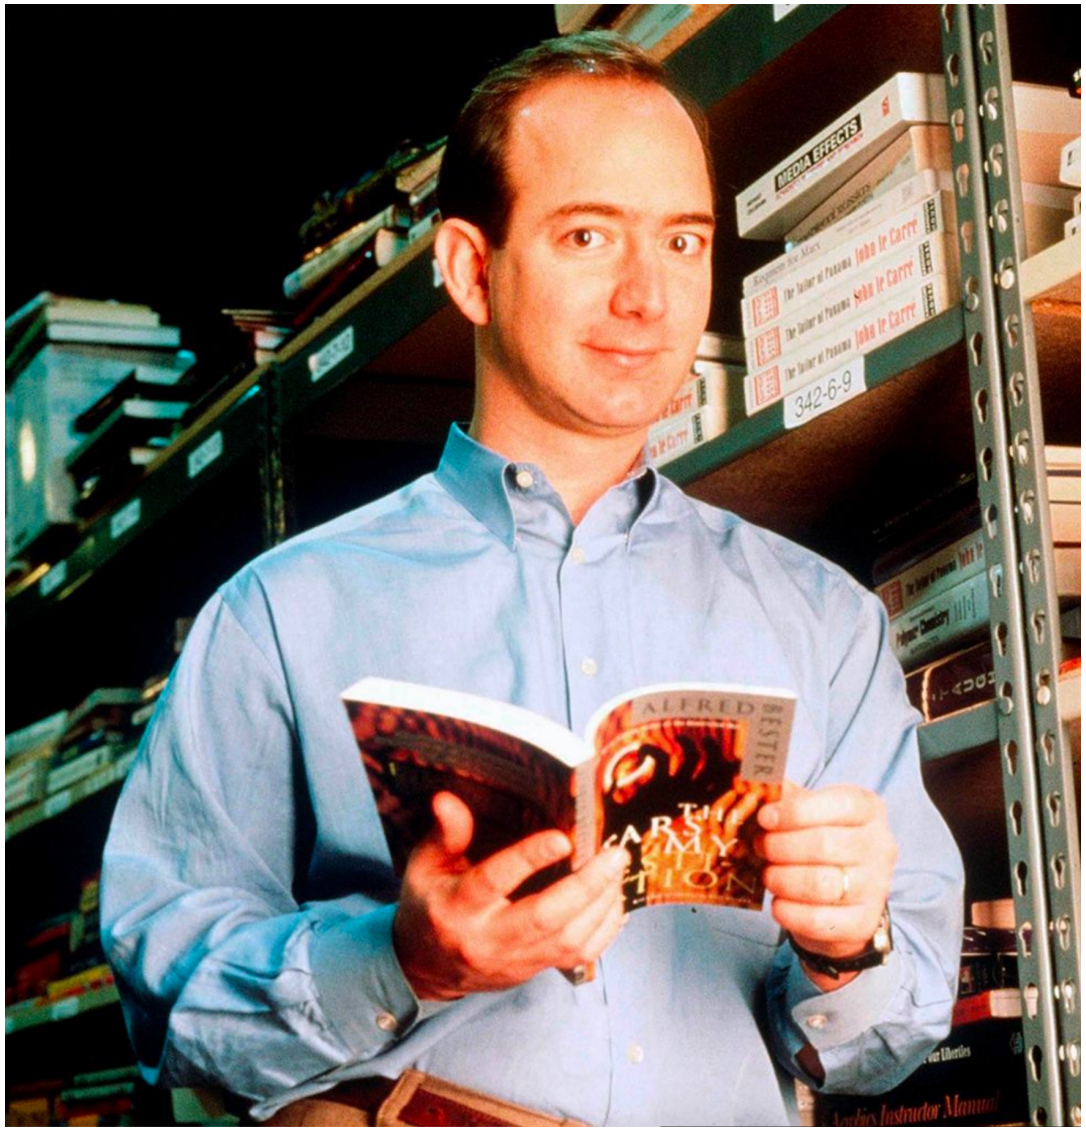


OpenEye SCIENTIFIC

# Founding of AWS

*”All teams will henceforth expose their data and functionality through **service interfaces**…*

*…the team must plan and design to be able to **expose** the interface to developers in the outside world. No exceptions.*

*Anyone who doesn’t do this will be fired.*

# Orion High Level Architecture

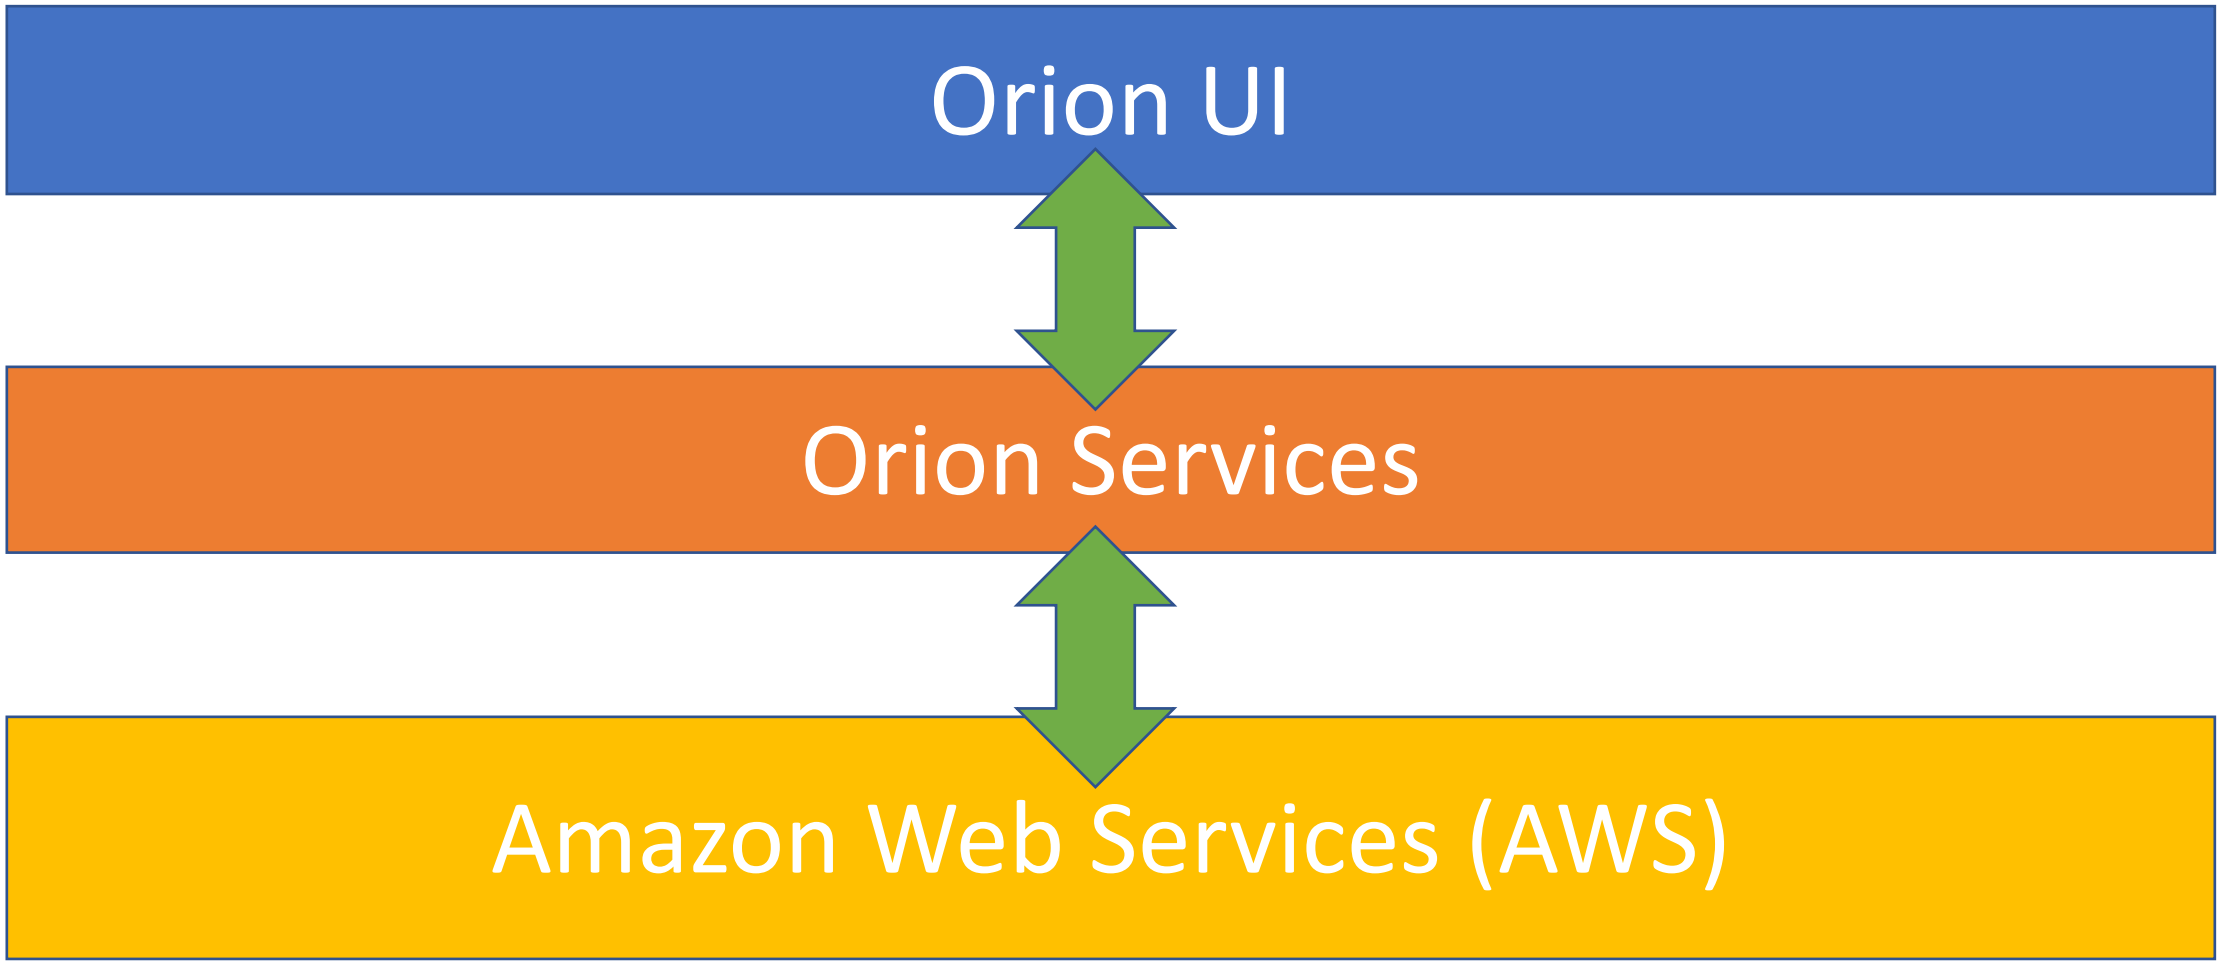

# Service: Write vs. Buy

- The choice is not always clear
  - “We should only use a service when it brings tremendous value…” – Jack D.
- Example, when to buy: Elastic Compute Cloud (EC2) service
  - By using EC2 we benefit from AWS economy of scale for security, reliability, etc.
- Example, when to write: MaaS, FastROCS, MMDS
- Orion Scheduler:
  - No AWS service for fair share, cost aware scheduling, cyclic workflows, etc.

# What AWS Services does Orion use?

| AWS Service | Purpose |
| --- | --- |
| EC2 (elastic compute cloud) | Web servers, workers, etc. |
| S3 (simple storage service) | Persistent storage of Orion data: files, collections, backups, … |
| Aurora (Postgres) Database | Any data with a relationship: users, accounts, permissions, datasets, jobs, accounting, … |
| Elasticache (Redis) Database | Caching, push notifications |
| Cloudwatch | Logging, Metrics |
| Cloudformation | Deployment & Orchestration |
| VPC (virtual private cloud) | Networking & isolation |
| IAM (identity & access management) | Permissions, credentials |
| SES (simple email service) | Sending email |
| … | |

# What Services does Orion provide?

- Data Services
  - Reading/writing project data (Files, Datasets, Collections, …)
  - Sharing & Permissions
  - Project management
- Administration
  - Users, roles, org. data, secrets, packages, …
- Jobs
  - Starting, stopping, monitoring

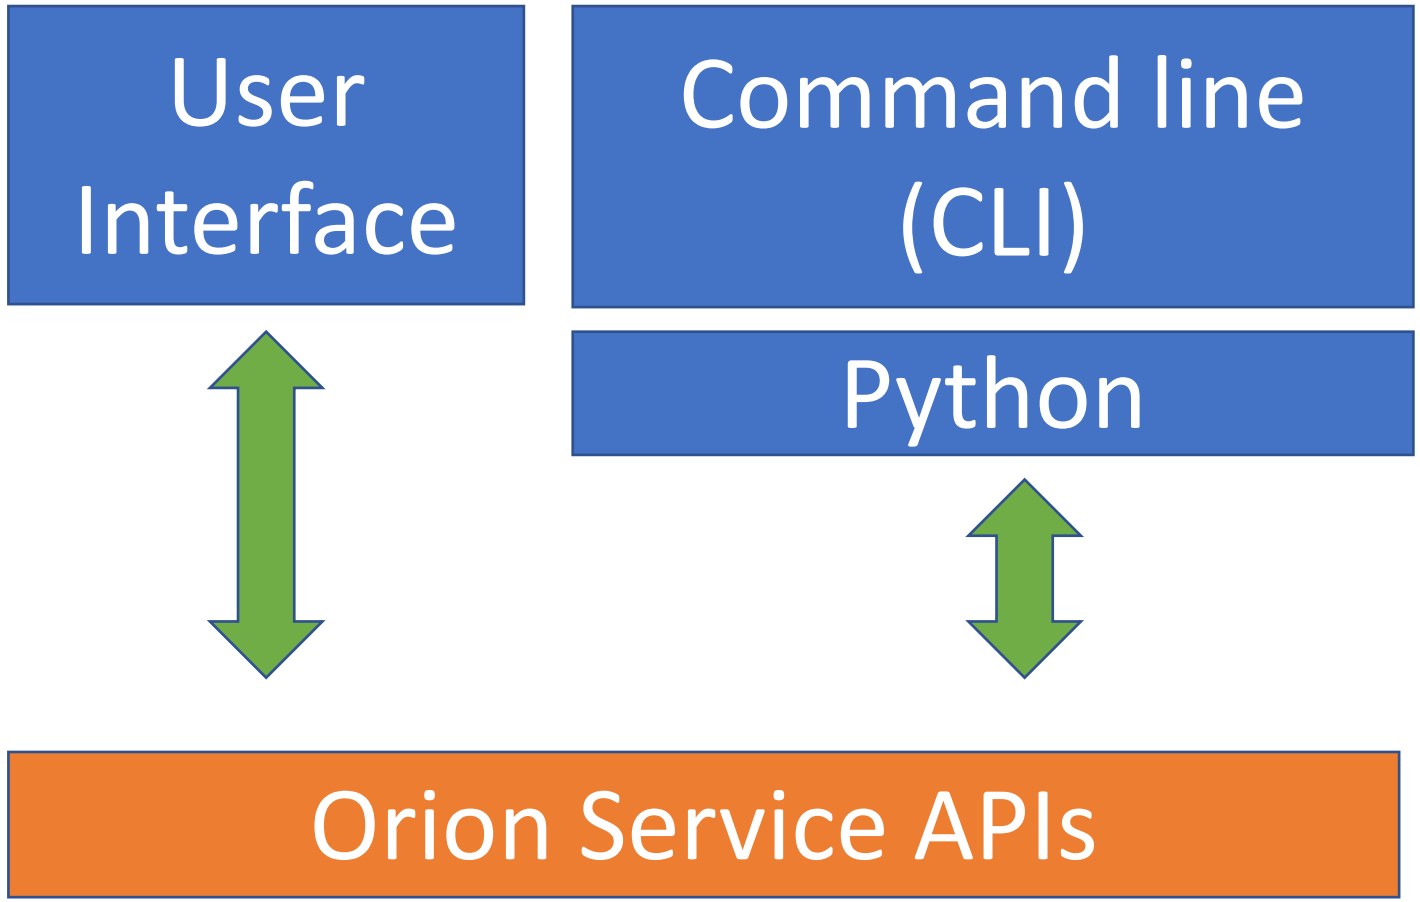

# Risks of using Amazon services

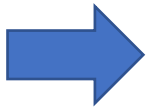

- Service disruptions
  - Can directly impact Orion…and most of the Internet!
- Service behavior changes
  - Can be unannounced, subtle, hard to diagnose
- Service Rate Limiting
  - Can impact Orion's performance

# How *not* to handle service disruptions

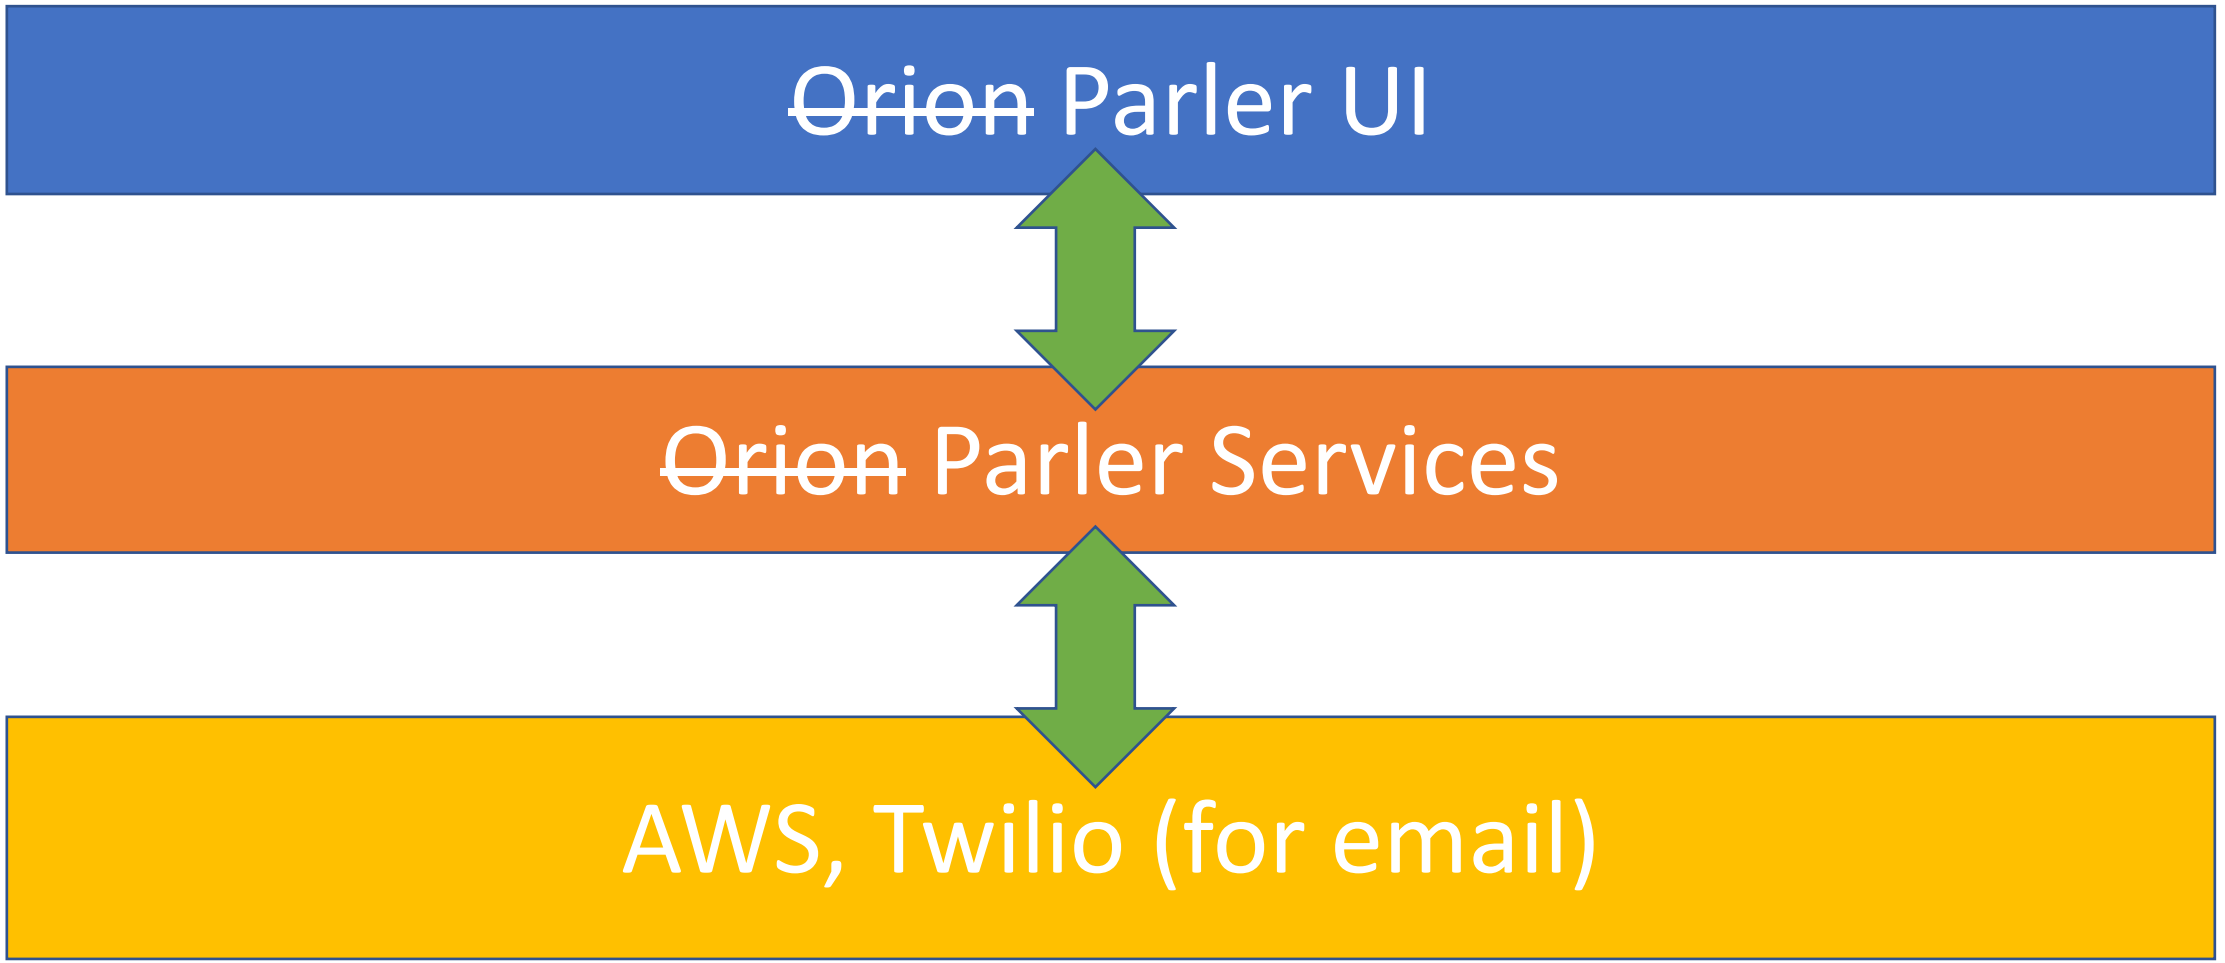
~~Orion~~ Parler UI
~~Orion~~ Parler Services
AWS, Twilio (for email)

# Parler Hack: A Series of Unfortunate Decisions

- Unlimited access to the entirety of Parler content via service API
  - Private messages, deleted posts, …
- Guessing the URLs for content was as trivial as counting
  - Insecure Direct Object Reference (IDOR)
  - Images & Videos were not scrubbed (GPS location, other info)
  - APIs did not require authentication
  - APIs were not rate limited
    - GET /posts/1/
    - GET /posts/2/
    - …
- Result: Most (all?) Parler data was exfiltrated

# Parler: ”Service Disruption"

- The logic for allowing a new user activation was flawed
- New accounts could be created faster than they could be blocked

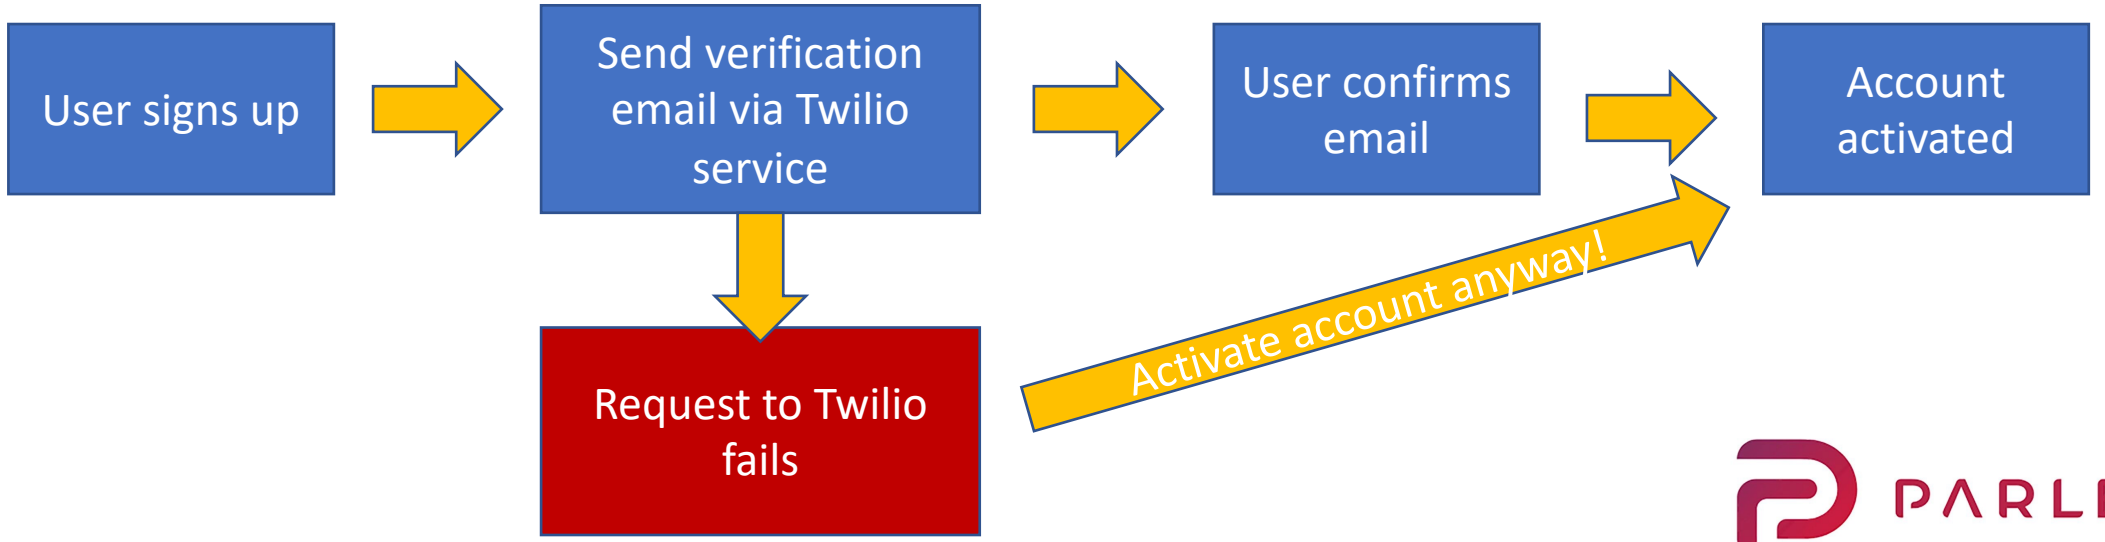


PARLER

# What happens to Orion during an AWS outage?

- Service disruptions may only impact a subset of Orion's functionality
- Ex: if Simple Email Service (SES) is unavailable: Multi-Factor Authentication (MFA) using email in Orion is unavailable

| Service | Purpose |
|---|---|
| EC2 (elastic compute cloud) | Web servers, workers, etc. |
| S3 (simple storage service) | Persistent storage of Orion data: files, collections, backups, … |
| Aurora (Postgres) Database | Any data with a relationship: users, accounts, permissions, datasets, jobs, accounting, … |
| Cloudwatch | Logging, Metrics |
| Cloudformation | Deployment & Orchestration |
| VPC (virtual private cloud) | Networking & isolation |
| IAM (identity & access management) | Permissions, credentials |
| SES (simple email service) | Sending email |
| … | |

# Major AWS Disruptions are Unavoidable

## AWS Post-Event Summaries

The following is a list of post-event summaries from major service events that impacted AWS service availability:

- Summary of the Amazon Kinesis Event in the Northern Virginia (US-EAST-1) Region, November, 25th 2020
- Summary of the Amazon EC2 and Amazon EBS Service Event in the Tokyo (AP-NORTHEAST-1) Region, August 23, 2019
- Summary of the Amazon EC2 DNS Resolution Issues in the Asia Pacific (Seoul) Region (AP-NORTHEAST-2), November 24, 2018.
- Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region, February 28, 2017.
- Summary of the Amazon DynamoDB Service Disruption and Related Impacts in the US-East Region, September 20, 2015.
- Summary of the Amazon EC2, Amazon EBS, and Amazon RDS Service Event in the EU West Region, August 7, 2014.
- Summary of the Amazon SimpleDB Service Disruption, June 13, 2014.
- Summary of the December 17th event in the South America Region (SA-EAST-1), December 20, 2013.
- Summary of the December 24, 2012 Amazon ELB Service Event in the US-East Region, December 24, 2012.
- Summary of the October 22, 2012 AWS Service Event in the US-East Region, October 22, 2012.
- Summary of the AWS Service Event in the US East Region, July 2, 2012.
- Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region, April 29, 2011.
- Summary of the AWS Service Event in the Sydney Region, June 8, 2011.

# November 25th 2020 AWS Outage

TECH | AMAZON

## Prolonged AWS outage takes down a big chunk of the internet

*AWS has been experiencing an outage for hours*

By Jay Peters | @jaypeters | Updated Nov 25, 2020, 5:39pm EST

Outage impacted 1Password, Adobe, Autodesk, Glassdoor, Pocket, Roku, Orion, …

Mitigation is possible, but very difficult

# Risks of using Amazon services

- Service disruptions
  - Can directly impact Orion…and most of the Internet!
- Service behavior changes
  - Can be unannounced, subtle, hard to diagnose
- Service Rate Limiting
  - Can impact Orion's performance

# 2020 AWS Cloudwatch Issue

- Nov. 25th AWS outage was for a service which powers Cloudwatch inside AWS
- Orion uses Cloudwatch for logs, including job logs
- Some time after Nov. 25th the behavior of Cloudwatch changes in a way that caused requests for job logs to timeout
- Jobs in Orion still ran, logs were still written, but couldn’t be read

# Risks of using Amazon services

- Service disruptions
  - Can directly impact Orion…and most of the Internet!
- Service behavior changes
  - Can be unannounced, subtle, hard to diagnose
- Service Rate Limiting
  - Can impact Orion's performance

# Rate limiting: Switchboard's busy signal

- Enforced limit on request rate to a service
- Ensure stability of a service & fulfill SLA
- Helps in capacity planning where a service's physical limitations are known
- Helps prevent abuse, Denial of Service (DoS) attacks
- ***Every single Amazon service is rate limited***
  - ***Because it's critical to their business!***

# Rate Limiting 101: Token Bucket

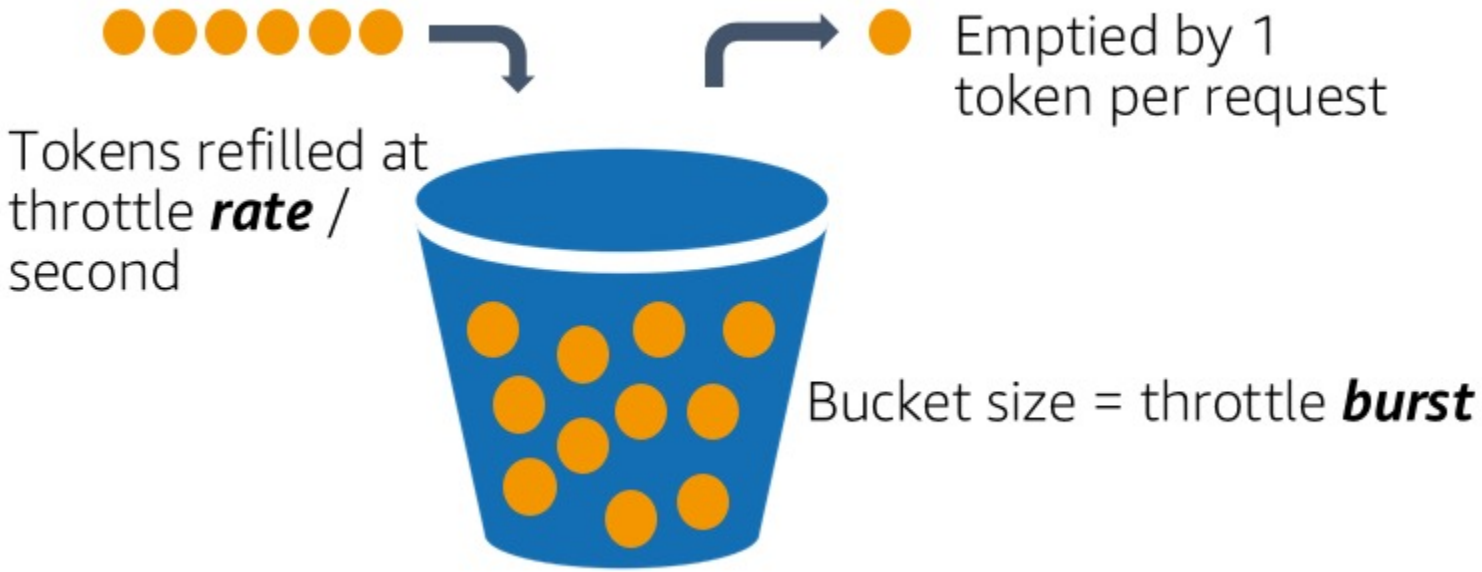


If you send a request and your bucket is empty, your request is denied (throttled)

# What happens when Orion is rate limited?

- If a synchronous request is throttled: return error
- If an asynchronous request is throttled: wait & try again later
  - Wait for how long?
    - set a timer for 1s; call again
    - set a timer for 2s; call again
    - …repeat until you get through
  - Binary exponential backoff (BEB)

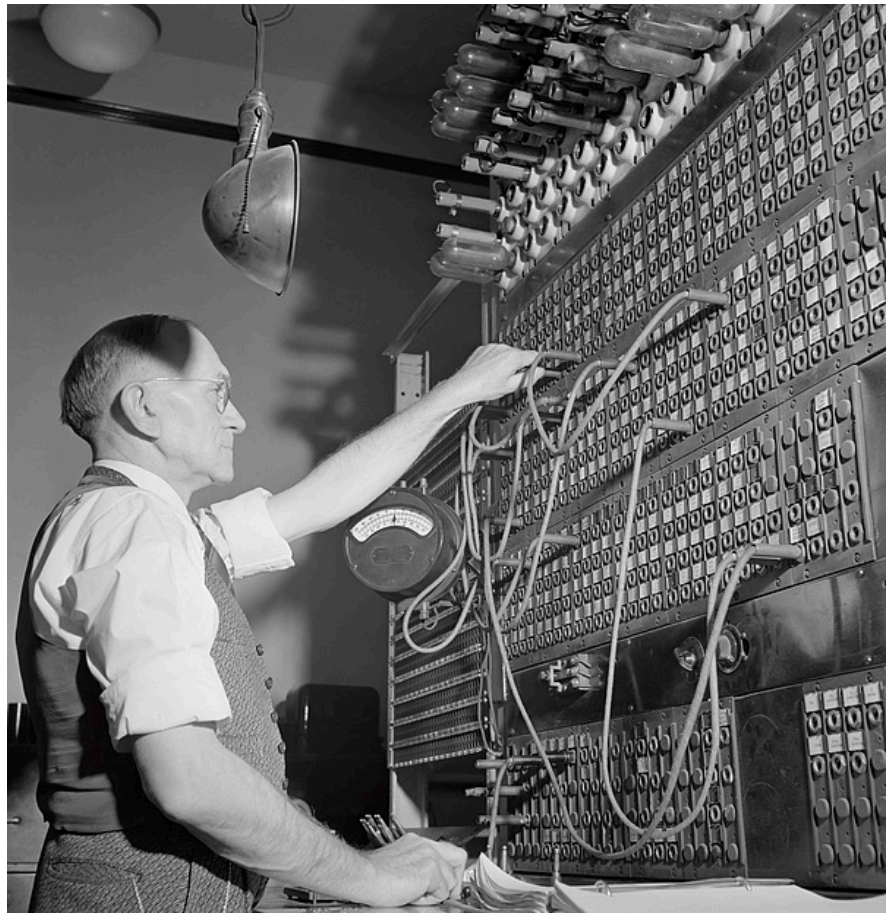

# Naïve backoff is not enough

## Uniform Backoff

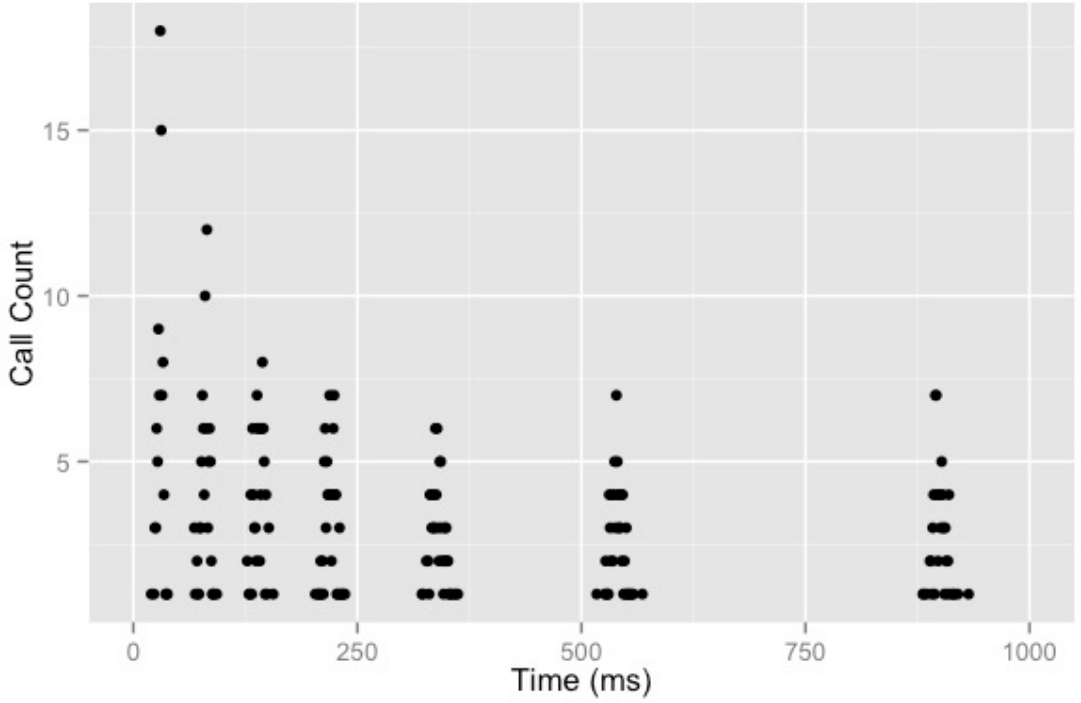


## Jittered Backoff

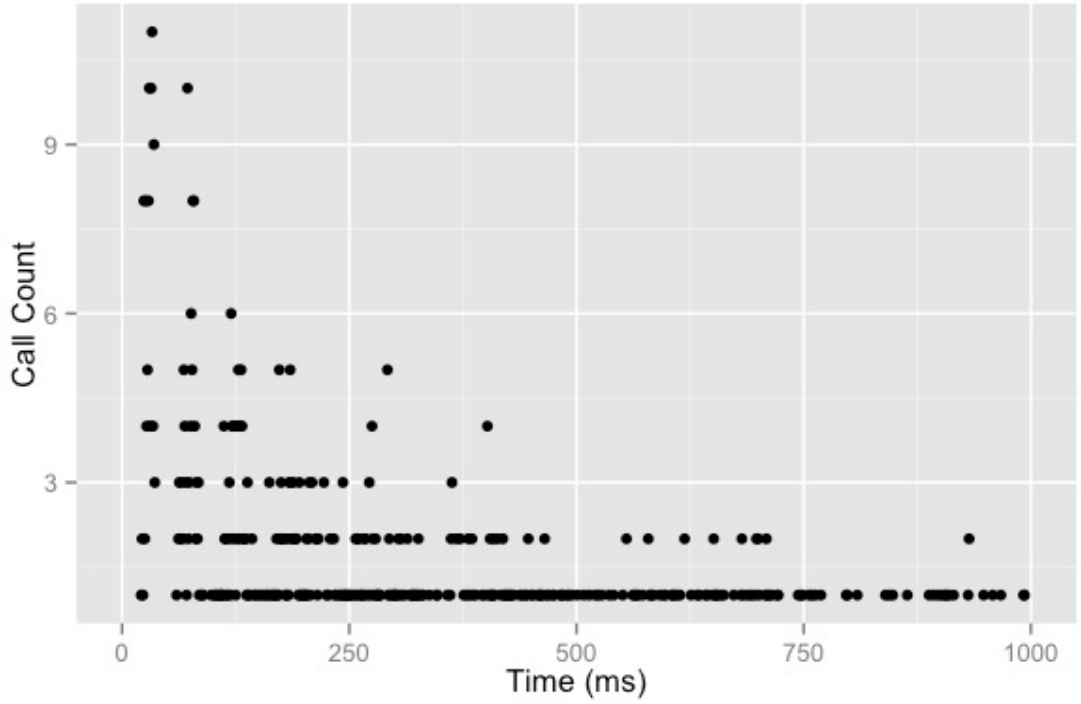


OpenEye SCIENTIFIC

# What happens without backoff?

- Each caller competes with N-1 others in the first round, N-2 in the second, …
- Without backoff this leads N^2 attempts!

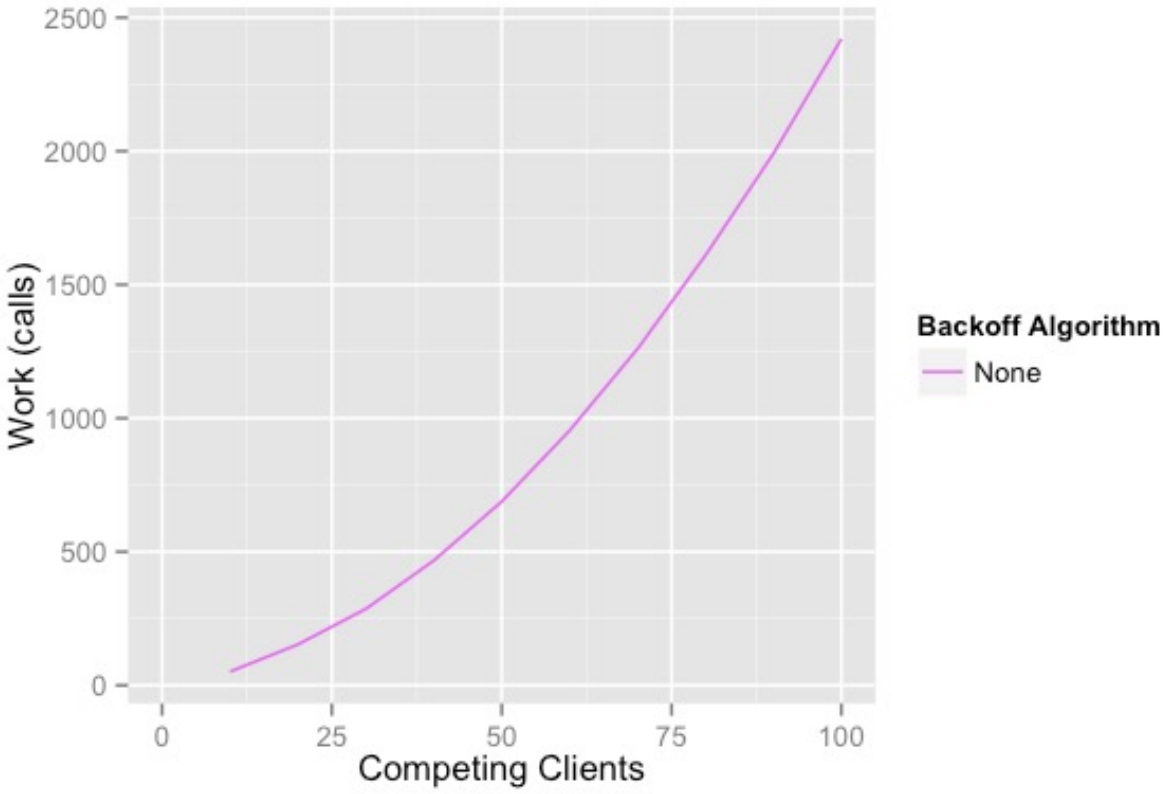

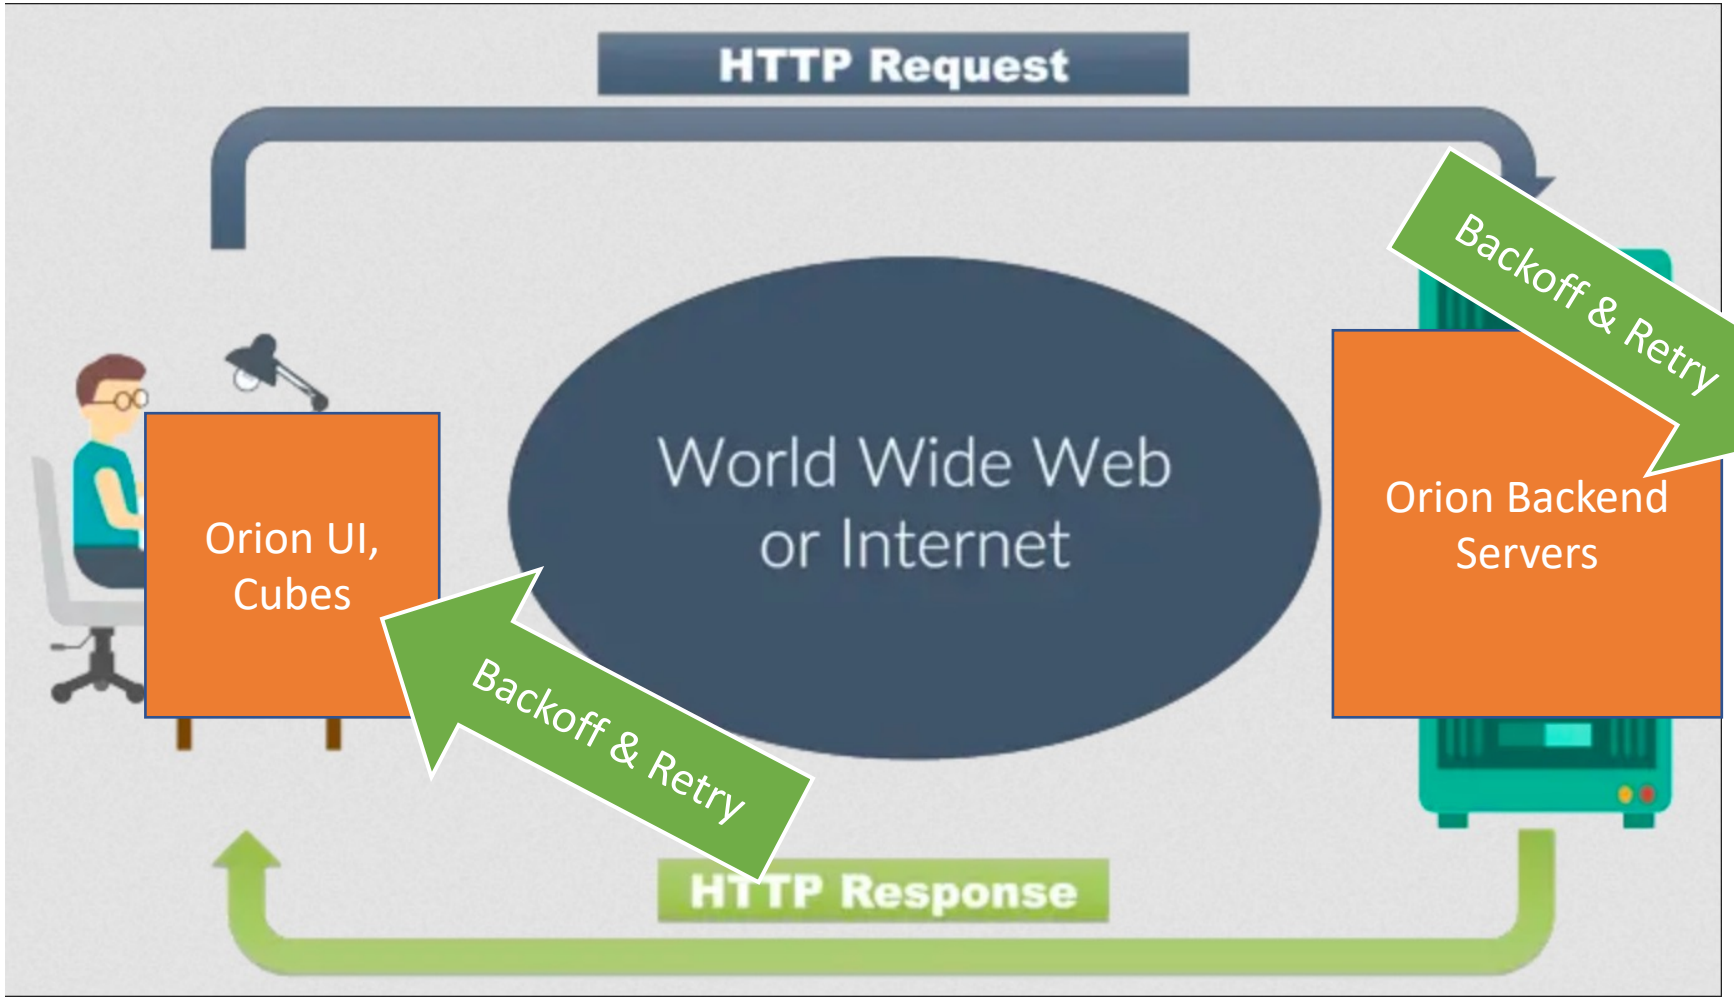
HTTP Request
World Wide Web
or Internet
Orion UI,
Cubes
Backoff & Retry
Backoff & Retry
HTTP Response
Orion Backend
Servers
Amazon Web
Services

# Summary

- A better understanding of
  - What a service is
  - How Orion is composed of services
  - What Amazon services Orion uses

- How Orion responds to service disruptions

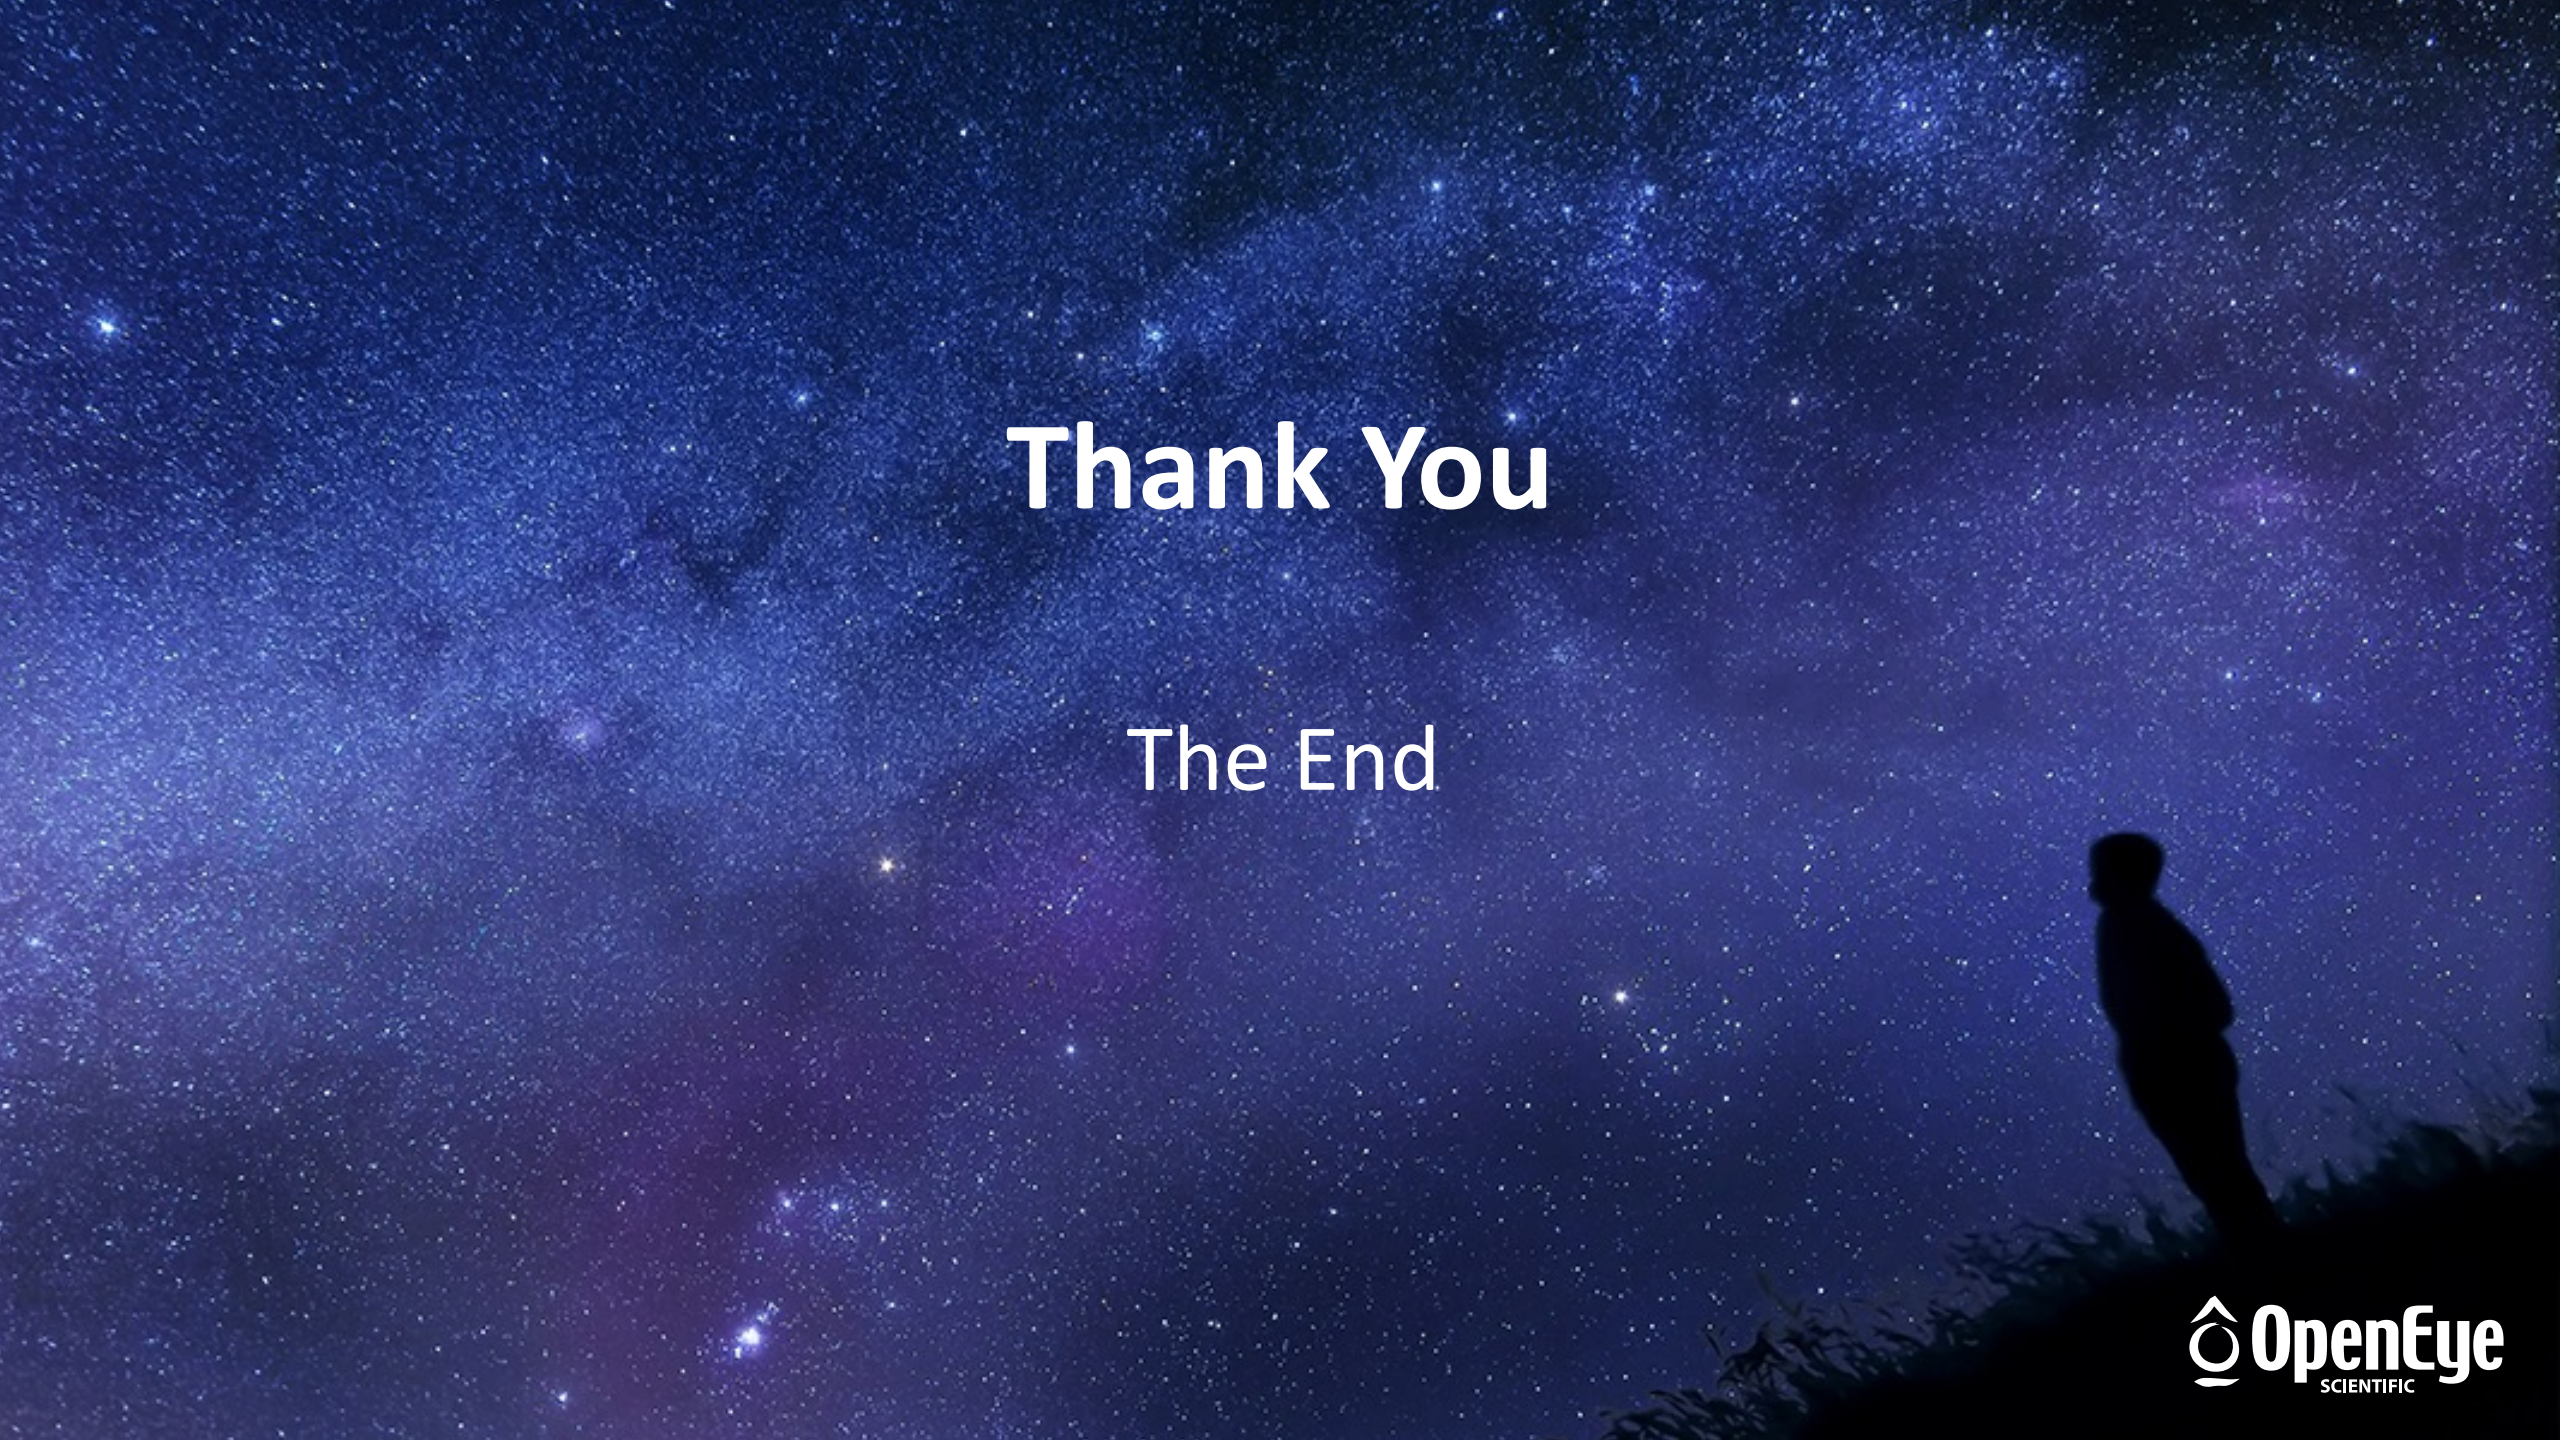
Thank You
The End
OpenEye
SCIENTIFIC

For more information, please contact:

**sales@eyesopen.com** | **info@eyesopen.com**

**www.eyesopen.com**

**+1-505-473-7385**

OpenEye SCIENTIFIC