AWS re:Invent

JPL: Auto Scaling Advantage

Advanced Rapid Imaging and Analysis

Hook Hua[1], Gerald Manipon[1], Michael Starch[1], Lan Dang[1], Justin Linick[1], Sang-Ho Yun[1], Gian Franco Sacco[1], Piyush Agram[1], Susan Owen[1], Brian Bue[1],, Eric Fielding[1], Paul Lundgren[1], Angelyn Moore[1], Paul Rosen[1], Zhen Liu[1], Eric Gurrola[1], Tom Farr[1], Vincent Realmuto[1], Frank Webb[1], Mark Simons[2], Pietro Mililo[1]

[1] Jet Propulsion Laboratory

[2] California Institute of Technology

Thursday, December 1, 2016

© 2016 California Institute of Technology. Government sponsorship acknowledged.

Reference herein to any specific commercial product, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

© 2016, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

amazon web services

Small Number of Ground Observations Calibrate a Wide Swath Flood Map Derived from Satellite SAR Data

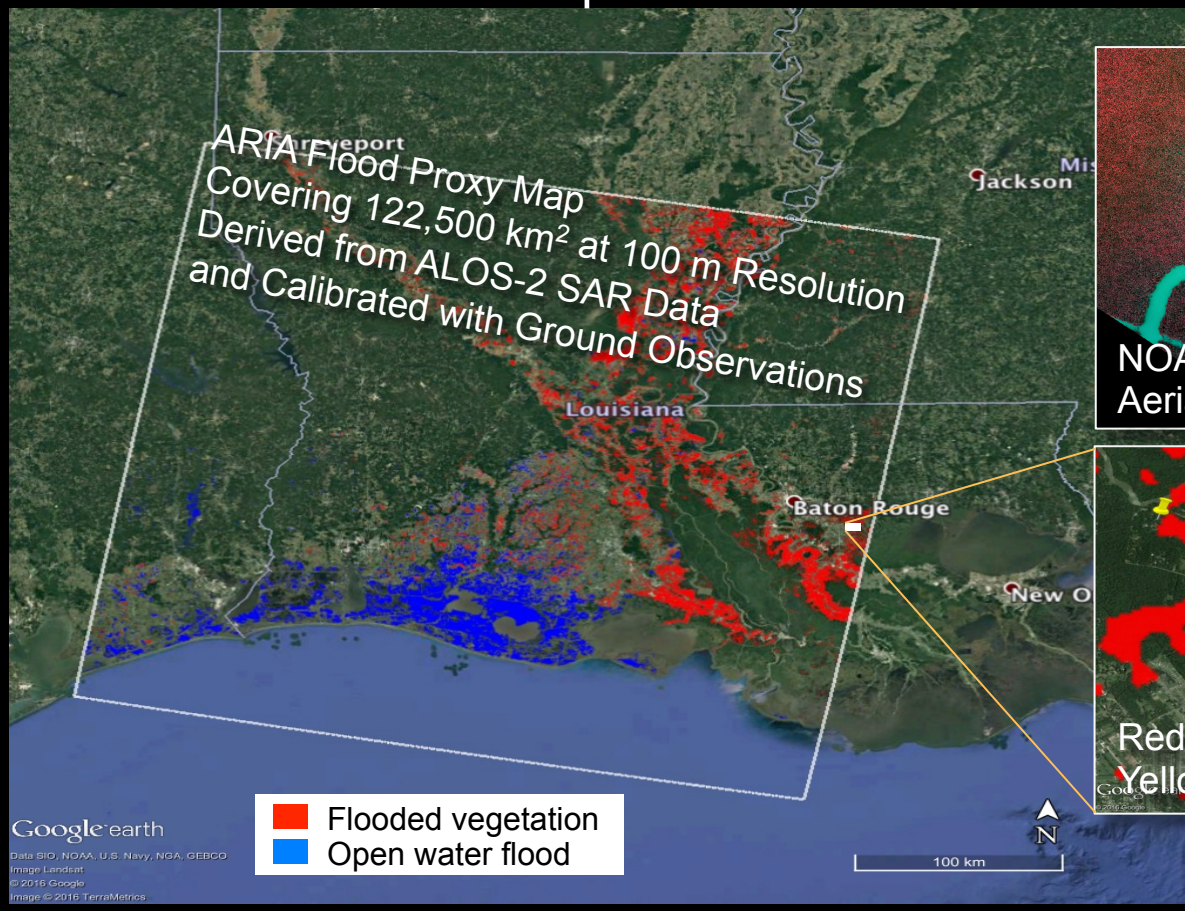

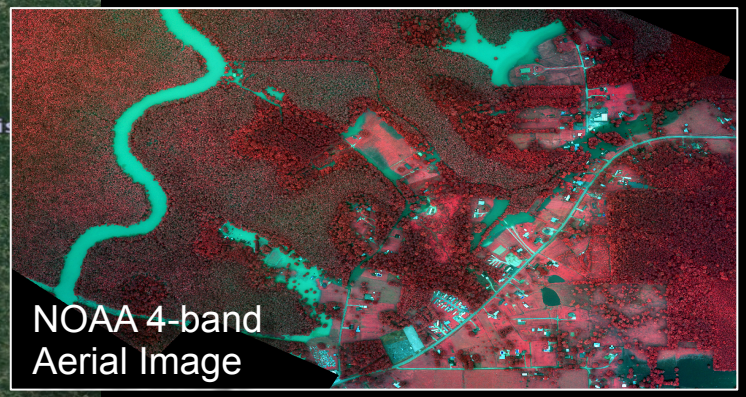

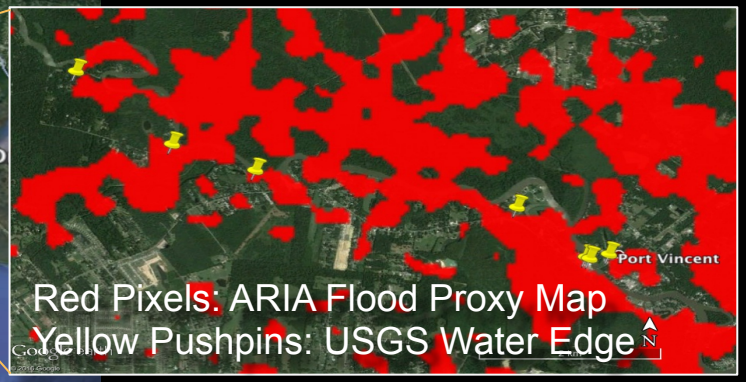

The Challenge of Leveraging Remote Sensing for Disaster Response

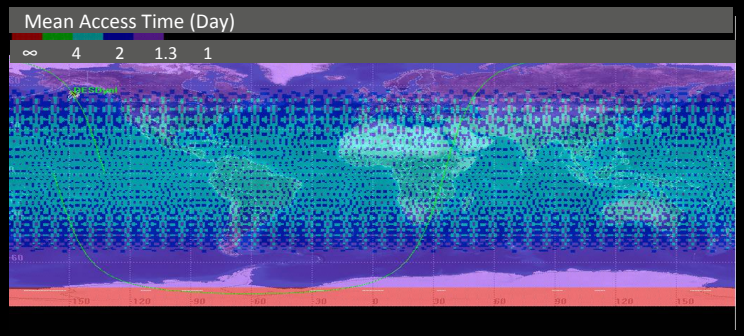

Automated data system are required to analyze large quantities of data from NASA NISAR, other satellite missions, and rapidly expanding GPS networks.

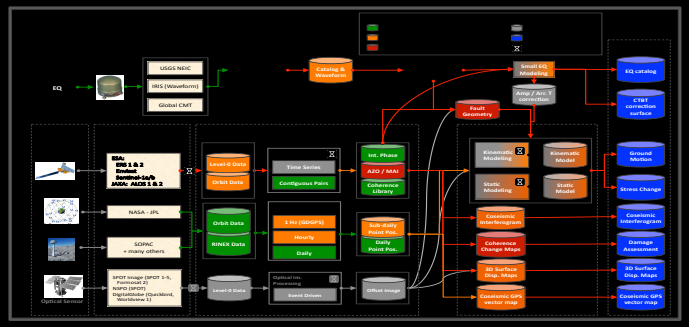

Going from Artisan to Automation: Use system engineering approach to translate specialized data analysis into operational capability.

Demonstrate response to hazards with standardized set of data products for decision & policy makers.

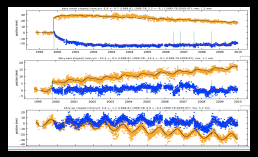

Temporal Records of Ground Deformation

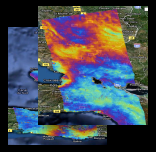

Spatial Maps of Ground Deformation

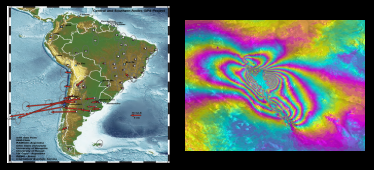

Coseismic Ground Deformation

Coseismic Damage

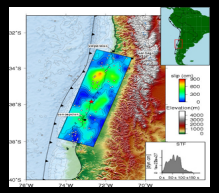

Earthquake Models

Advanced Rapid Imaging and Analysis (ARIA)

"We have high hopes for ARIA."

Keiko Saito
World Bank

JPL

Integrating
Space Geodesy
Seismology
Modeling

Seismological Laboratory
California Institute of Technology

"This is exactly the kind of product we are looking for."

Anne Rosinski
Calif. Geol. Survey

Radar Sensors

GPS Networks

Seismic Networks

Optical Sensors

Automated Data Collection & Processing

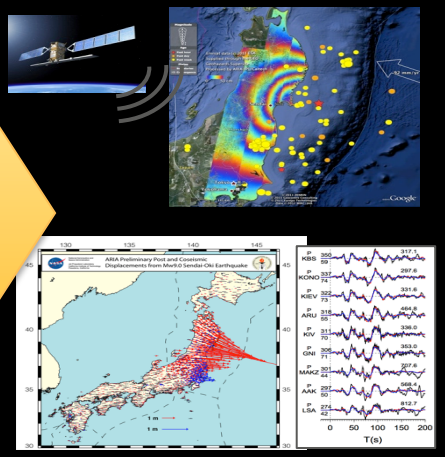

Monitoring & Near Real-Time Assessment

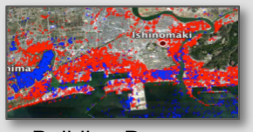

Building Damage & Inundation

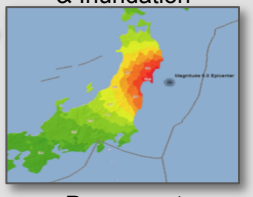

Permanent Ground Deformation

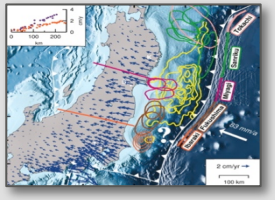

High-Resolution Hazard Assessment from Fault Models

Examples from the 2011 M9.0 Tohoku-Oki (Japan) earthquake

Scalable Processing, Data Management, and Monitoring

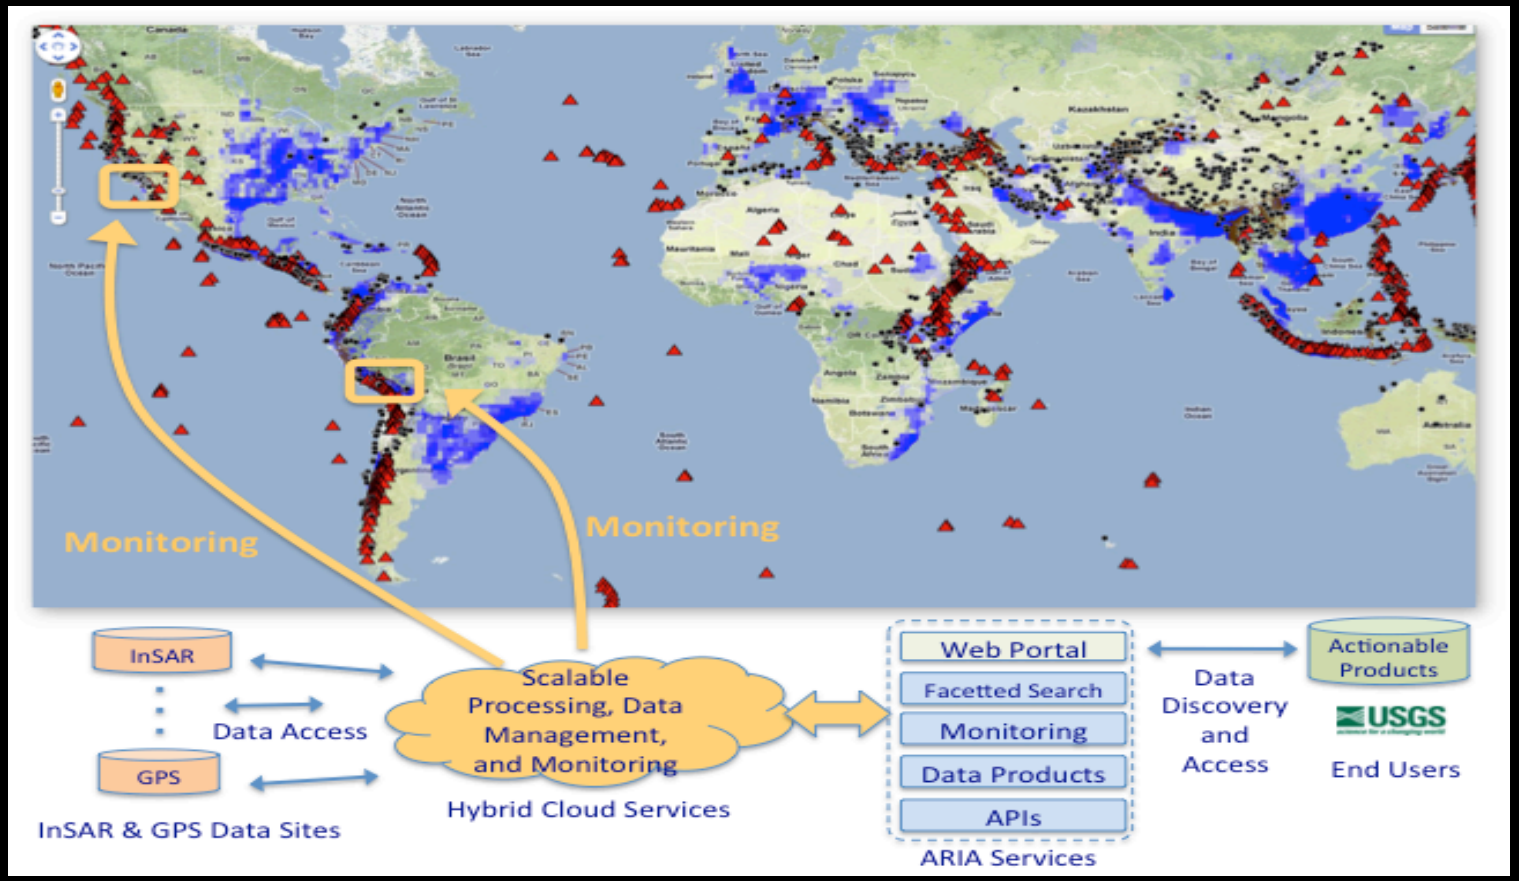

Example Analysis Scenarios

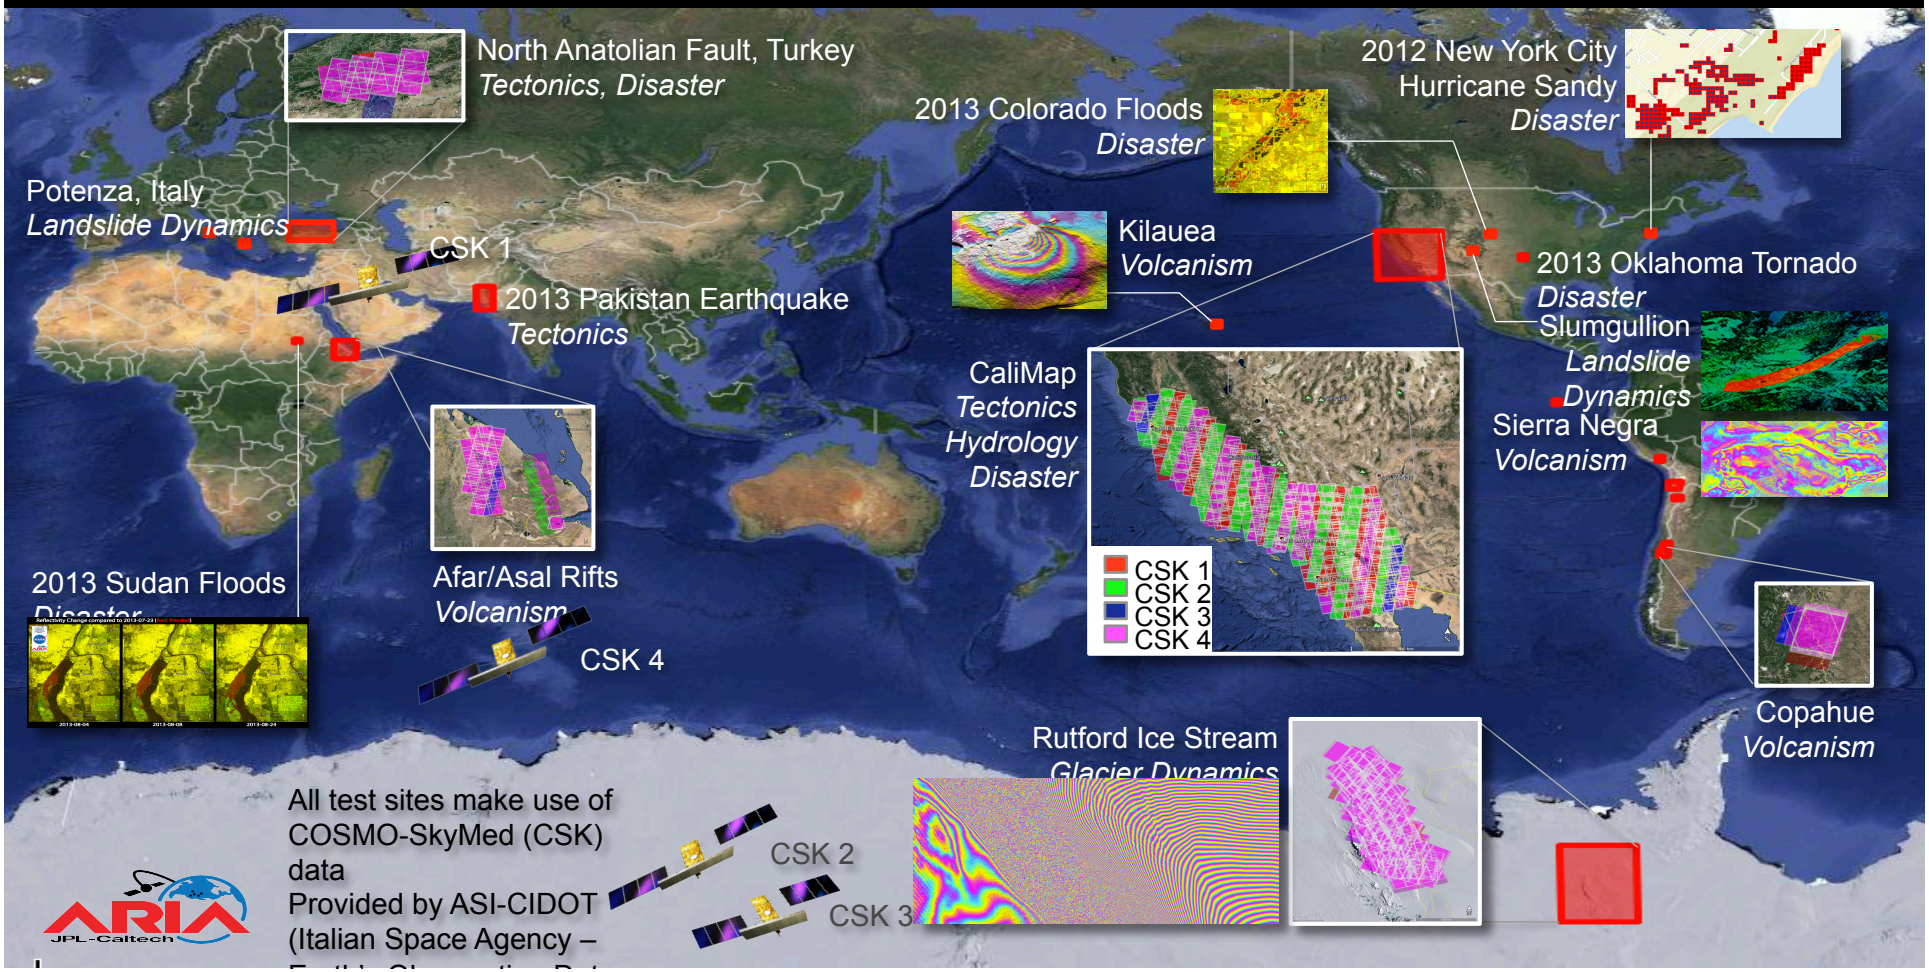

A SAR Science Data Processing Foundry

Growing number of SAR instruments available to scientists

- NASA Instruments: UAVSAR, AirMOSS, EcoSAR
- ASF's ERS/ALOS/RADARSAT/Sentinel-1 archives
- WInSAR's Envisat archive
- ESA Sentinel-1 A/B
- ASI COSMO SkyMed archive

Growing volumes of SAR Instrument data remain unprocessed to higher level science products

SAR data processing requires physics and computing knowledge often too specialized for investigators

Ongoing SAR processing costs include

- On-ramping SAR instrument-specific processing algorithms
- SAR science data processing
- Data storage and movement

The Foundry aims to standardize to the extent possible SAR processing approaches across sensors, algorithms, and computing platforms

- *Minimize cost*
- *Maximize science return*

Sentinel-1A

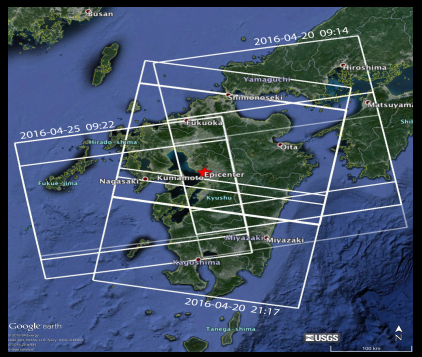

COSMO-SkyMed

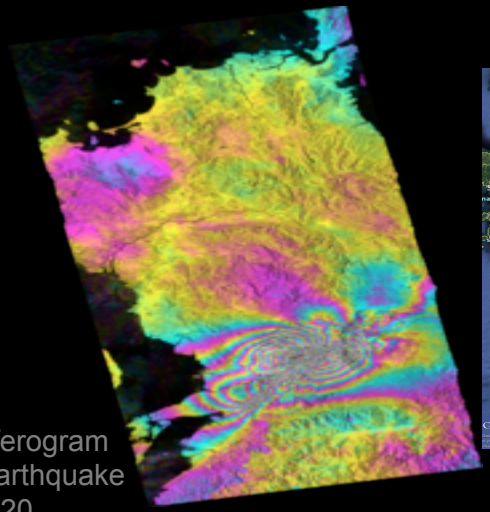

Sentinel-1A interferogram over Kumamoto Earthquake April 4 and 20

ALOS-2

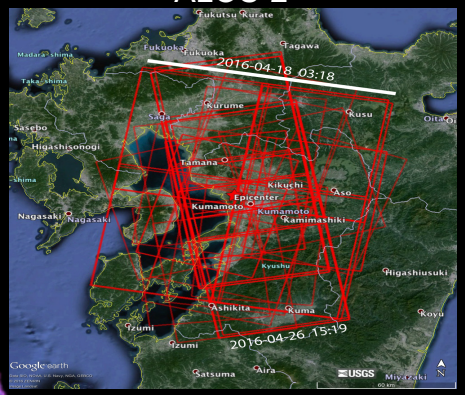

Big Data

HYBRID CLOUD SCIENCE DATA SYSTEM

Process Control & Management in AWS

- Sample pipelined PGE

- Each major PGE processing step is deployed in its own AWS auto-scaling group (ASG)

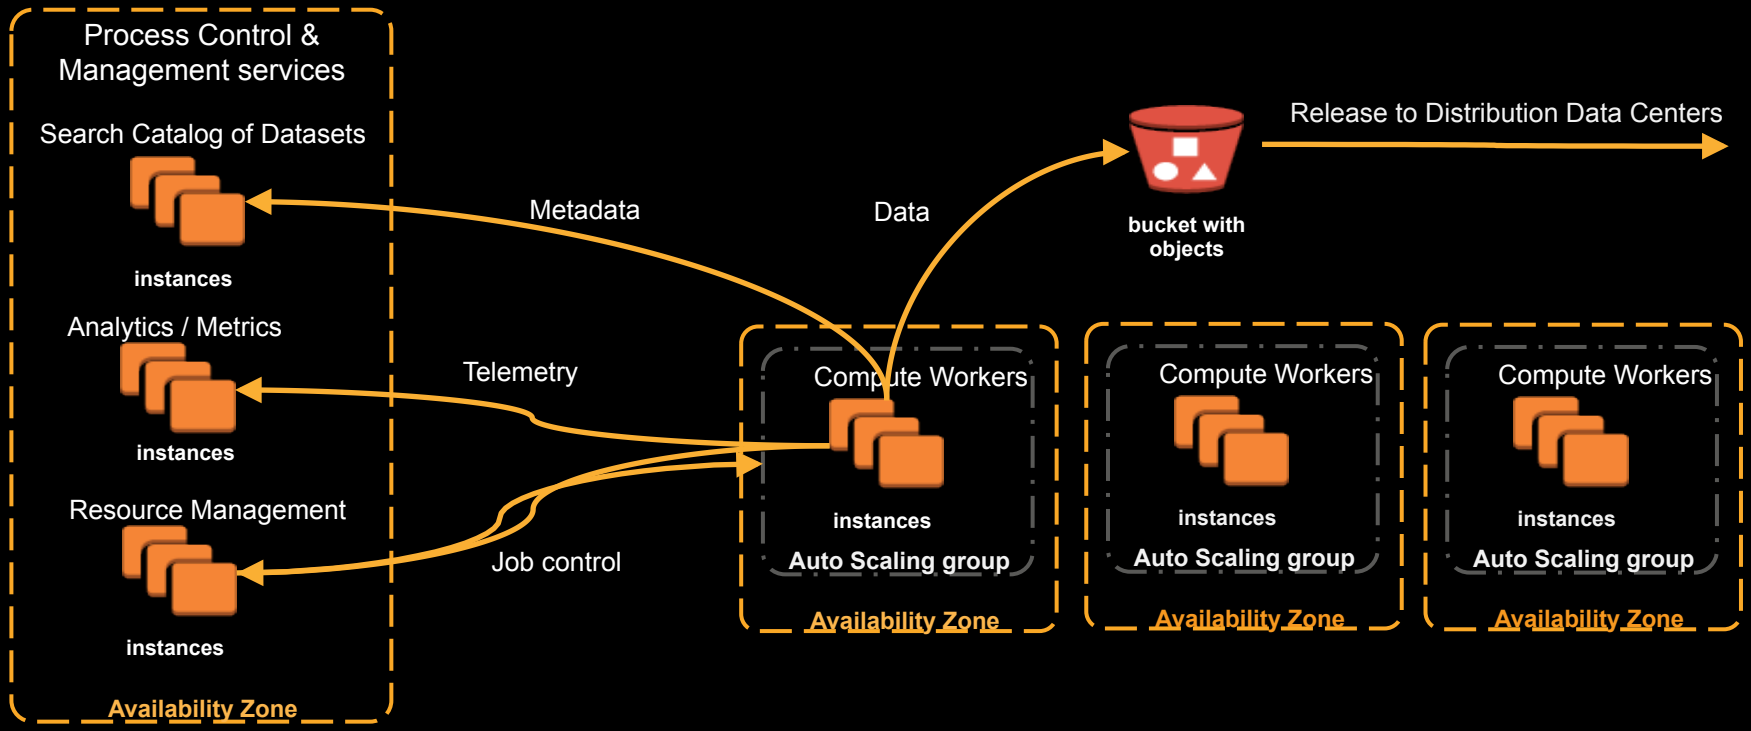

Auto-Scaling Science Data System

The size of the science data system compute nodes can automatically grow/shrink based on processing demand

Auto-scaling tests to over 100,000 ***simultaneous processors***

Auto-scaling group policies

- Auto-scaling size = 21
- Metric alarm of queue size > 20
- Auto-scaling rest period of 60-seconds

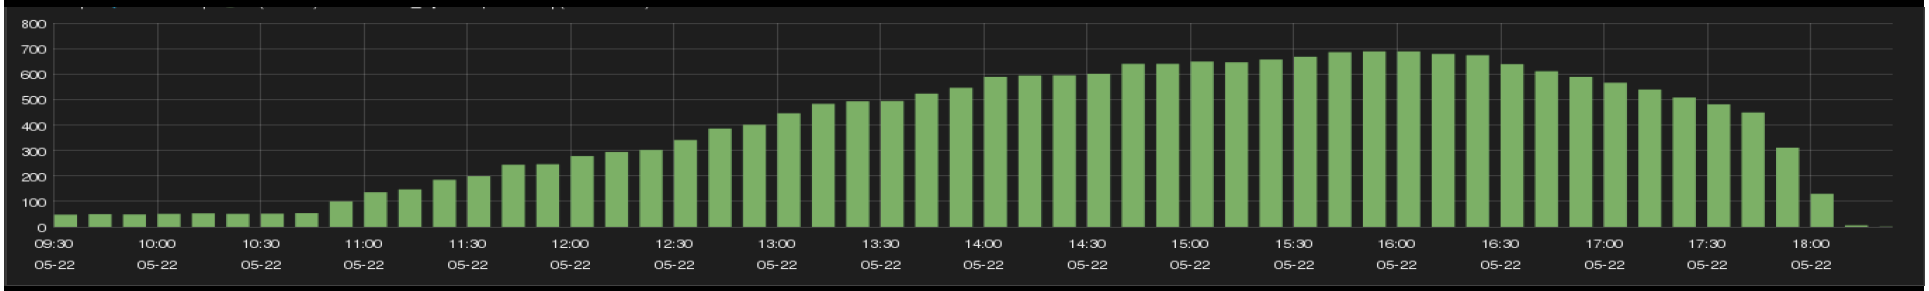

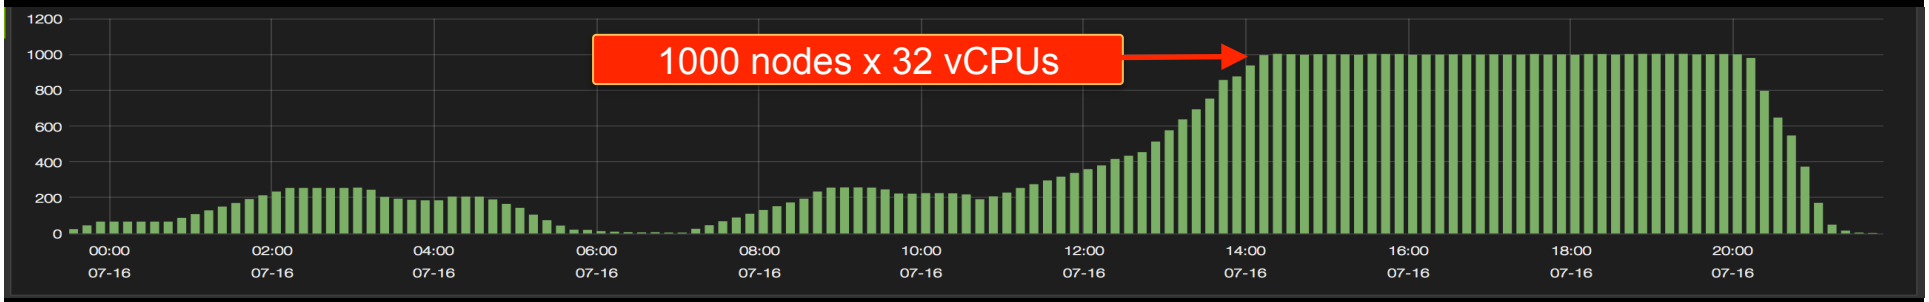

Hybrid Cloud Science Data Systems

OCO-2 L2 Full Physics processing operational in AWS

- Processing of L2 full physics data products in Amazon cloud across multiple regions
- Scaled up **thousands** of compute nodes
- Demonstrated capability of higher internal data throughput rates than NISAR needs

compute nodes over time

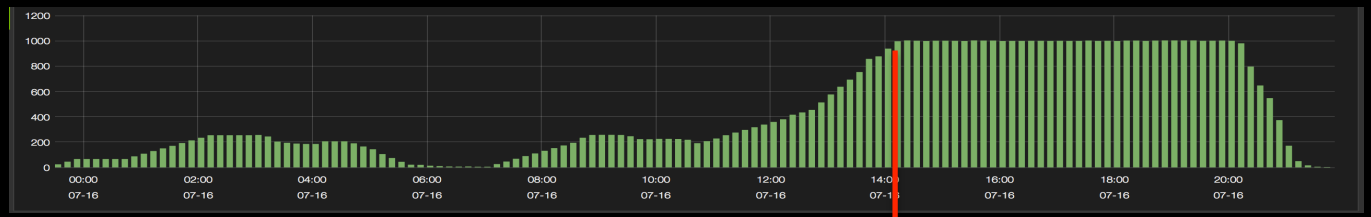

Per node transfer rate over time

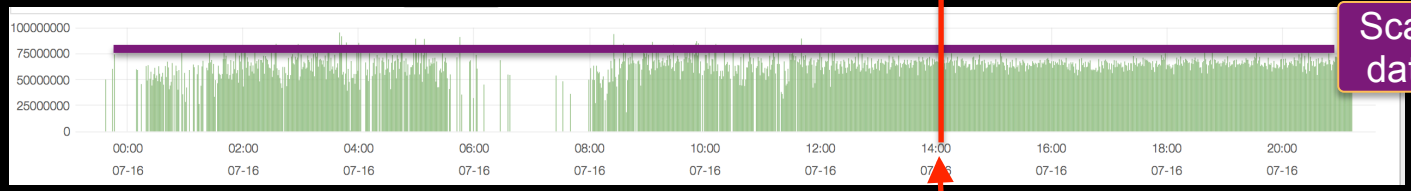

Scalable internal data throughput

@ 32,000 PGEs on 1,000 nodes

Availability Zone Load Rebalancing

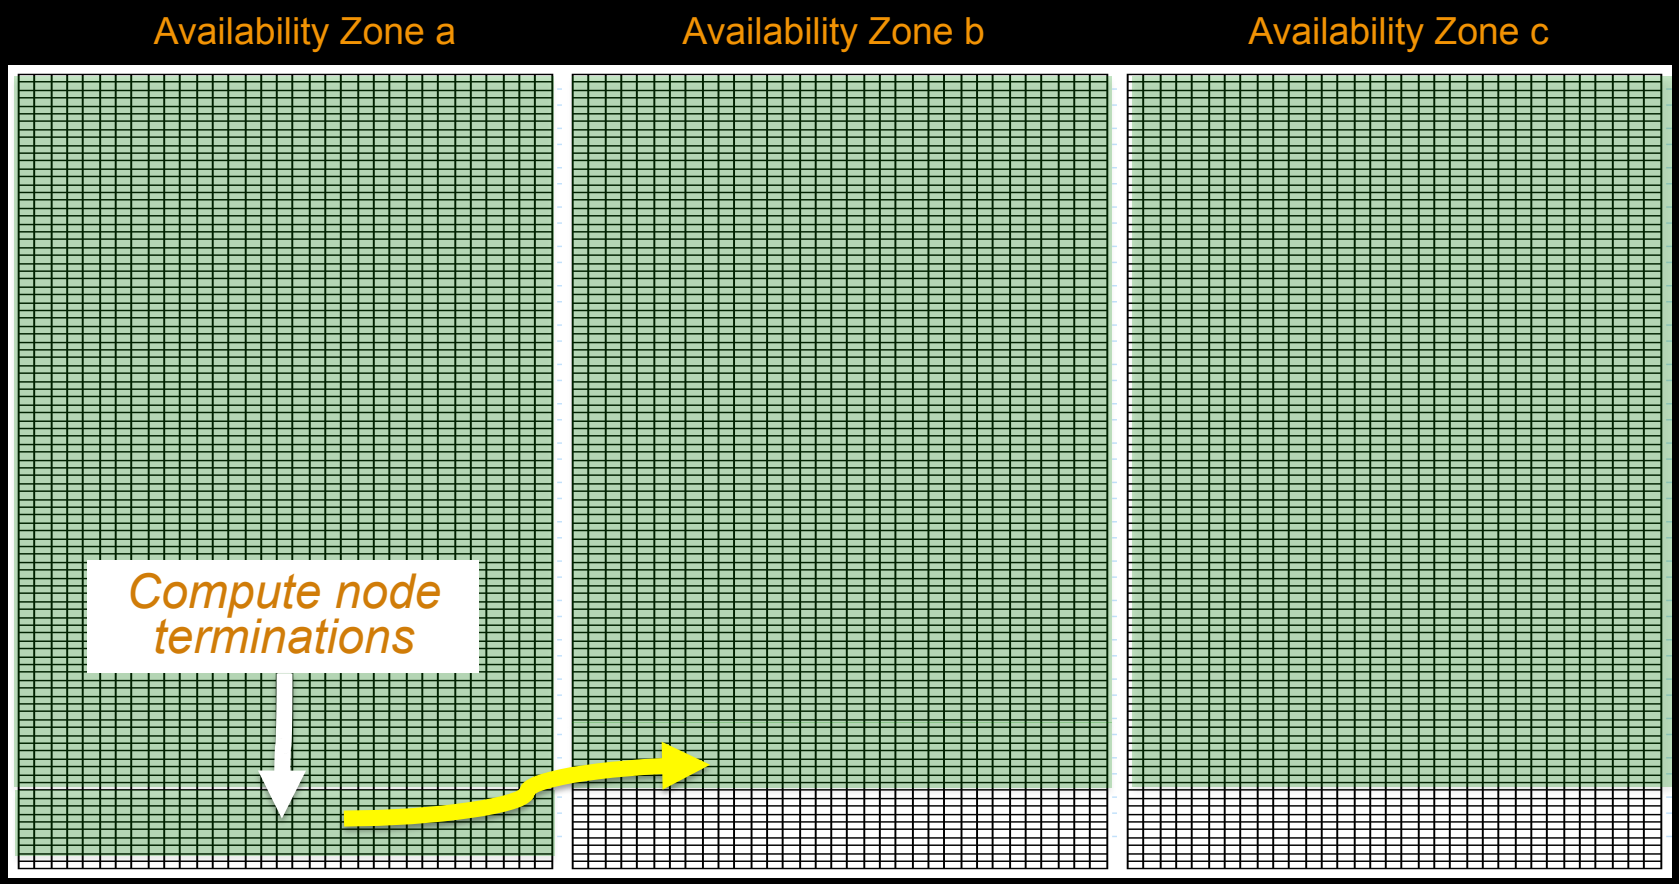

High-Resiliency

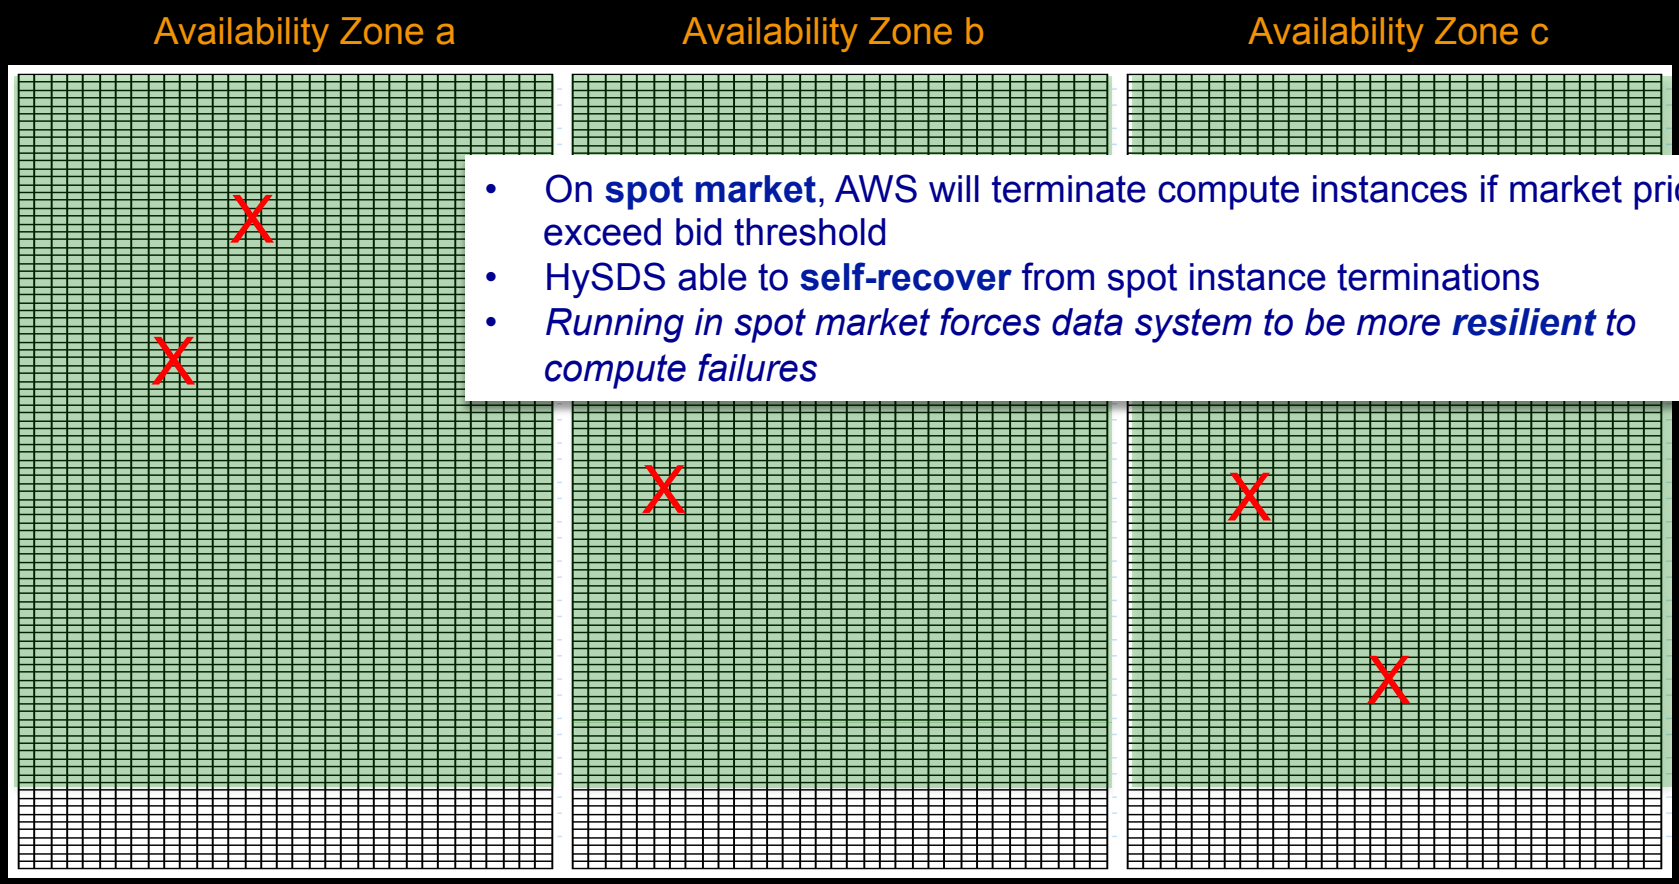

- On **spot market**, AWS will terminate compute instances if market prices exceed bid threshold
- HySDS able to **self-recover** from spot instance terminations
- *Running in spot market forces data system to be more **resilient** to compute failures*

Prioritized Rapid Analysis

URGENT RESPONSE

Monitoring & Urgent Response

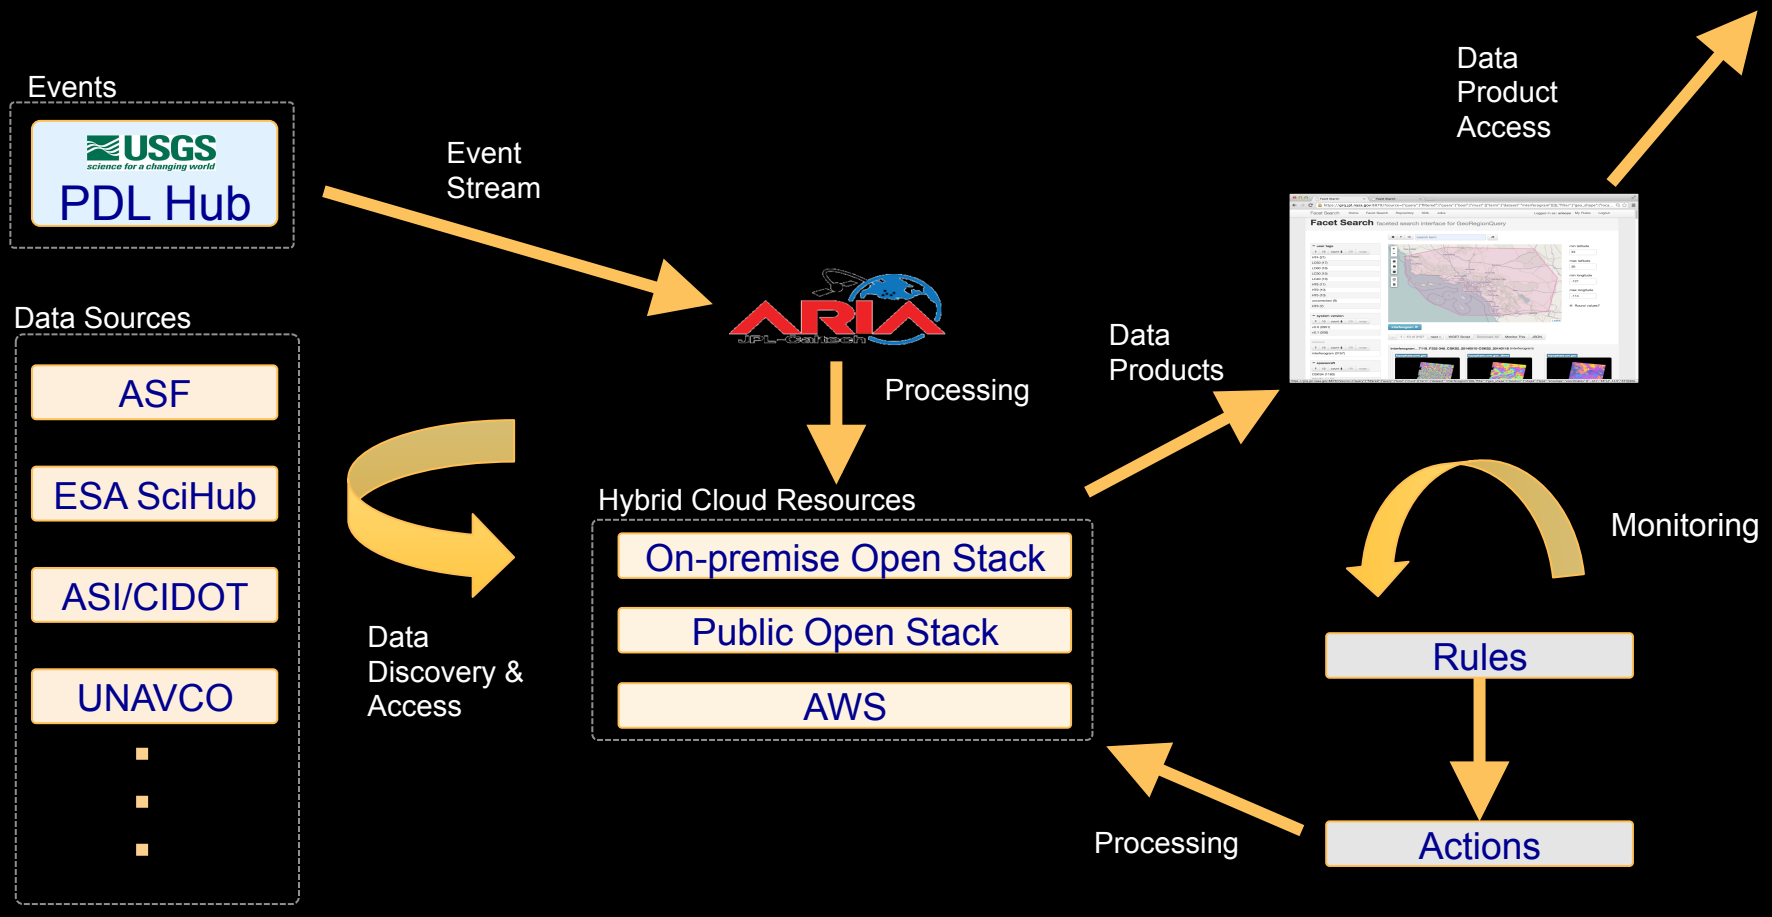

ARIA Automated Amatrice Italy Earthquake (August 23, 2016) Response Interferogram

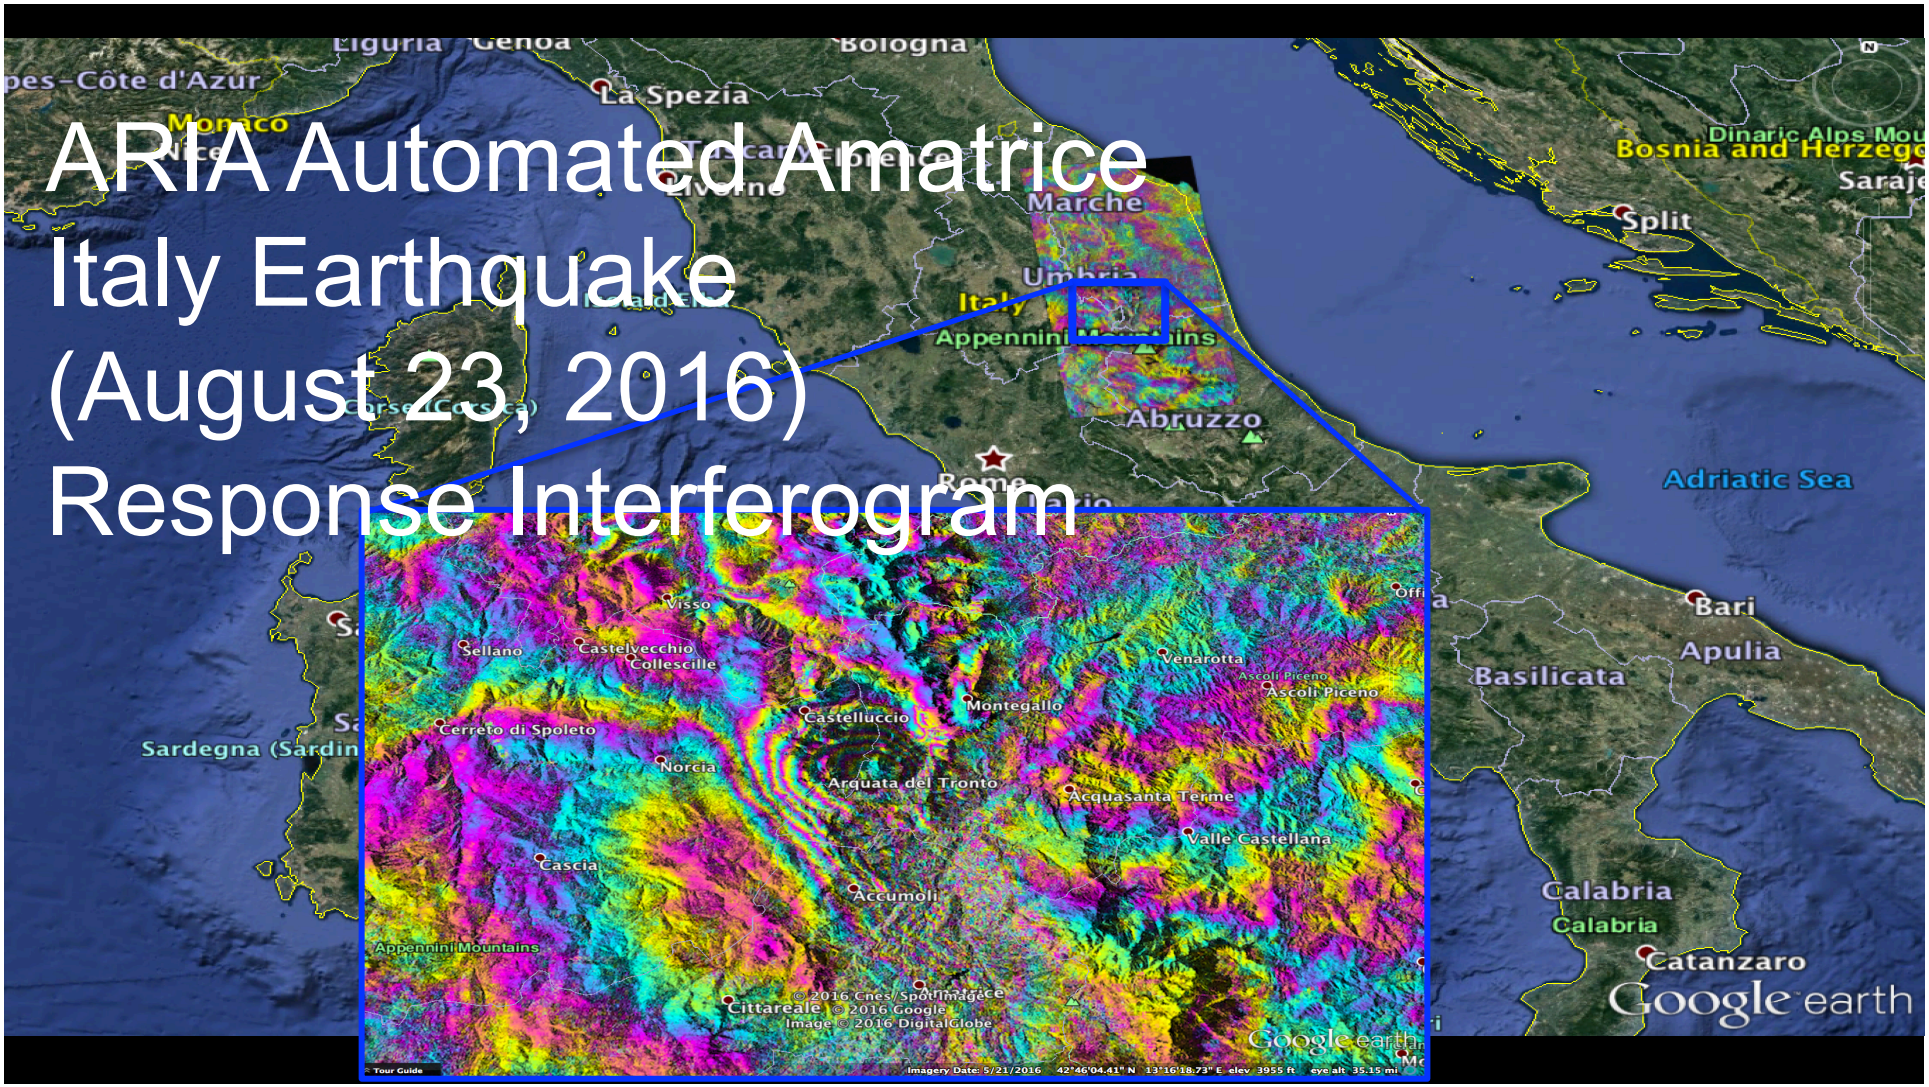

CLOUD ECONOMICS

Spot Market

					US-West-2 (Oregon)							
					Hourly Costs				Per vCPU Costs			
instance	vCPU	memory	memory-cpu ratio	disks	on-demand ($/hr)	reserved 1-yr upfront ($/hr)	reserved 3-yr upfront ($/hr)	spot linux ($/hr)	on-demand ($/cpu/hr)	reserved 1-yr upfront ($/cpu/hr)	reserved 3-yr upfront ($/cpu/hr)	spot linux ($/cpu/hr)
m2.4xlarge	8	68.4	8.55	2 x 840	$1.0780	$0.4087	$0.2441	$0.1000	$0.1348	$0.0511	$0.0306	$0.0125
cc2.8xlarge	32	60.5	1.89	4 x 840	$2.0000	$0.9131	$0.6137	$0.2705	$0.0625	$0.0285	$0.0191	$0.0085
m3.2xlarge	8	30.0	3.75	SSD 2 x 80	$0.6160	$0.3750	$0.2300	$0.0700	$0.0770	$0.0469	$0.0288	$0.0088
c3.8xlarge	32	60.0	1.88	SSD 2 x 320	$1.6800	$0.9920	$0.6280	$2.4001	$0.0525	$0.0310	$0.0196	$0.0750
r3.8xlarge	32	244.0	7.63	SSD 2 x 320	$2.8000	$1.4860	$0.9820	$2.8000	$0.0875	$0.0464	$0.0307	$0.0875
c3.xlarge	4	7.5	1.88	SSD 2 x 40	$0.2310	$0.1370	$0.0870	$0.0353	$0.0578	$0.0343	$0.0218	$0.0088

- Major cost savings (**75%-90%** savings over on-demand)…if can use **spot instances**
- On **spot market**, AWS will **terminate** compute instances if market prices exceed bid threshold
- HySDS able to **self-recover** from spot instance terminations
- *Running in spot market forces data system to be more **resilient** to compute failures*

“Market Maker”

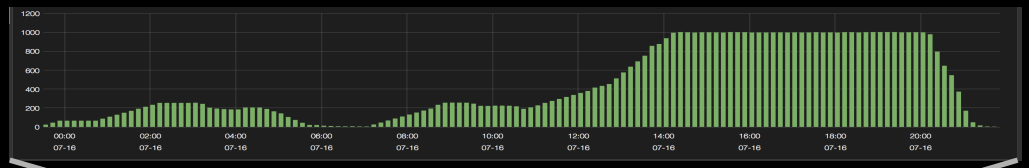

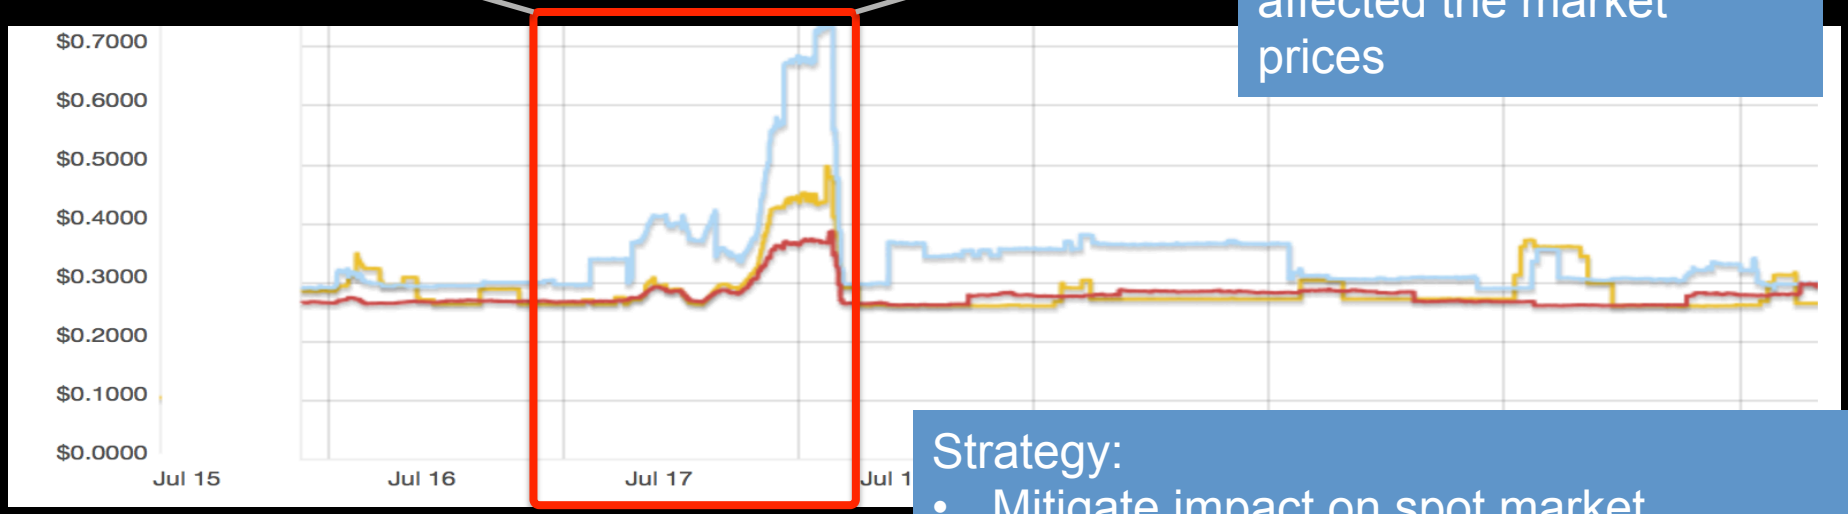

This OCO-2 production run of 1000 x 32vCPUs affected the market prices

Strategy:

- Mitigate impact on spot market
- Diversification of resources
- “Spot fleet”

Looking Ahead

Next Generation Missions

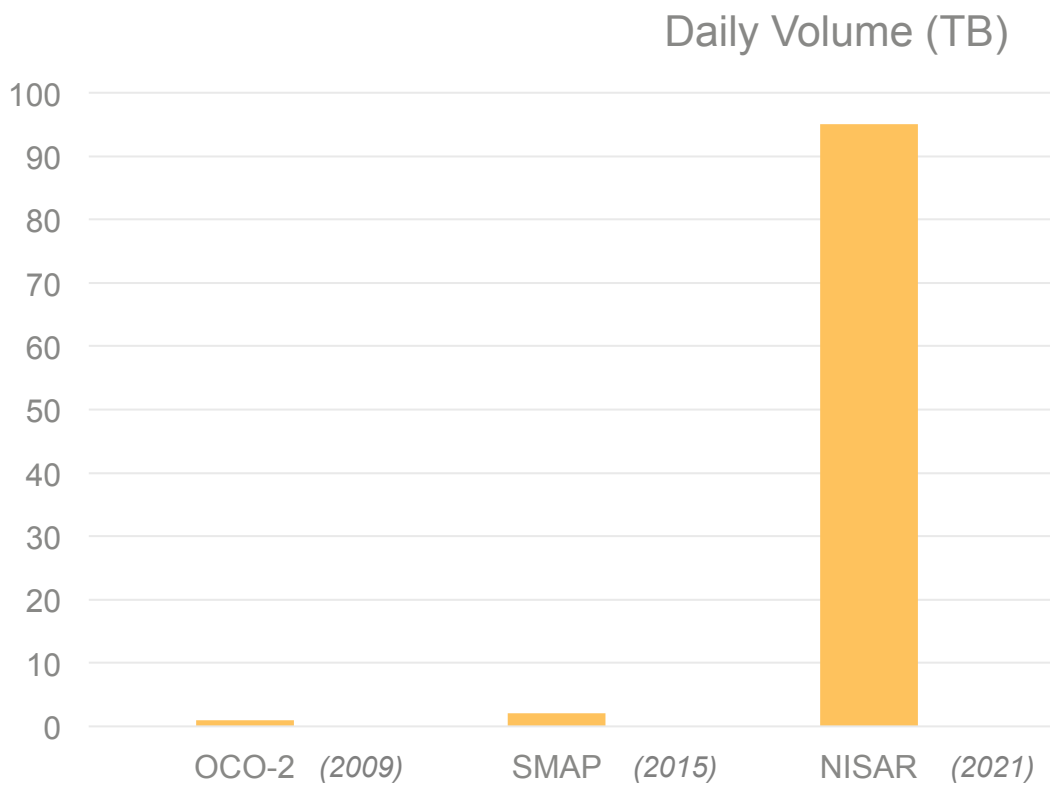

- The volume of data produced is larger than previous missions
- Data storage, processing, movement, and **costs** are the biggest challenges!

Major Systems of Each Earth Science Remote Sensing Mission

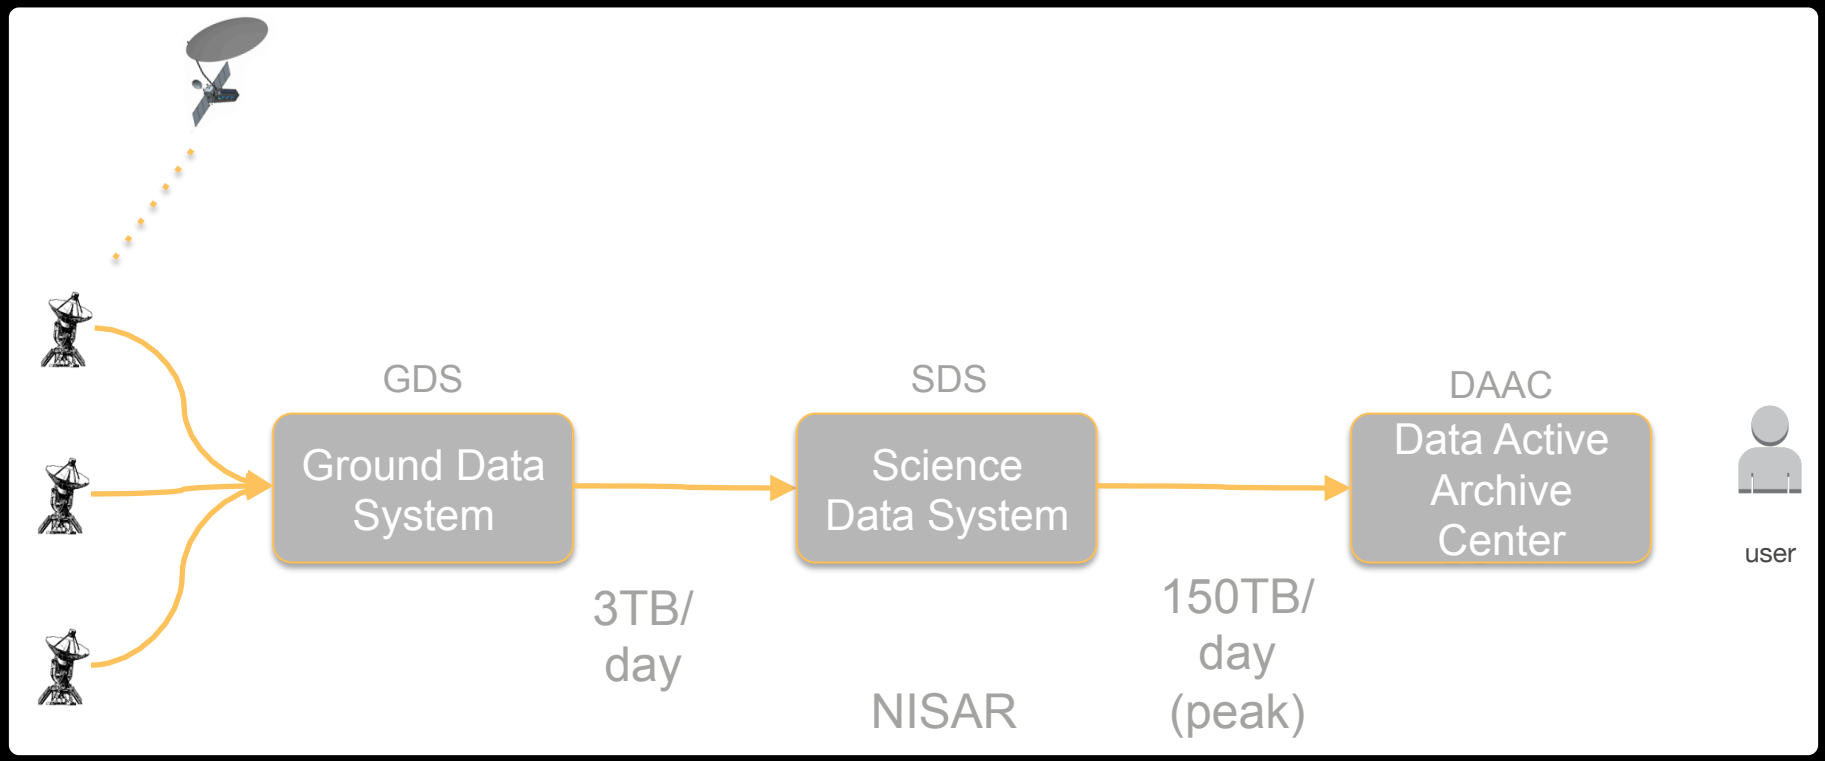

Towards a Cloud-based "Data Lake"

- It's about *collocation*!
- Minimize data movement
- Maximize user services
- Run on public cloud provider or at an on-premise data center

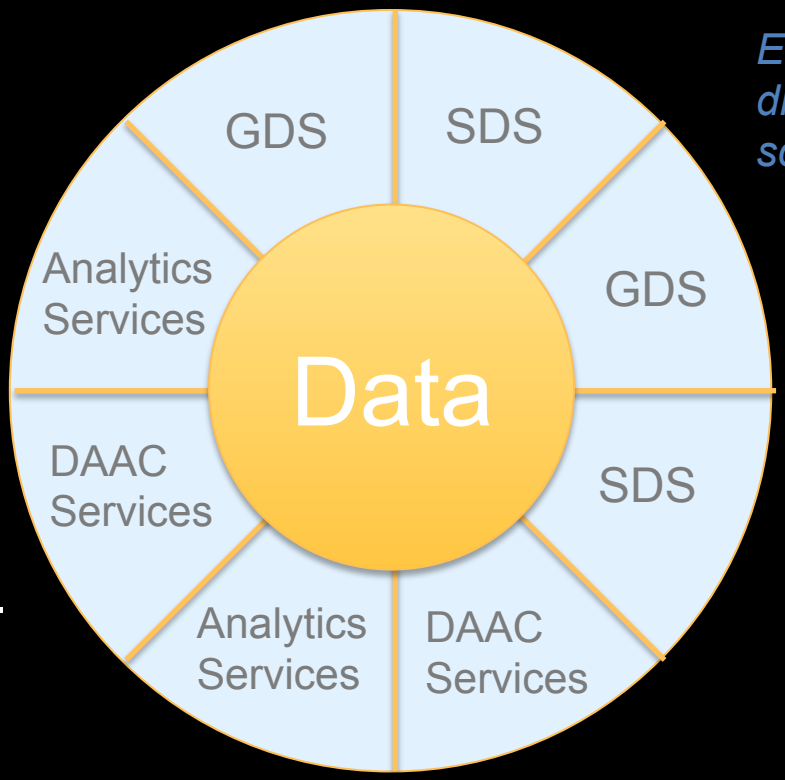

Enabling multi-disciplinary data science analysis

AWS re:Invent

Thank you!

amazon web services