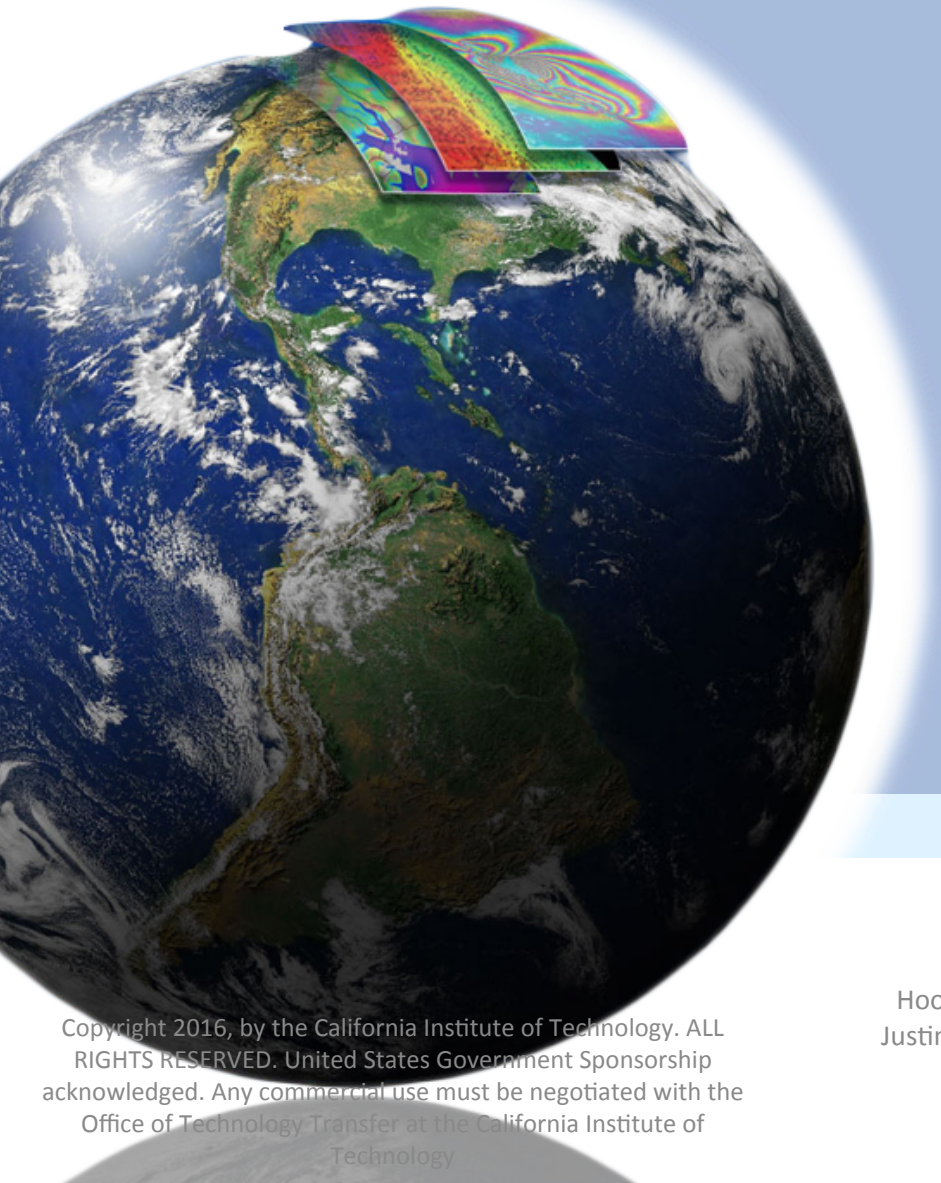

SAR Science Data Processing (SDP) Foundry

2016 Earth Science Technology Forum (ESTF2016)
Annapolis, MD
Wednesday, June 15, 2016

Hook Hua, Paul Rosen, Susan Owen, Gerald Manipon, Michael Starch, Lan Dang, Justin Linick, Piyush Agram, Gian Franco Sacco, Eric Fielding, Eric Gurrola, Tom Farr
[1] Jet Propulsion Laboratory
[2] California Institute of Technology

Motivation

- Growing volumes of SAR Instrument data remain unprocessed to higher level science products
  - 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
- Historically, SAR data products are hand built
  - Dependent upon specialized knowledge by investigator
  - Expensive, laborious and slow to create (days)
  - Unable to scale processing capacity
- Ongoing SAR processing costs
  - On-ramping SAR instrument-specific processing algorithms
  - SAR science data processing
  - Data storage and movement
- *Optimization opportunities for consolidation of SAR processing approach*

Sentinel-1A

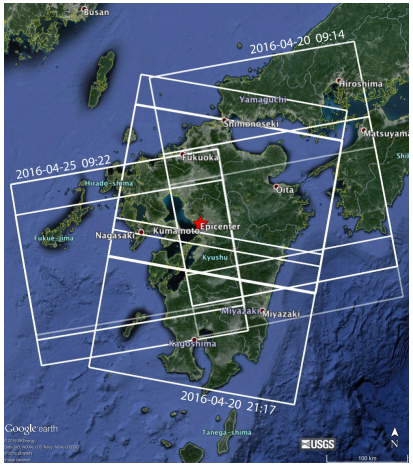

CSK

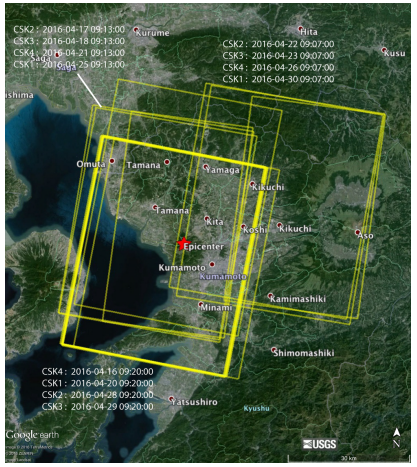

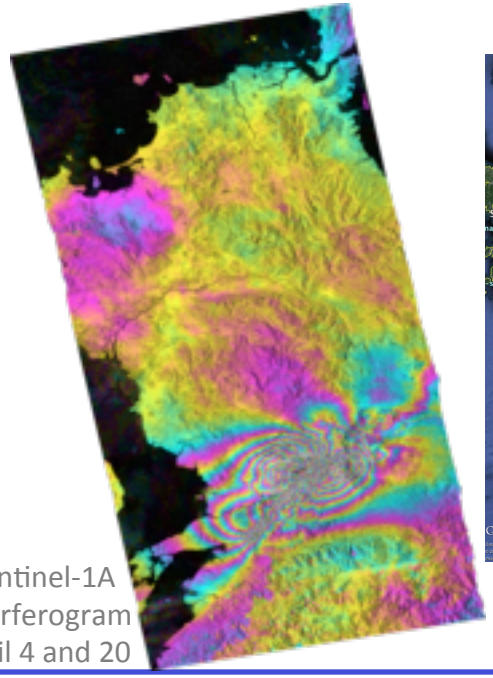

Sentinel-1A interferogram April 4 and 20

ALOS-2

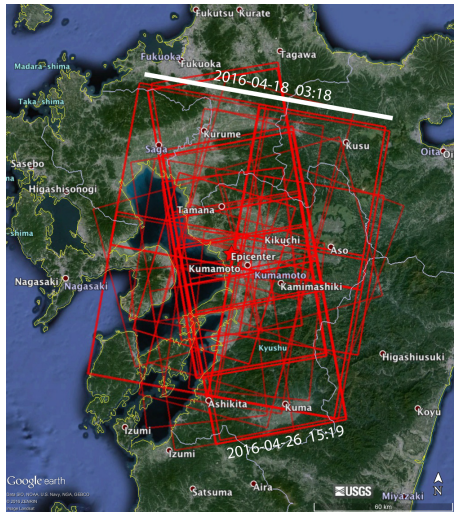

NASA SAR Science Data Processing (SDP) Foundry Enables Science

ESTF 2016

- The focus of this talk is on the **ARIA science data system** behind the SAR SDP Foundry.
- Eric Gurrola will be presenting on the **InSAR Scientific Computing Environment (ISCE)** component tomorrow @ 10:00am

Instrument Data Sources

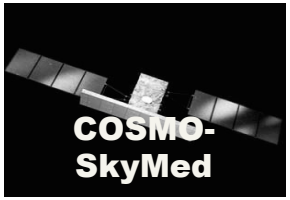

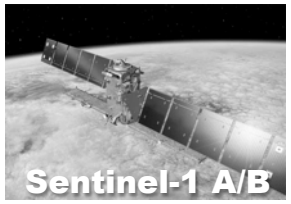

The Cloud-Based SDP Foundry

ARIA
JPL-Caltech

Advanced Rapid Imaging & Analysis for Monitoring Hazards

ISCE

InSar Scientific Computing Environment on the Cloud

Processed Science Products

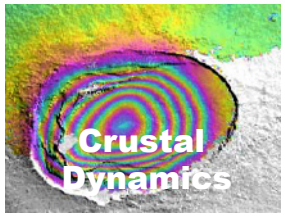

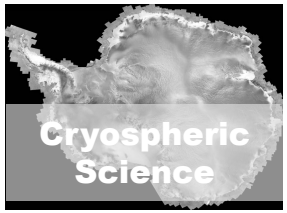

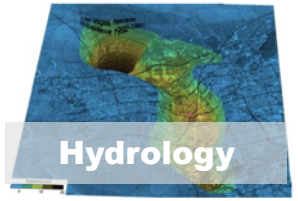

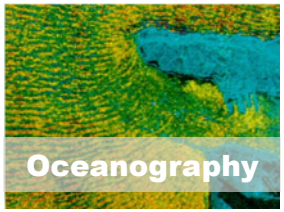

Foundry Components

- Leverage Hybrid-cloud Science Data System
- Leverage InSAR Scientific Computing Environment (ISCE)

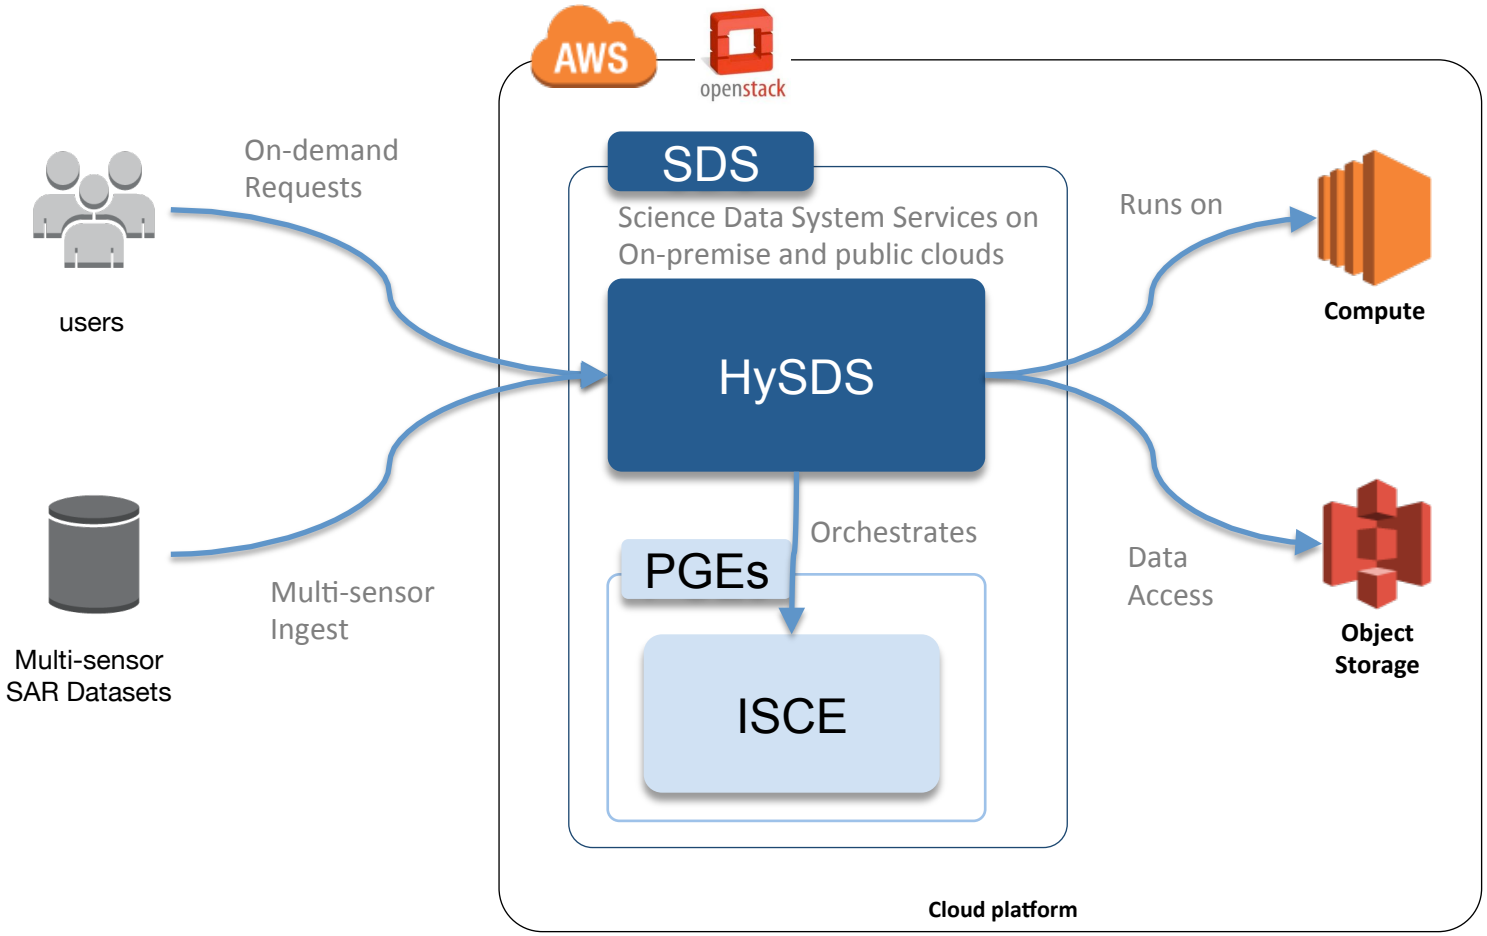

SAR SDP Foundry Interfaces

- Science Data Products — Faceted search and browse

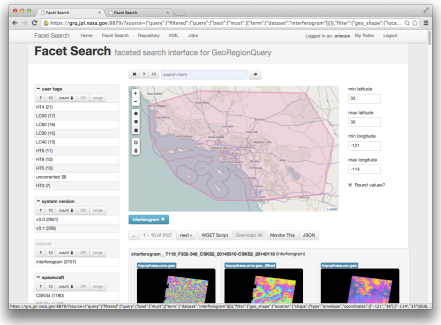

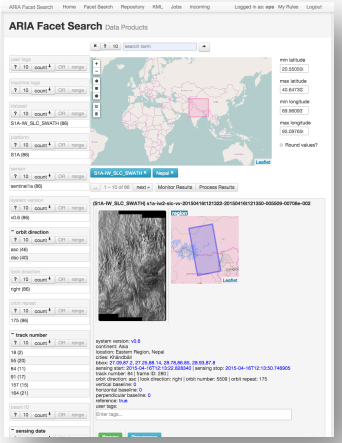

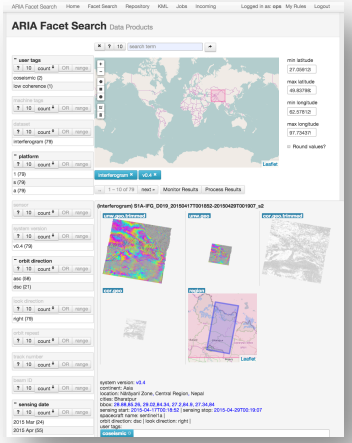

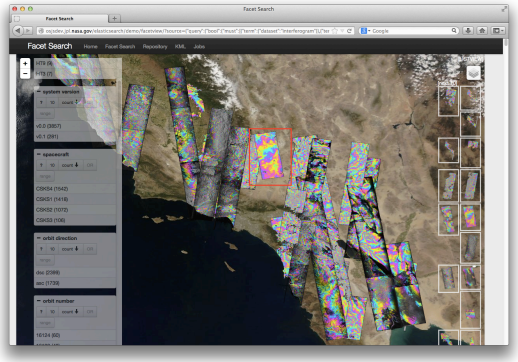

- Science Data System — Faceted Metrics and PROV-ES provenance

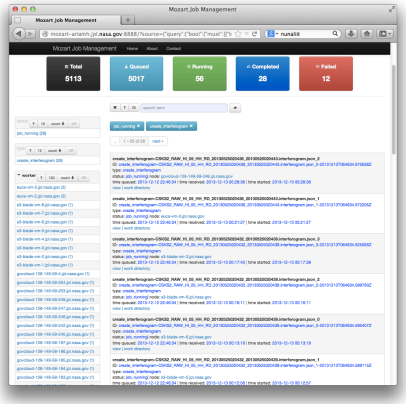

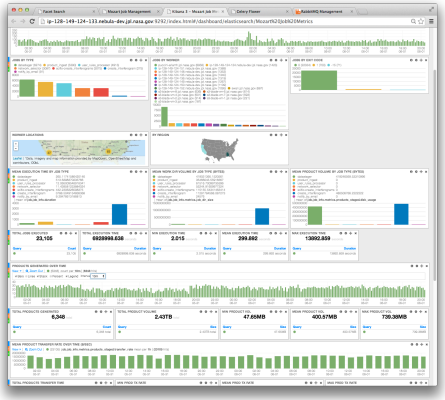

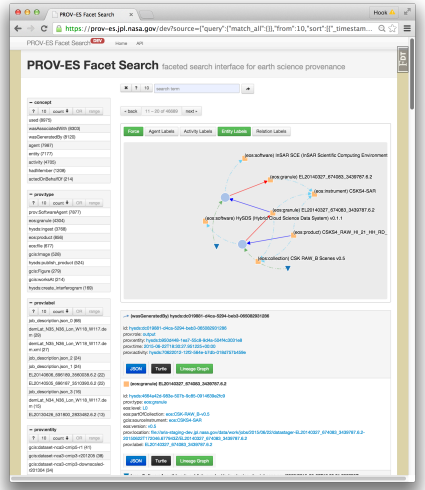

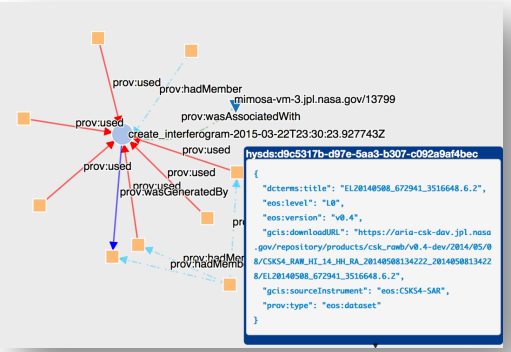

Key Aspects of SAR Foundry

- ESTO-funded Foundry Technology accelerates Science Access
  - On-ramped multi-sensor **data sources** permits on-demand ingest of L0/L1 data products
  - On-ramped Science Products **processing** and metadata, as defined by Science Teams
- On-ramping to a *cloud-based SDP*
  - On-demand Cloud Computing (via AWS) enables PI **pay-as-you-go** processing
- **Characterize** the on-ramping
  - Models the **performance, cost, and scope** of adding science workflows to the Foundry to allow NASA to intelligently develop new processing capabilities
- Custom processing
  - **Custom data processing costs paid for by customers**.
- Processing environment is published and community-accepted
  - **ISCE**: continual improvements to SAR processing algorithms
- Processing improvements are shared among the science communities
  - **HySDS**: scalable SDP with reliable and optimal use of AWS
- **Open Data**: Science Data Products can become available to the communities regardless of who funded their production
  - Consistent with 2004 InSAR Working Group Workshop Summary Report (10/20/2004)

ON-RAMPING

On-ramping SAR Processing

- Identified steps to on-ramp processing support of a new SAR instrument datasets

Science Data System
setup discovery worker (data-provider specific)
setup access worker (sling, may be re-usable)
setup ingest/extractor worker
data formating (e.g. HDF5)
data preparation
GIS compliancy
publishing dataset to Faceted Search, ES, storage
setup worker for workflow
setup worker for network selection
setup worker for create interferogram
update faceted view templates

Algorithm & Processing
Starting Point
Ending Point
metadata extractors for L0 dataset
metadata extractors for L1 dataset
metadata extractors for L2 dataset
pipeline for processing from raw to interferogram
script prototyping
Framework modifications
settings modifications (to get sensor data to run well; e.g. unwrapping parameters)
makeraw modifications
formslc modifications
topo modifications
resample modifications
geocode modifications
offset estimation modifications
dem and water mask estimation modifications
correlation estimation modifications
unwrapping modifications
network selection
giant: time series

Integration & Testing
integrate network selection
integrate isce scripts into interferogram creation workflow
setup production trigger rules for ingest of L0/L1 data
setup production trigger rules for processing to L2 interferograms
deploy on dev cloud SDS
run tests with golden data set
benchmark performance to find optimal instance types
delpoy on ops cloud SDS

Data Ingest: High-Efficiency Data Slinging

- Tsunami-UDP
  - accelerated file transfer protocols that implement *TCP for control* and *UDP for data*
  - *"increased network efficiency by replacing TCP's packet acknowledgement-based congestion control mechanisms with an alternative model that leverages UDP and is focused on the efficiency of data transfer on lossy or variable-latency networks."*
  - Efficient for high-latency global-scale data transfers
- Sling large datasets over tsunami-udp over global scales to destination
- Alternative: *Amazon-specific S3 Transfer Acceleration*

Developed Osaka framework

- generalization of data access to cloud storage
  - S3 (AWS)
  - Blobs (Azure)
  - WebDAV (http)
  - file (local)
- all HySDS workers have support for Osaka data access for push-pull

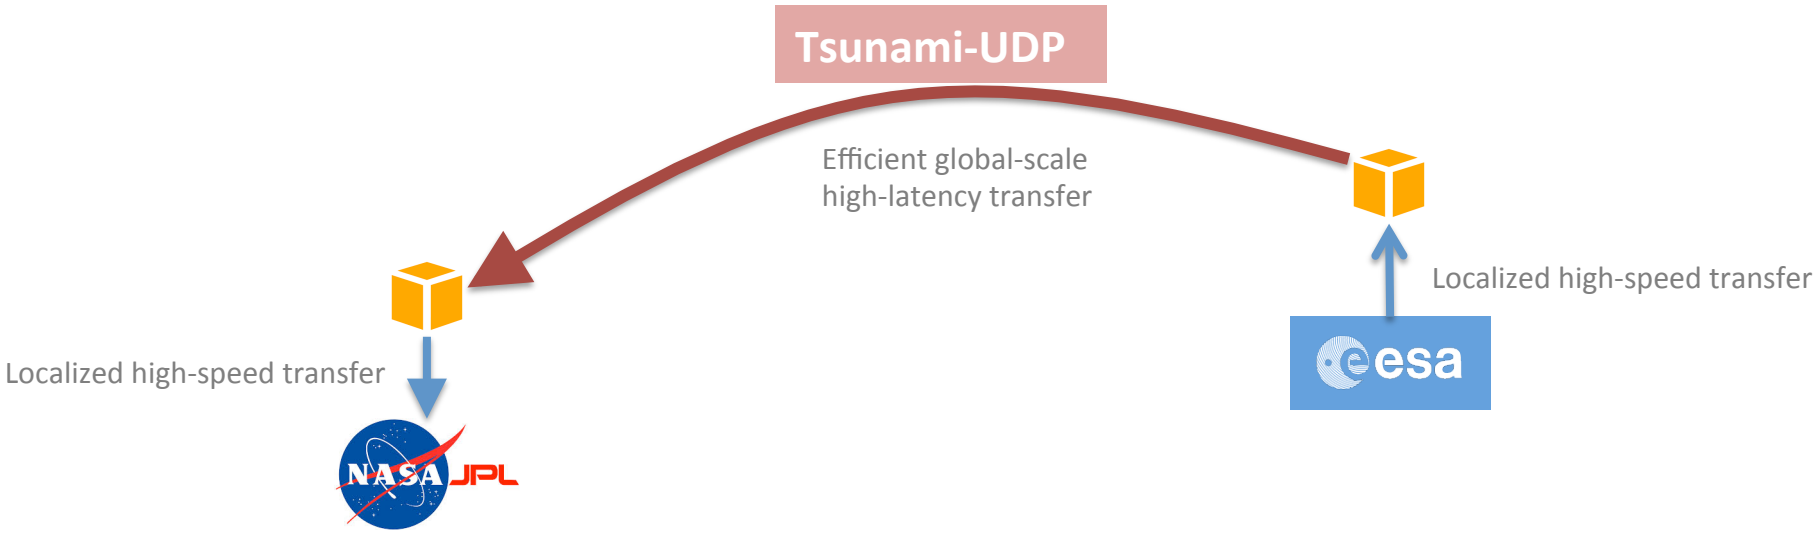

Interferogram Processing Workflow (1 of 2)

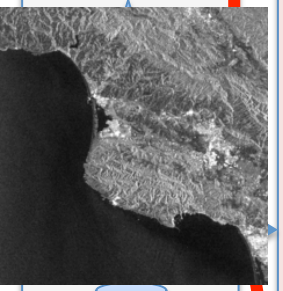

raw data ingest

$T_nFS_nD_1$
$T_nFS_nD_2$
$T_nFS_nD_3$
$T_nFS_nD_4$
$T_nFS_nD_5$
$T_nFS_nD_6$
$T_nFS_n D_{m+1}$
$T_nFS_n D_{m+2}$
$T_nFS_n D_{m+3}$

ISCE Network Selection

Optimize
- Spatial
- Temporal
- Doppler

Baselines

Execute after some months accumulating data

Peg_{12} Dop_{12} → Int-gen → I_{12},A_{12} → tropo Map → tc I_{12},A_{12} → unwrap → tc,u I_{12},A_{12}

Peg_{13} Dop_{13} → int-gen → I_{13},A_{13} → tropo Map → tc I_{13},A_{13} → unwrap → tc,u I_{13},A_{13}

Peg_{24} Dop_{24} → int-gen → I_{24},A_{24} → tropo Map → tc I_{24},A_{24} → unwrap → tc,u I_{24},A_{24}

Peg_{34} Dop_{34} → int-gen → I_{34},A_{34} → tropo Map → tc I_{34},A_{34} → unwrap → tc,u I_{34},A_{34}

Peg_{35} Dop_{35} → int-gen → I_{35},A_{35} → tropo Map → tc I_{35},A_{35} → unwrap → tc,u I_{35},A_{35}

Peg_{46} Dop_{46} → int-gen → I_{46},A_{46} → tropo Map → tc I_{46},A_{46} → unwrap → tc,u I_{46},A_{46}

int

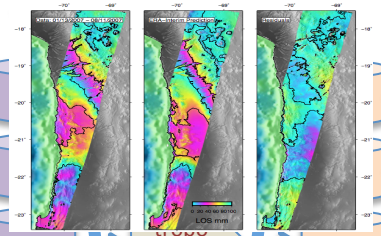

Peg_{tu} Dop_{tu} → int-gen → I_{tu},A_{tu} → unwrap → tc,u I_{tu},A_{tu}

$Peg_{t'u'}$ $Dop_{t'u'}$ → int-gen → $I_{t'u'},A_{t'u'}$ → unwrap → tc,u $I_{t'u'},A_{t'u'}$

$Peg_{t''u''}$ $Dop_{t''u''}$ → int-gen → $I_{..},A_{..}$ → tropo Map → tc $I_{..},A_{..}$ → unwrap → tc,u $I_{..},A_{..}$

$Peg_{t'''u'''}$ $Dop_{t'''u'''}$ → int-gen → $I_{...},A_{...}$ → tropo Map → tc $I_{...},A_{...}$ → unwrap → tc,u $I_{...},A_{...}$

$Peg_{t''''u''''}$ $Dop_{t''''u''''}$ → int-gen → $I_{....},A_{....}$ → tropo Map → tc $I_{....},A_{....}$ → unwrap → tc,u $I_{....},A_{....}$

Cont'd Next Page

Cont'd Next Page

Interferogram Processing Workflow (2 of 2)

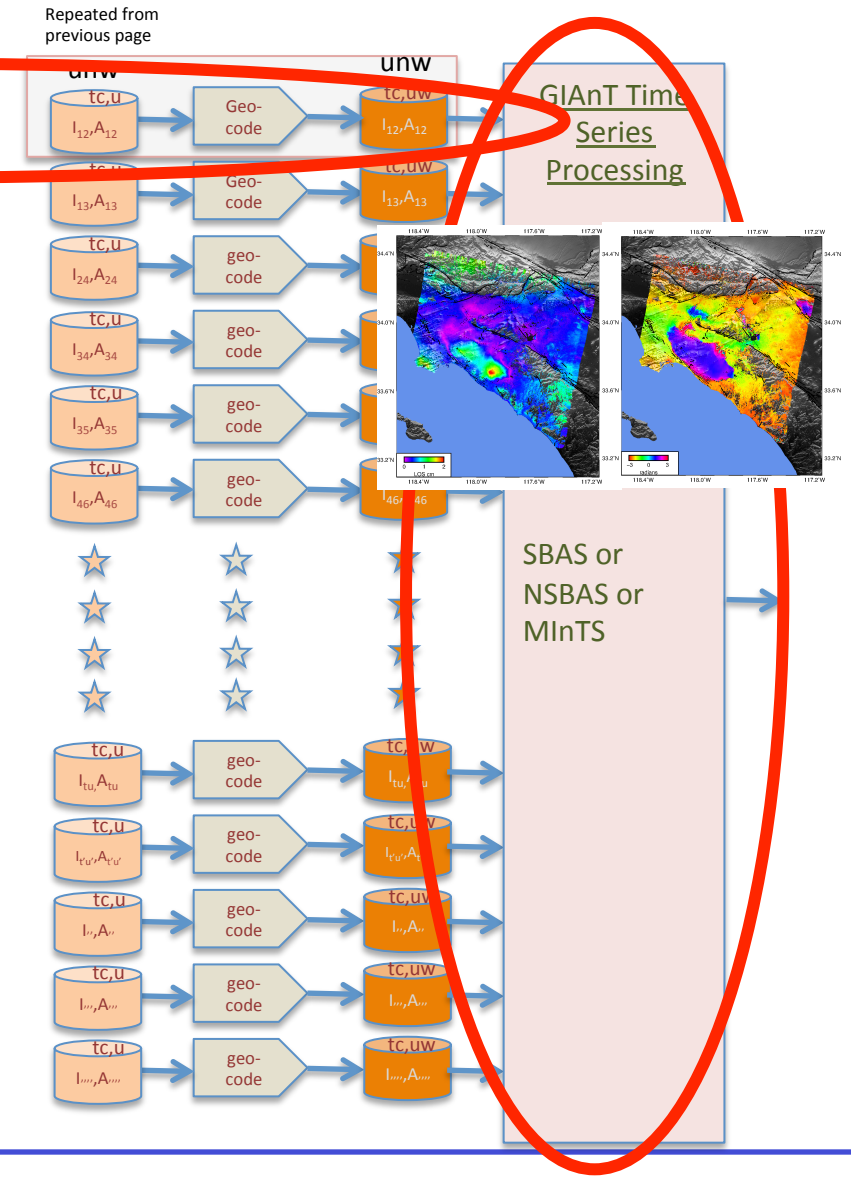

- Each horizontal row is a **scalable processing pipeline in the cloud**
- **Consolidating** the SAR on-ramping
  - Additional SAR sensor support is **added into this generic pipeline**
  - SDS orchestrates and re-uses similar pipelines
- Paying customers control **processing customizations** needed for their data sets
- ISCE team working on newer pipeline approach (“*option 2*”)

Multi-sensor SAR

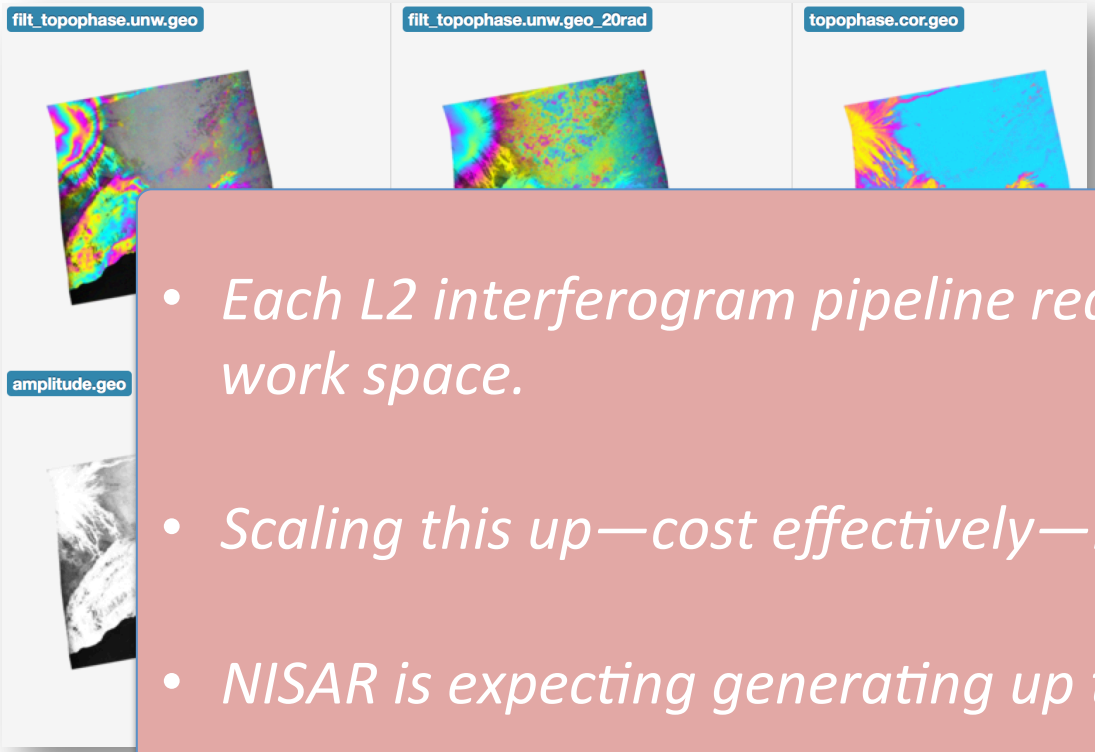

- *Each L2 interferogram pipeline requires 60GB-100GB of work space.*
- *Scaling this up—cost effectively—is key*
- *NISAR is expecting generating up to ~150TB per day*

CSKS1 over Kilauea vc

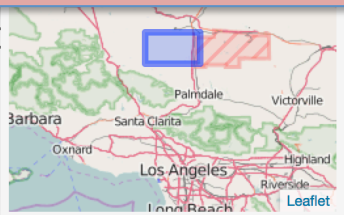

UAVSAR flight over San Andreas

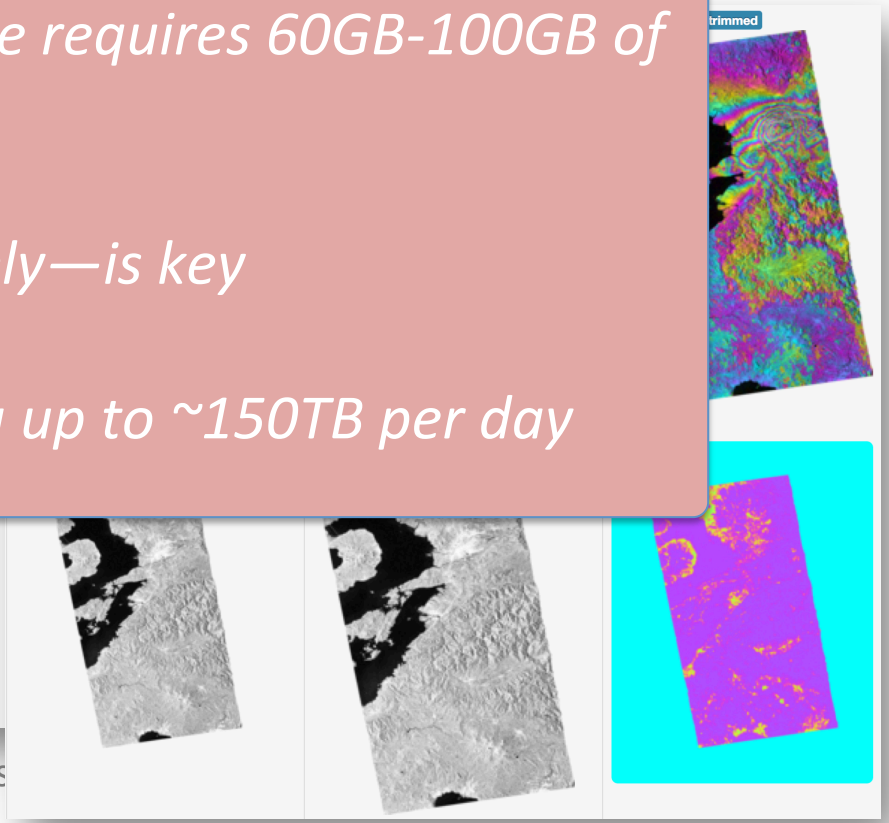

Sentinel-1A over Kumamoto earthquake, Japan

CONSOLIDATED LARGE-SCALE SAR PROCESSING

High-Level Cloud-based SDS View

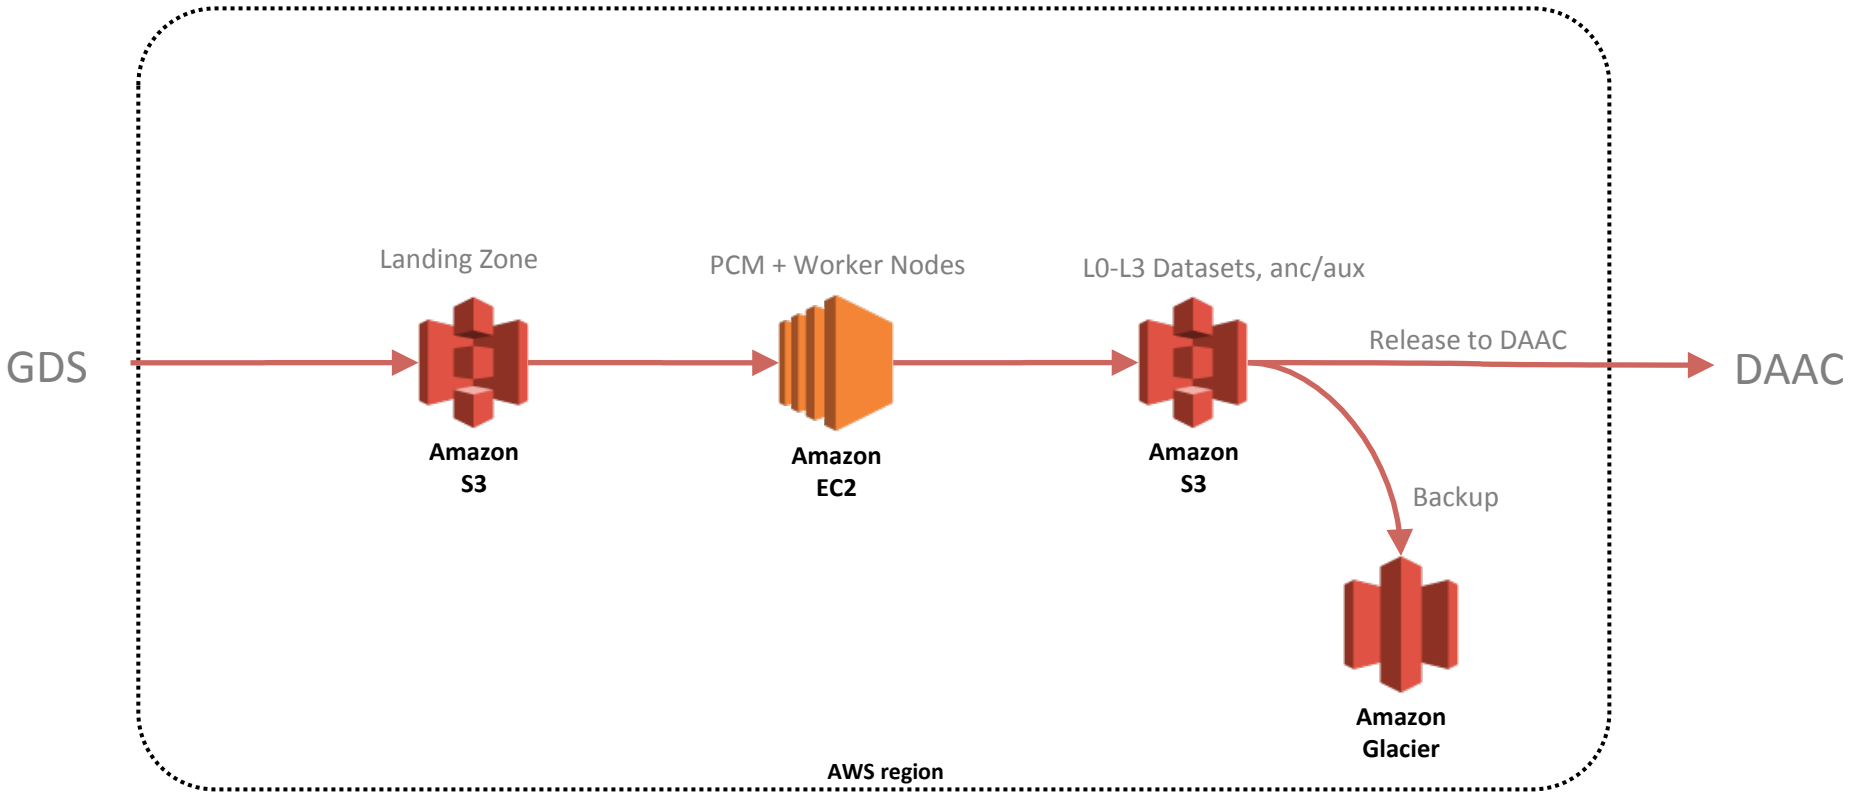

Process Control & Management

- Recall the pipeline:

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

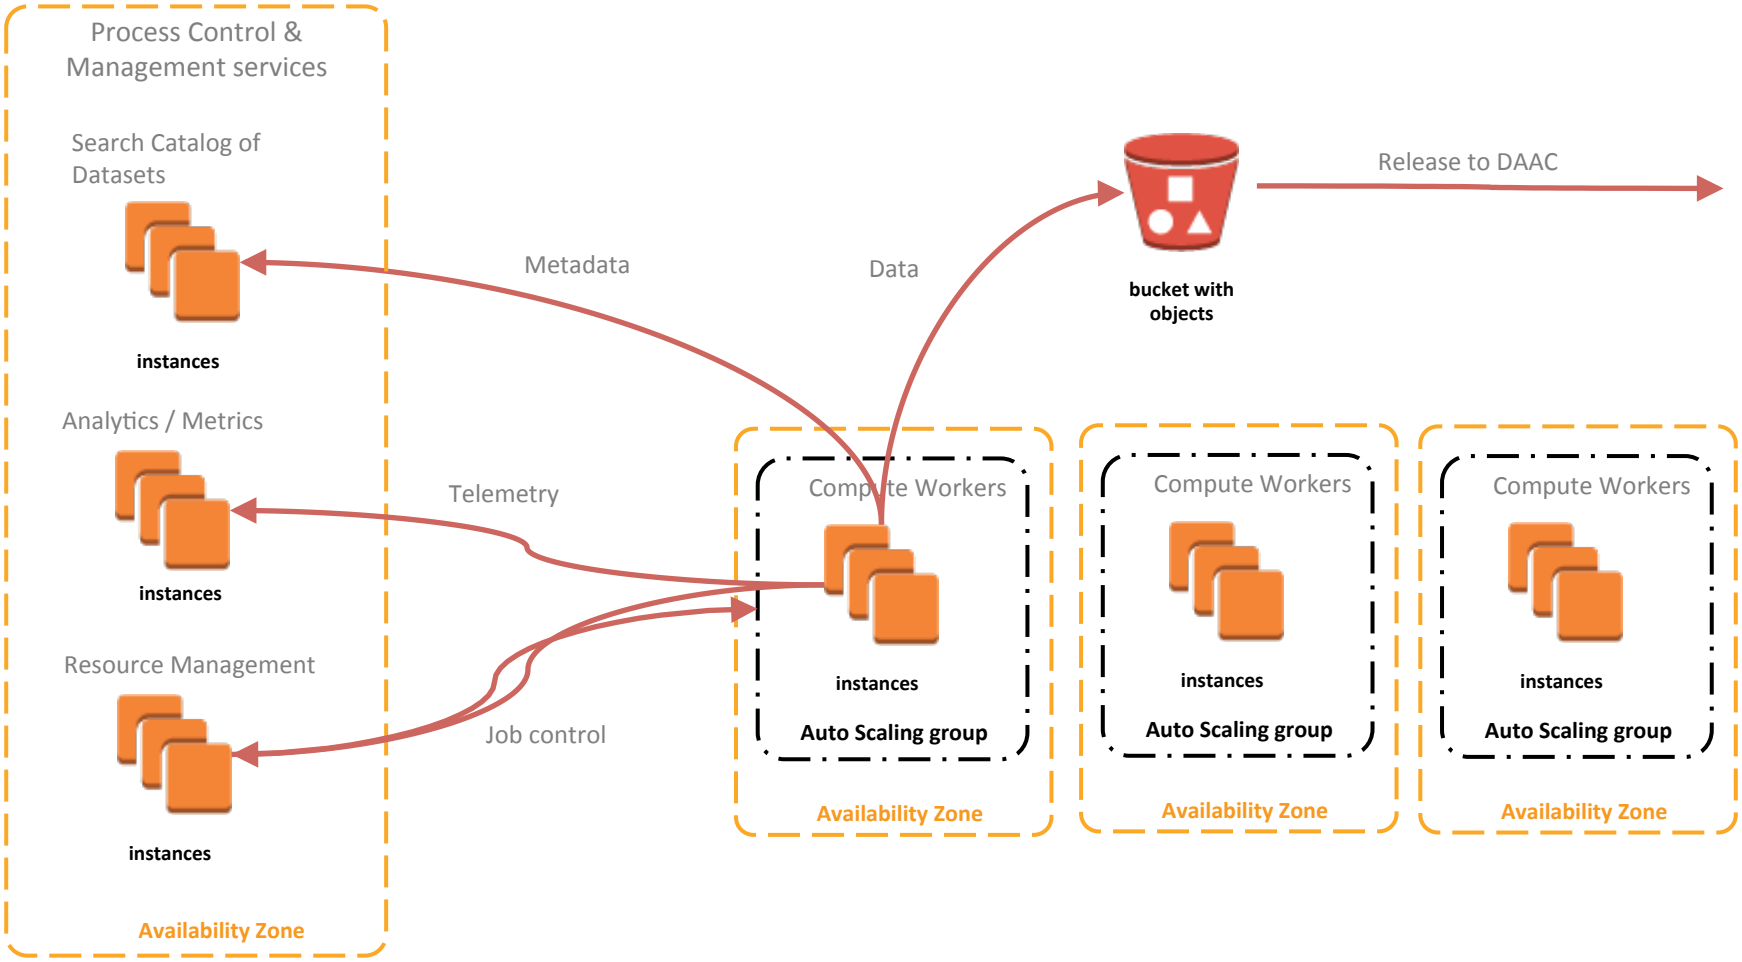

Customer-specific Processing Cloud Billing

- Custom processing pipelines are mapped to their own auto-scaling groups
  - Requestor-specific work queues drive the ASGs
- **"Requestor pays"** for their own specific SAR processing jobs
  1. **Via tagging**
     - Each pipeline's auto-scaling groups are tagged appropriately
     - Usage tags are later billed to requestor
  2. **Via AWS accounts**
     - Compute workers in their auto-scaling group can be under the requestor's own AWS account
- NASA managed clouds
  - AMCE
  - NGAP

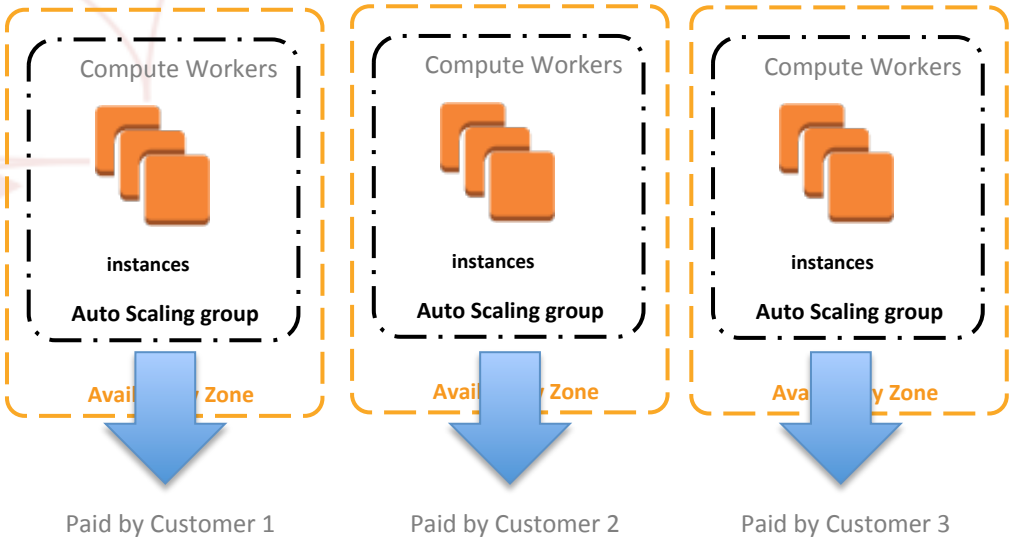

Cloud Security

- SAR processing compute can be locked down in security group
  - Currently **is EAR99 export controlled.** Restricted to **GovCloud** region
- Process, control, and management services in **separate security group** with user access

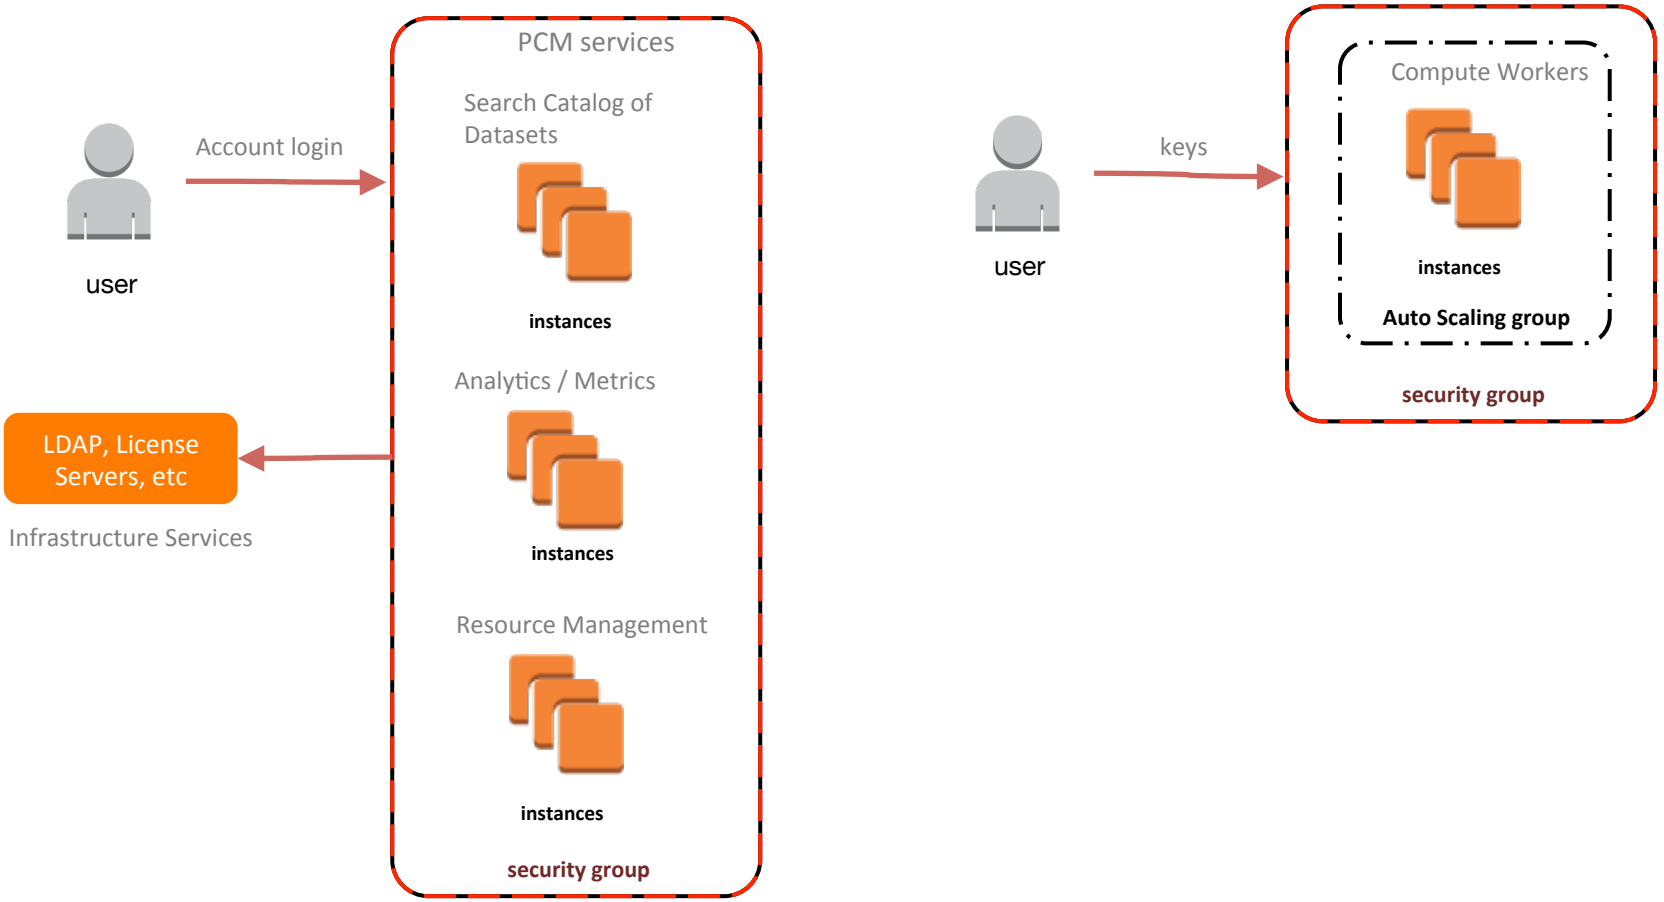

Hybrid SDS Variant

- Optional hybrid approach to reduce longer-term operational costs
- Run process, control, and management services on-premises
- Bursting out to AWS for processing only when needed. Ideal for bulk processing campaigns.

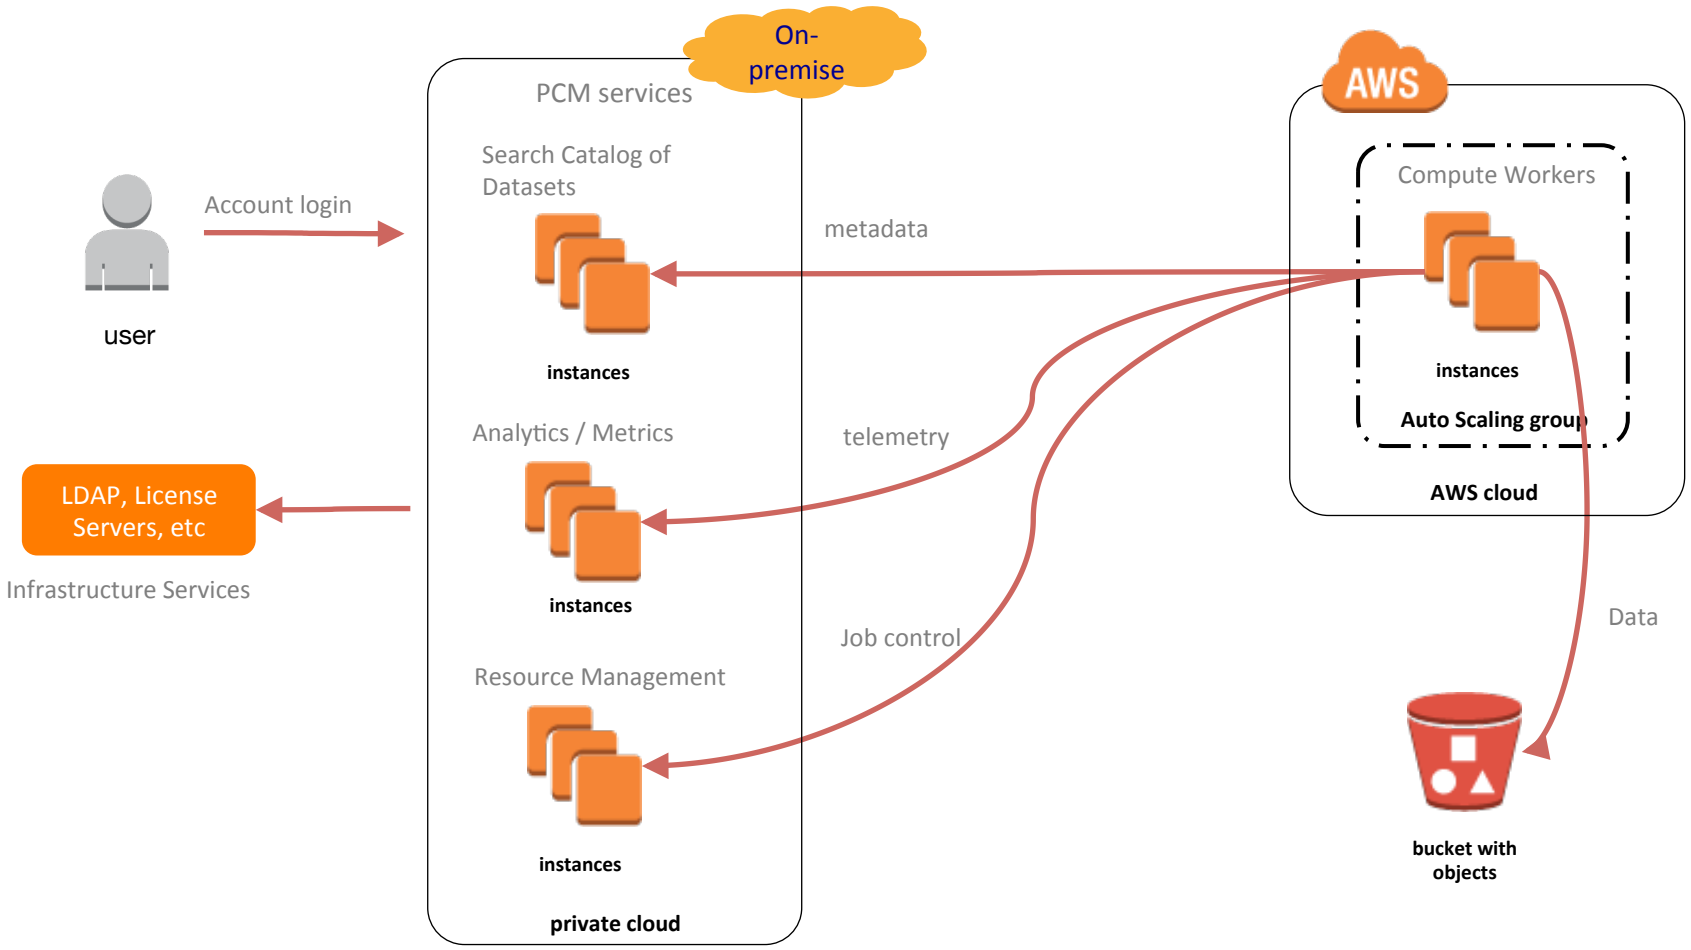

"Containerizing" Processing Steps

- **Containerizing**
  - Encapsulating analysis steps into more portable and self-contained Docker Containers
- **Agility**
  - Foster agility through rapid development and deployment of analysis steps
- **Portability**
  - Deploy analysis steps in private and public clouds
- **Scalability**
  - Large-scale deployment of Containers to compute fleet
- **Provenance**
  - Archive PGE Containers in AWS/S3
  - Reproduce all existing and prior versions of data analysis and production
  - *"use what you store, and store what you use"*
  - Re-run analysis by data system and DAAC

*"Docker containers wrap up a piece of software in a complete filesystem that contains everything it needs to run: **code, runtime, system tools, system libraries** – anything you can install on a server. This guarantees that it will always run the same, regardless of the environment it is running in."*

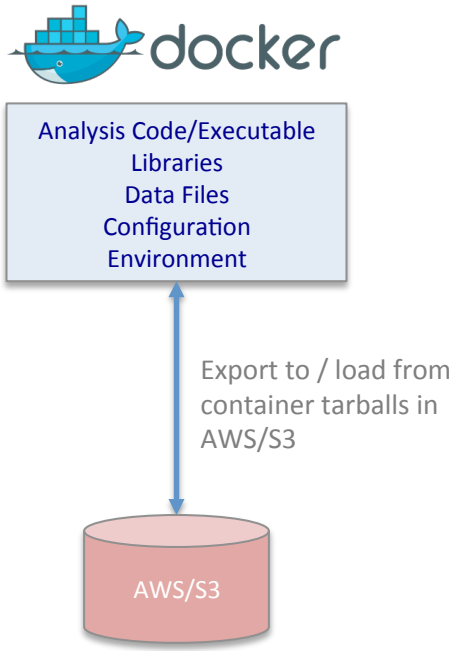

ISCE in Compute Instances

- On-ramping new SAR sensor processing support involves:
  - Updating **ISCE processing steps**
  - Integrating the new ISCE **PGE** modules as Containers into compute instance
  - Updating handling of new SAR dataset type by **PGE wrapper**

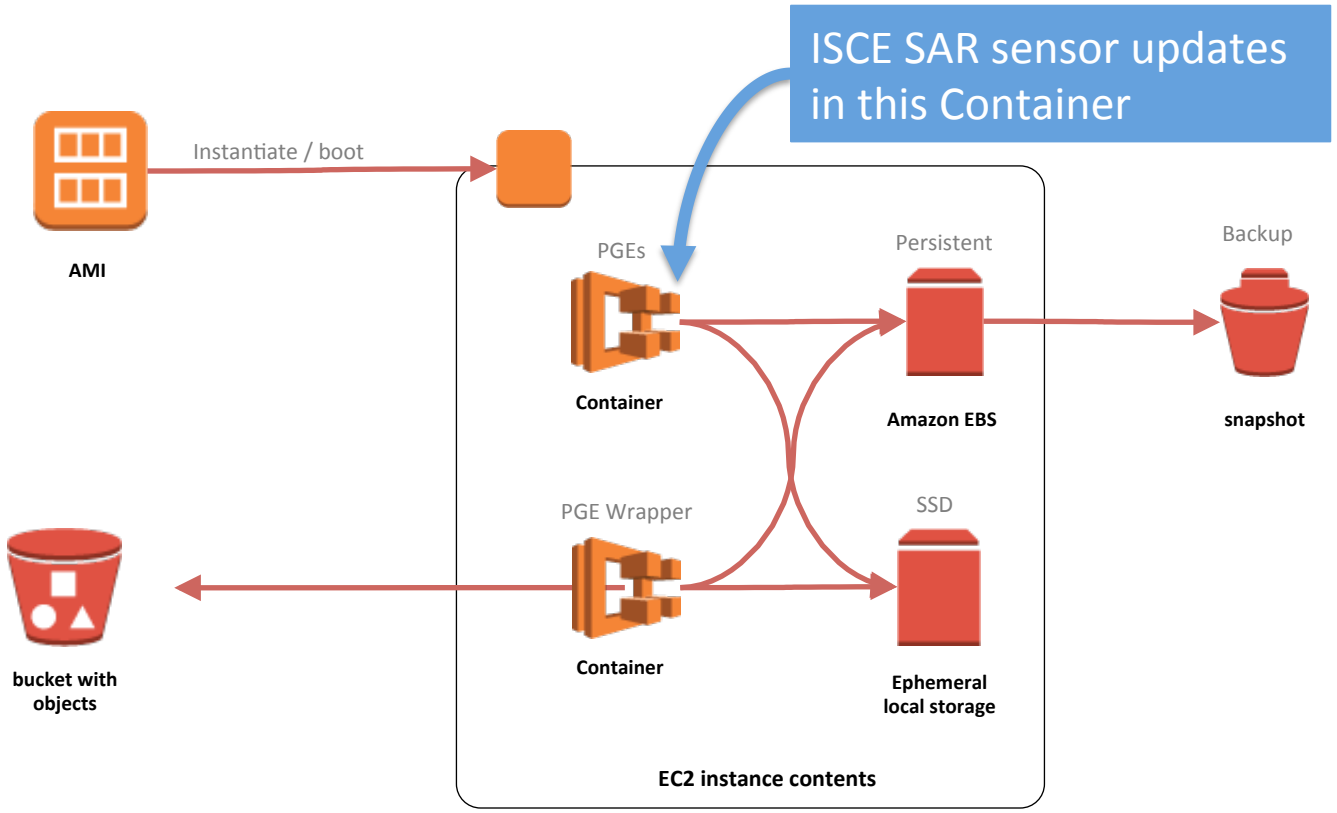

Consolidating Best Practices

LARGE SCALING CONSIDERATIONS

Auto-Scaling Science Data System

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

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

Auto-scaling tests over 100,000 x vCPUs

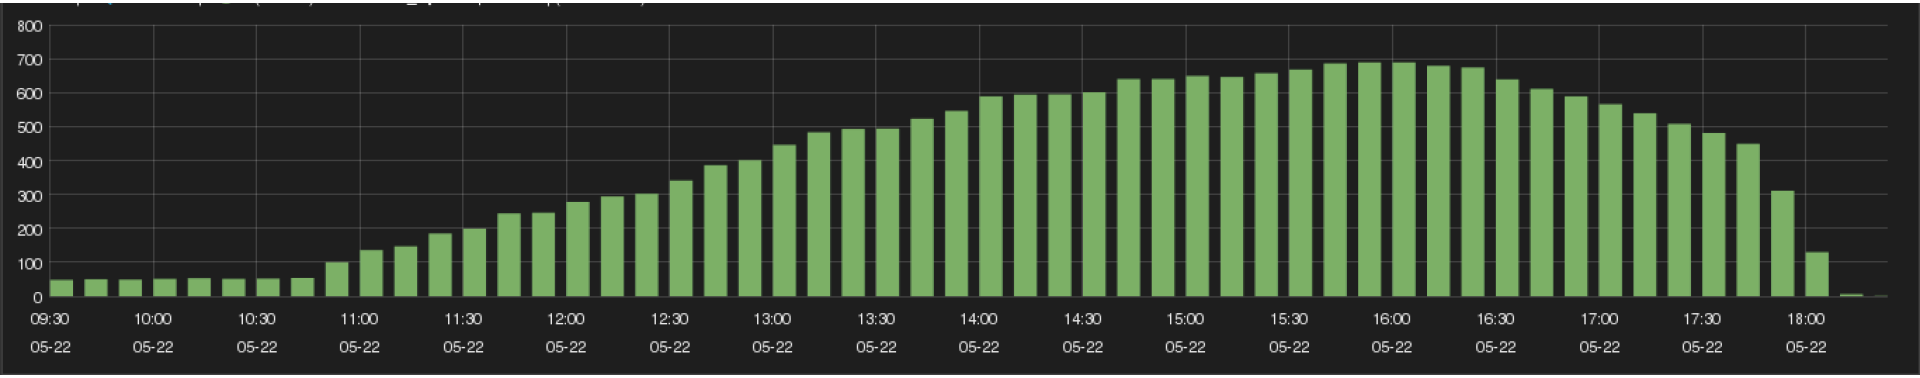

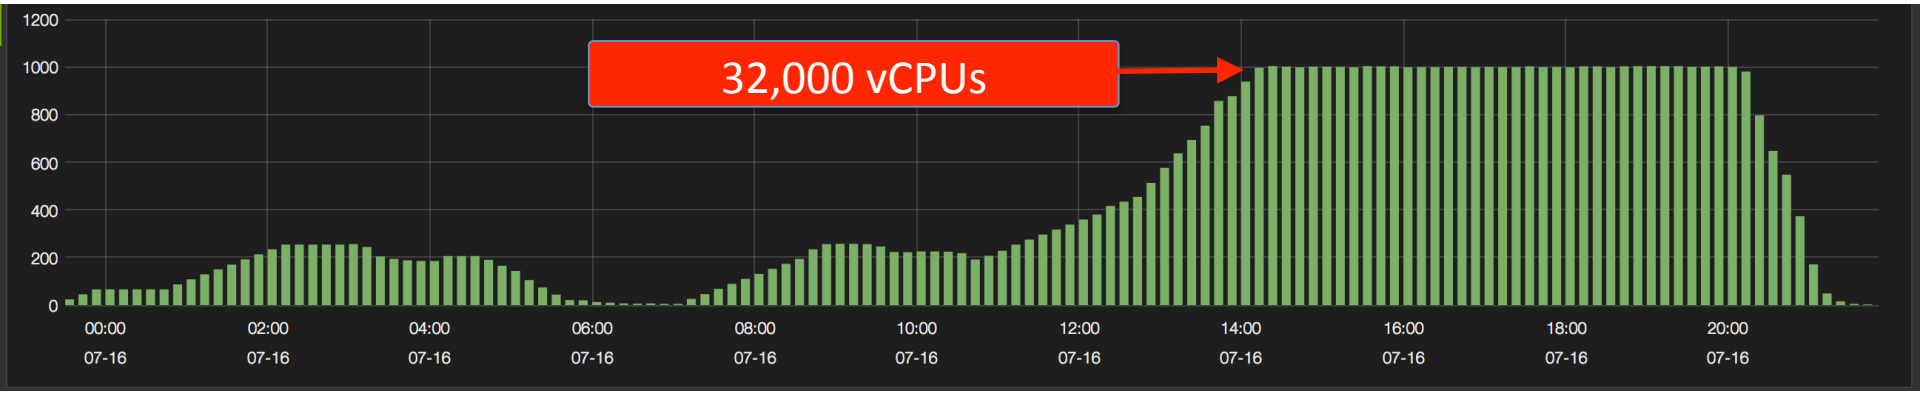

Availability Zone Load Rebalancing

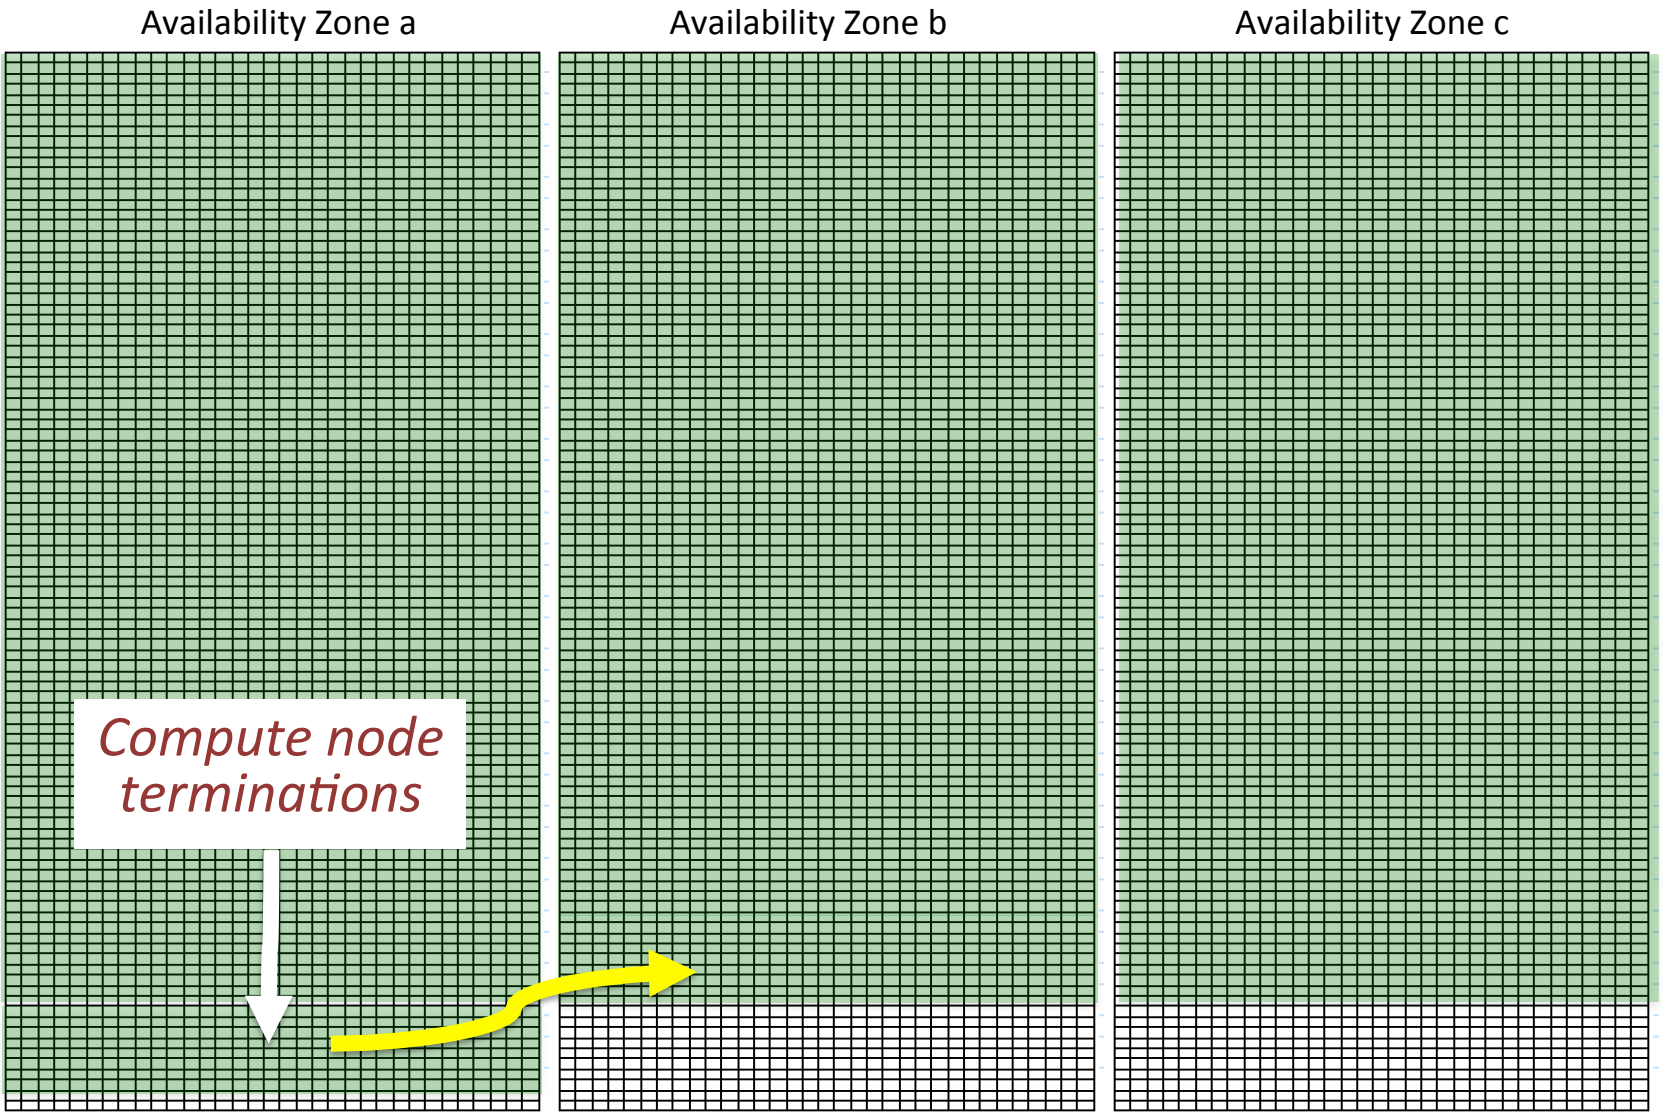

SLA Cloud Reliability

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

- *"The Amazon EC2 Service Level Agreement commitment is 99.95% availability for each Amazon EC2 Region."*
  - Expect about 1 failure per 2000 nodes
    - E.g. *"EC2 has detected degradation of the underlying hardware hosting one or more of your Amazon EC2 instances in the us-west-2 region. Due to this degradation, your instance(s) could already be unreachable."*
    - E.g. disk failures, network issues, etc.
- Plan for failures

"Thundering Herd"

- Fleet of compute instances hitting same services at same time
  - "*API rate limit exceeded*"
- Mitigation plan
  - Introduce **randomizations** to API calls
  - Distributes load on infrastructure
  - Move resources to S3/ S3-AI
- AWS already does some randomizations

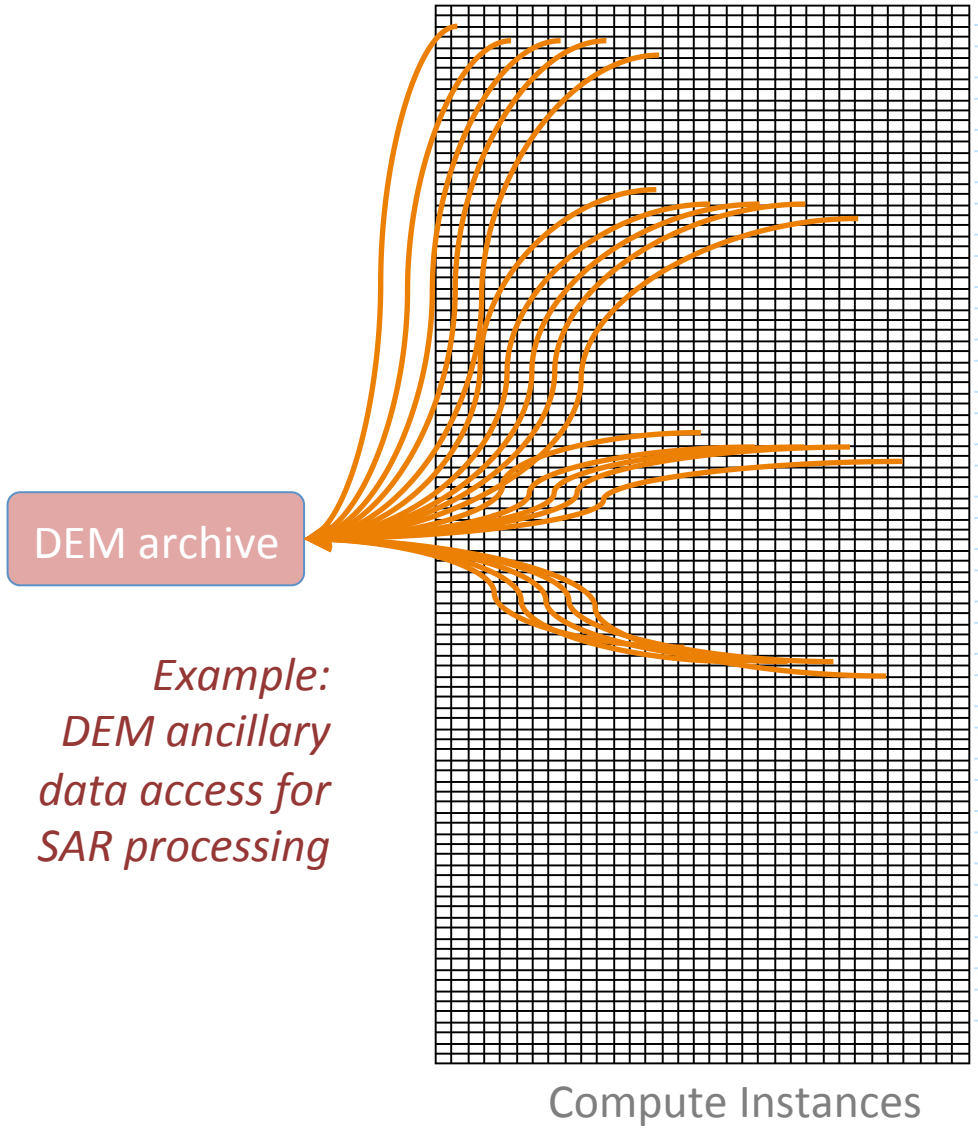

Example: DEM ancillary data access for SAR processing

More Cost-Efficient Use of Cloud Resources

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	$0.244	$0.1000	$0.1348	$0.0511	$0.030	$0.0125
cc2.8xlarge	32	60.5	1.89	4 x 840	$2.0000	$0.9131	$0.613	$0.2705	$0.0625	$0.0285	$0.019	$0.0085
m3.2xlarge	8	30.0	3.75	SSD 2 x 80	$0.6160	$0.3750	$0.230	$0.0700	$0.0770	$0.0469	$0.028	$0.0088
c3.8xlarge	32	60.0	1.88	SSD 2 x 320	$1.6800	$0.9920	$0.628	$2.4001	$0.0525	$0.0310	$0.019	$0.0750
r3.8xlarge	32	244.0	7.63	SSD 2 x 320	$2.8000	$1.4860	$0.982	$2.8000	$0.0875	$0.0464	$0.030	$0.0875
c3.xlarge	4	7.5	1.88	SSD 2 x 40	$0.2310	$0.1370	$0.087	$0.0353	$0.0578	$0.0343	$0.021	$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

High-Resiliency

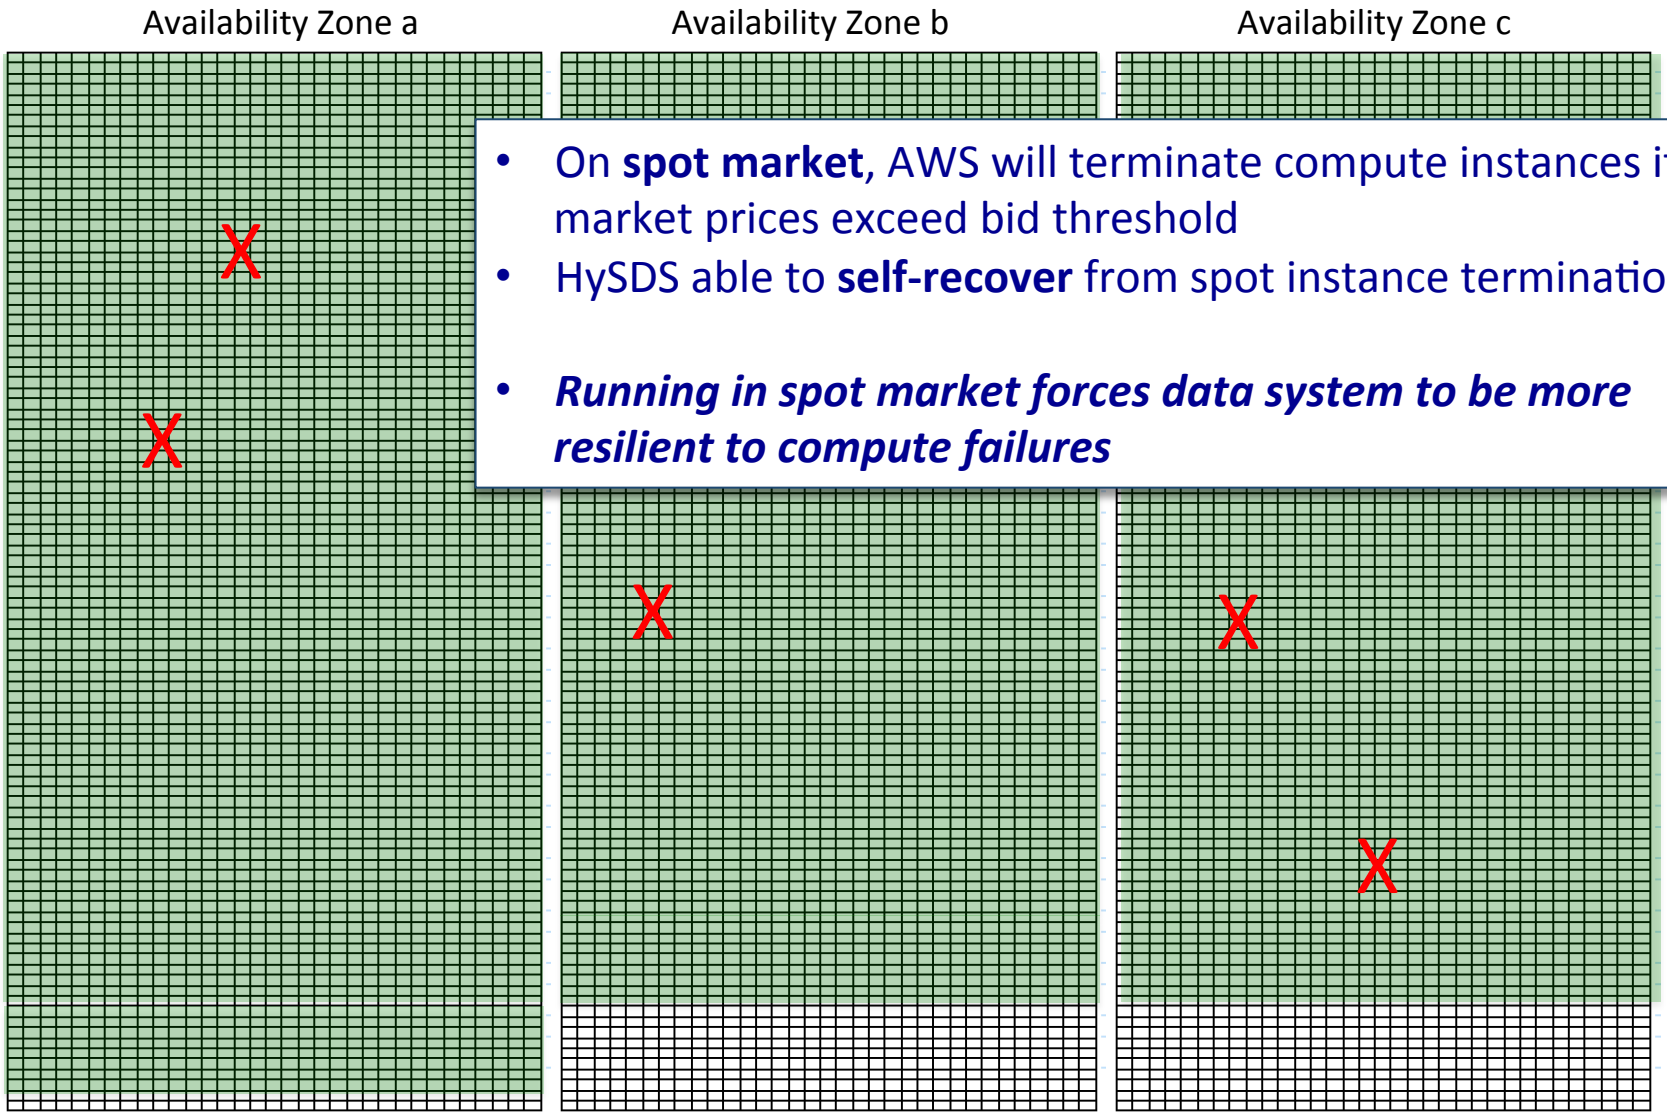

- On **spot market**, AWS will terminate compute instances if market prices exceed bid threshold
- HySDS able to **self-recover** from spot instance terminations
- ***Running in spot market forces data system to be more resilient to compute failures***

“Market Maker”

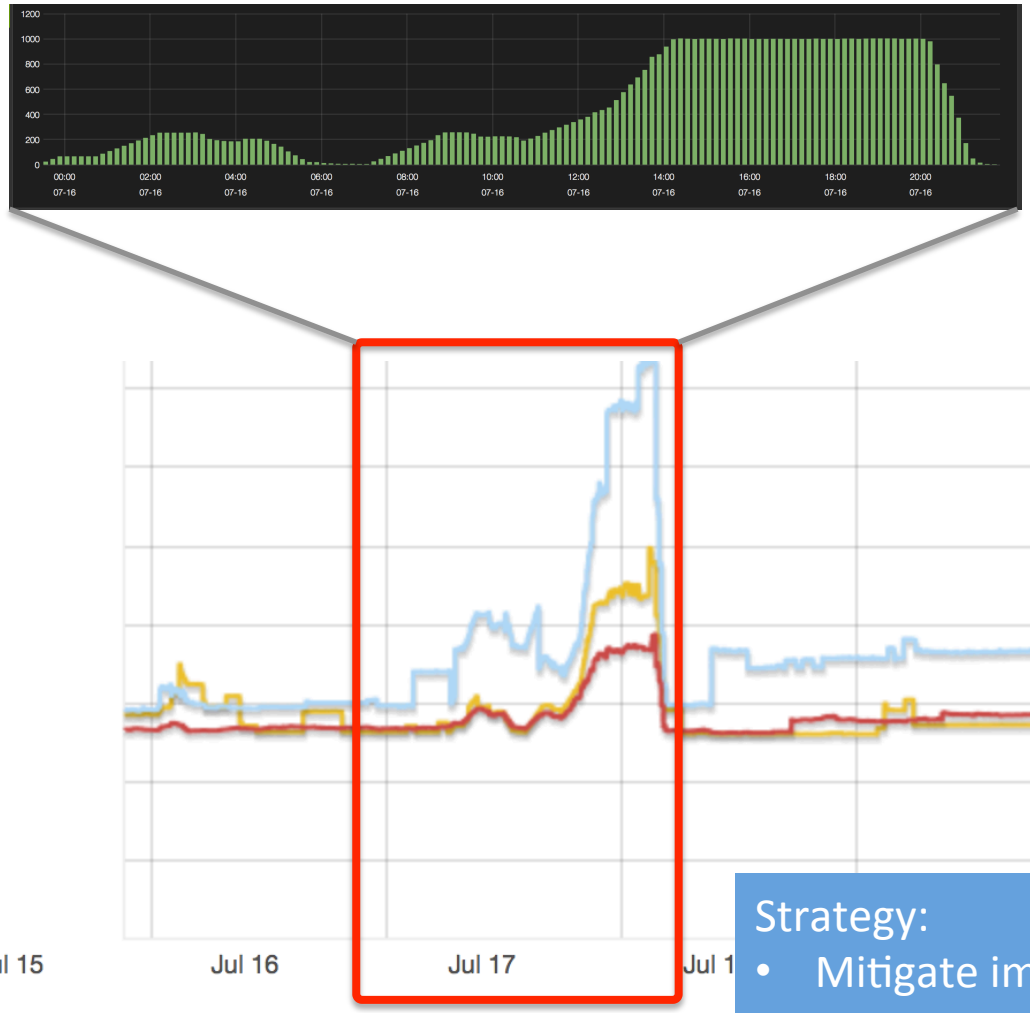

This production run of 32,000 vCPUs affected the market prices

Strategy:

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

“Other” Market Makers

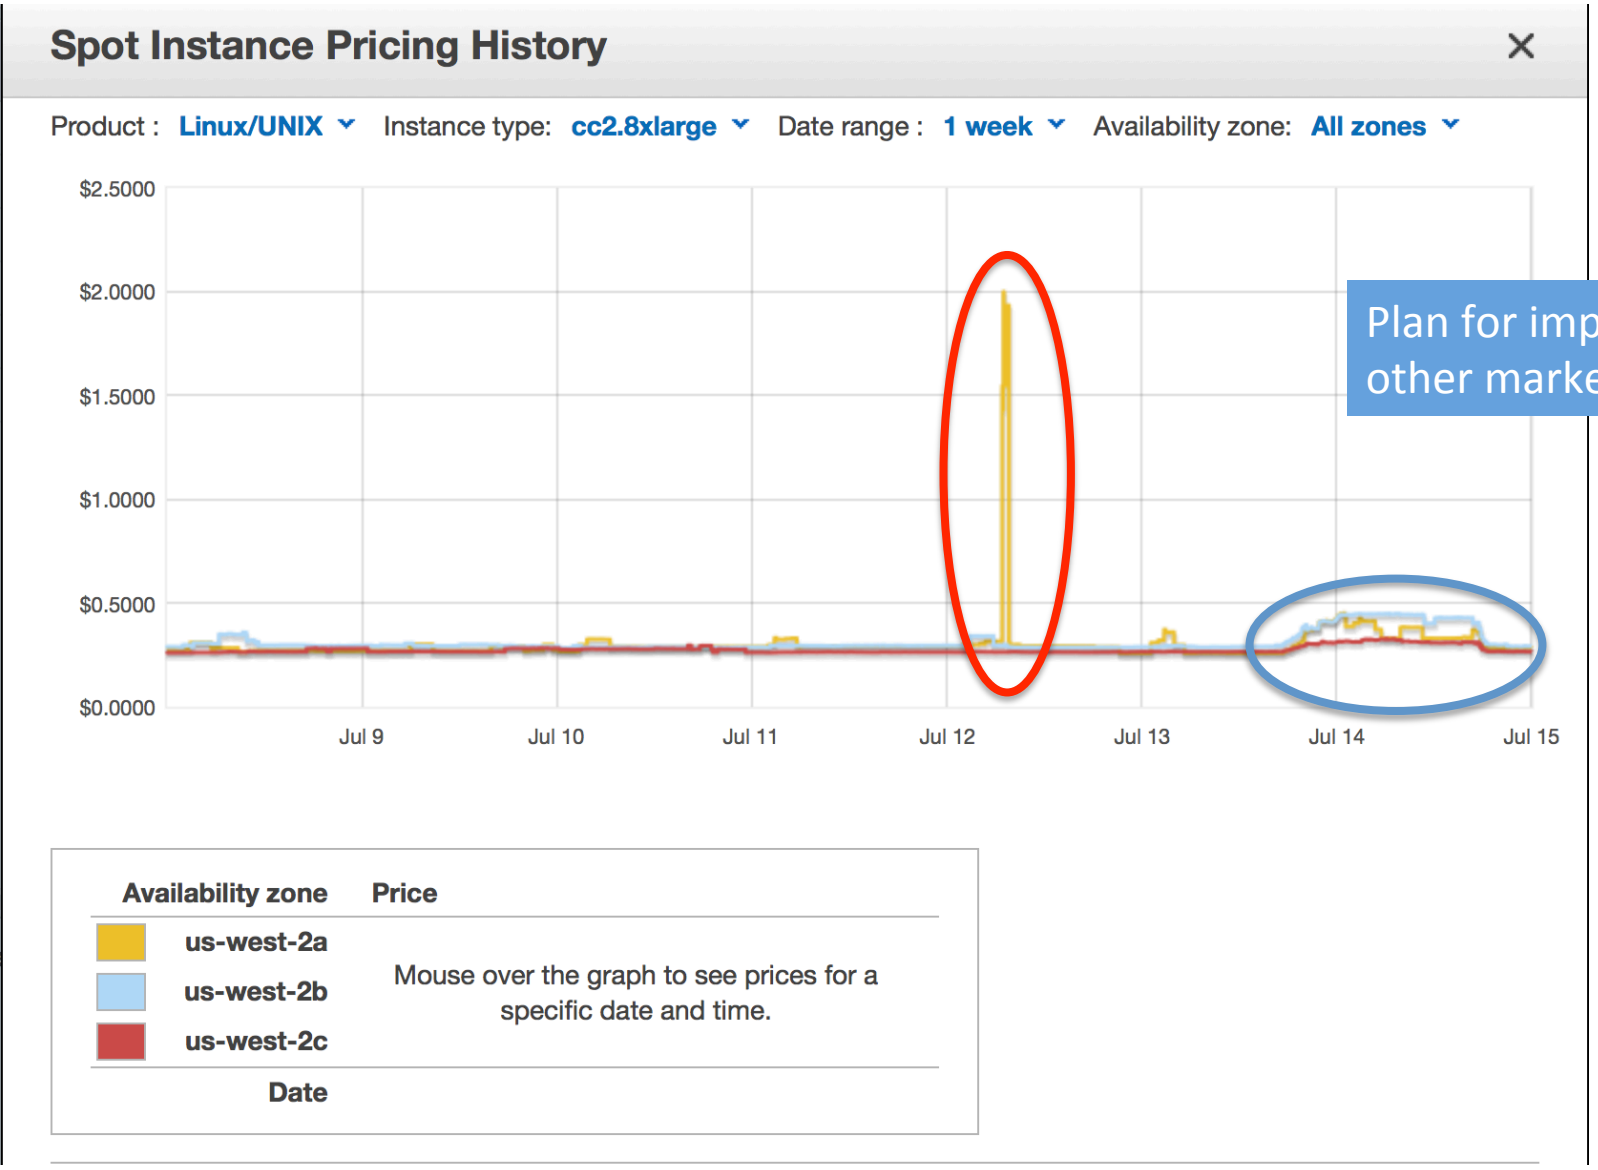

Plan for impacts by other market makers

COST MODEL

Example: Costing Sentinel-1A Processing

- Cost exercise example
  - Processing to 1000 x Sentinel-1A Level-2 phase unwrapped interferograms

For this estimate:

~$3 per L2 interferogram data product

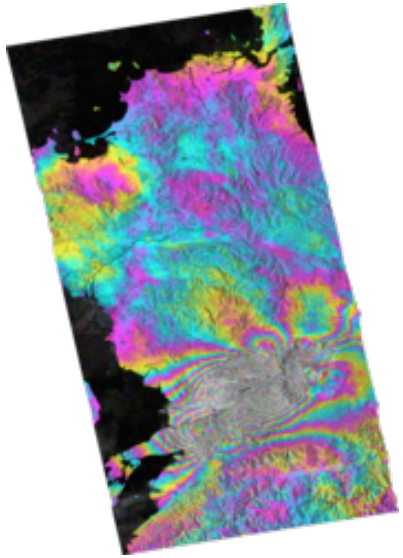

	Compute				
	per scene (m)	total (hr)	ec2 costs ($)	EBS storage (GB)	EBS costs ($)
L1 IW_SLC	5	83.3	$14.00	500	$6.94
L1 IW_SLC_SWATH	5	500.0	$84.00	500	$41.67
L2 interferogram	300	12500.0	$2,100.00	500	$1,041.67
	310.0	13083.3	$2,198.00	1500.00	$1,090.28

	Storage			Access			
	scenes	volume (GB)	s3-ia monthly costs ($)	data retrieved (%)	data egress (TB)	s3-ia data retrieva monthly costs ($)	egress monthly costs ($)
L1 IW_SLC	1000	283.0	$5.66	10%	0.03	$0.2	$4.39
L1 IW_SLC_SWATH	6000	1698.0	$33.96	10%	0.17	$1.7	$26.32
L2 interferogram	2500	6250.0	$125.00	100%	6.10	$62.5	$968.75
	9500	8231.0	$164.62		184.63	$64.4	$999.46

Why is this important?

IMPACTS

Analysis for California Department of Water Resources (DWR)

- California Drought Causing Valley Land to Sink
  - *"Sinking land, known as **subsidence**, has occurred for decades in California because of excessive groundwater pumping during drought conditions, but the new NASA data show the sinking is happening faster, putting infrastructure on the surface at growing risk of damage."*
- DWR interest in processing all of the Sentinel data for CA
  - Over 5500 Sentinel-1A scenes processed to L2 interferograms
- Tom Farr (JPL geologist)'s DWR-based funding paying for processing under **SAR SDP Foundry**
  - **Sentinel-1A on-ramped**
  - **UAVSAR on-ramped**

Areas of Interests for DWR study

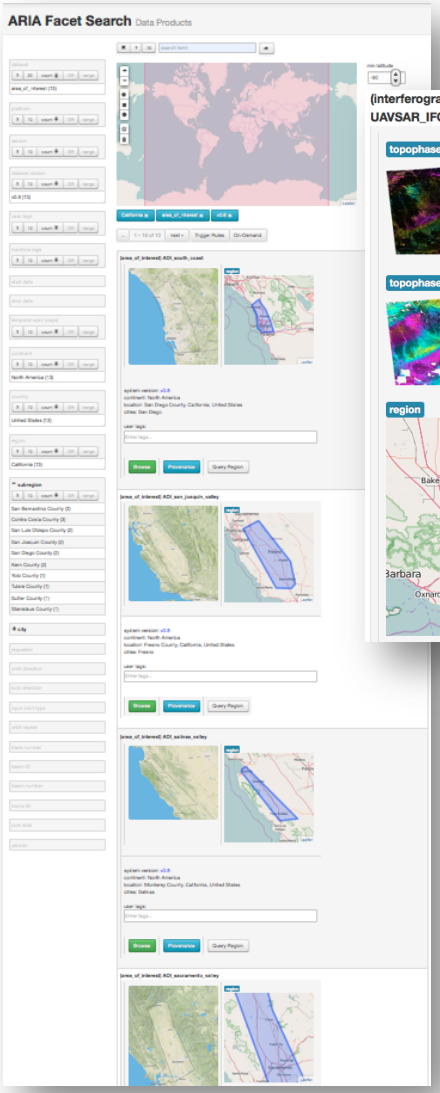

On-ramped UAVSAR Interferogram processing

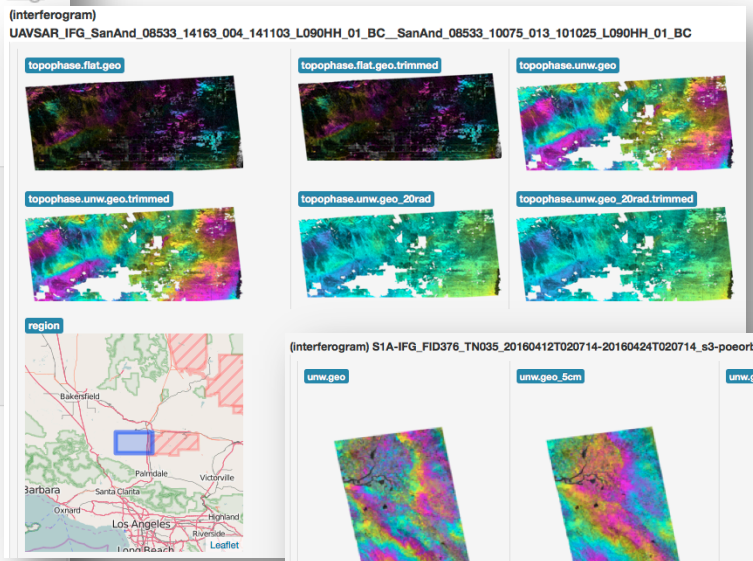

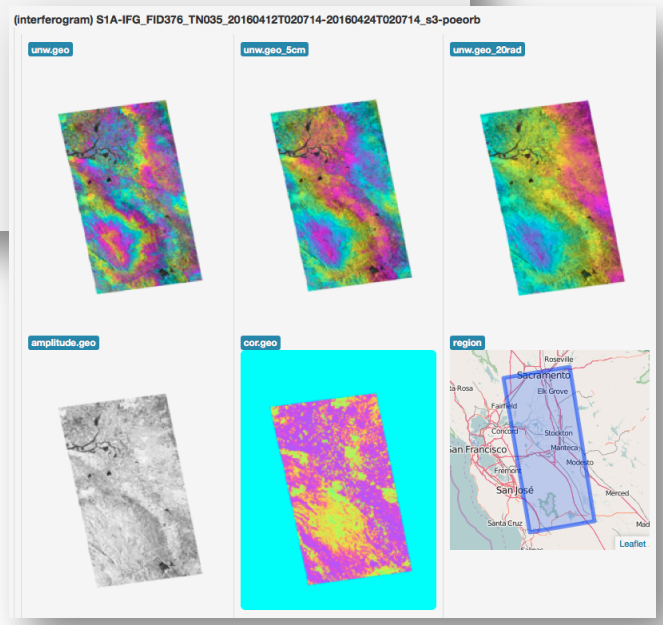

On-ramped Sentinel-1A Interferogram processing

Impacts to NISAR Level-3 Science Data Product Estimates

- **Foundry models** were useful in costing out various deployment scenarios for **on-demand** NISAR Level-3 science data products
  - Solid earth, biomass, ecosystems, and cryosphere
- A more detailed **monthly cost model** was developed
- Explored options for
  1. Storing L1 SLCs to produce L3
  2. Store only L0B, dynamically generate on-demand L1 SLCs to L3
  3. Pull L1 SLCs from DAAC to produce L3

SDS dependency on DAAC	Options	Cost Estimate
None	Store all needed SLCS (2.5PB) and produce L3	High
None	Store L0B, produce SLCs and produce L3	Medium
Ingest SLCs from DAAC	Ingest SLCs from DAAC and produce L3	Low
DAAC Co-located (SDS has direct access to DAAC storage)	Use SLC's on DAAC and produce L3	Lower

NASA SAR Science Data Processing (SDP) Foundry Enables Science

L-Band
P-Band
(JPL)

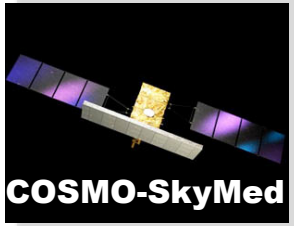

X-Band
(Italy)

P-Band
(GSFC)

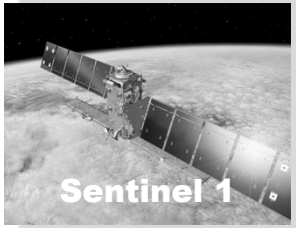

C-Band
(ESA)

X,C,S,L,P-Bands
(DLR)

The SAR SDP Foundry

ARIA
JPL-Caltech

Advanced Rapid Imaging & Analysis for Monitoring Hazards

ISCE

InSar Scientific Computing Environment on the Cloud

amazon web services

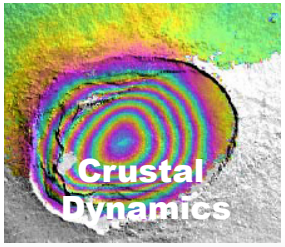

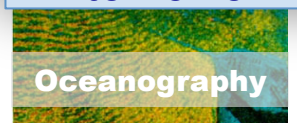

SAR data from NASA and other sources can be on-ramped into the cloud-based processing data system, the **SAR Science Data Processing (SDP)**

ESTF 2016

- This talk focused on the **ARIA science data system** behind the SAR SDP Foundry.
- Eric Gurrola will be presenting on the **InSAR Scientific Computing Environment (ISCE)** component tomorrow @ 10:00am

Once on-ramped, PI's can quickly process data they select, to the extent that they can afford it.

BACKUP

High-Level Cloud Approach for Science Data Processing

- **Stream input data into AWS S3 object storage**
- **Scale up compute nodes to run in AWS EC2**
- **Internal** SDS data throughput needs are scalable via cloud architecture
  - **Compute** instances can scale up to demand
  - Object storage can scale up **data volume** and **aggregate data throughput** to demand
- Asynchronously move results from AWS S3 back to on-premise facility
- Architectural components can be **collocated**

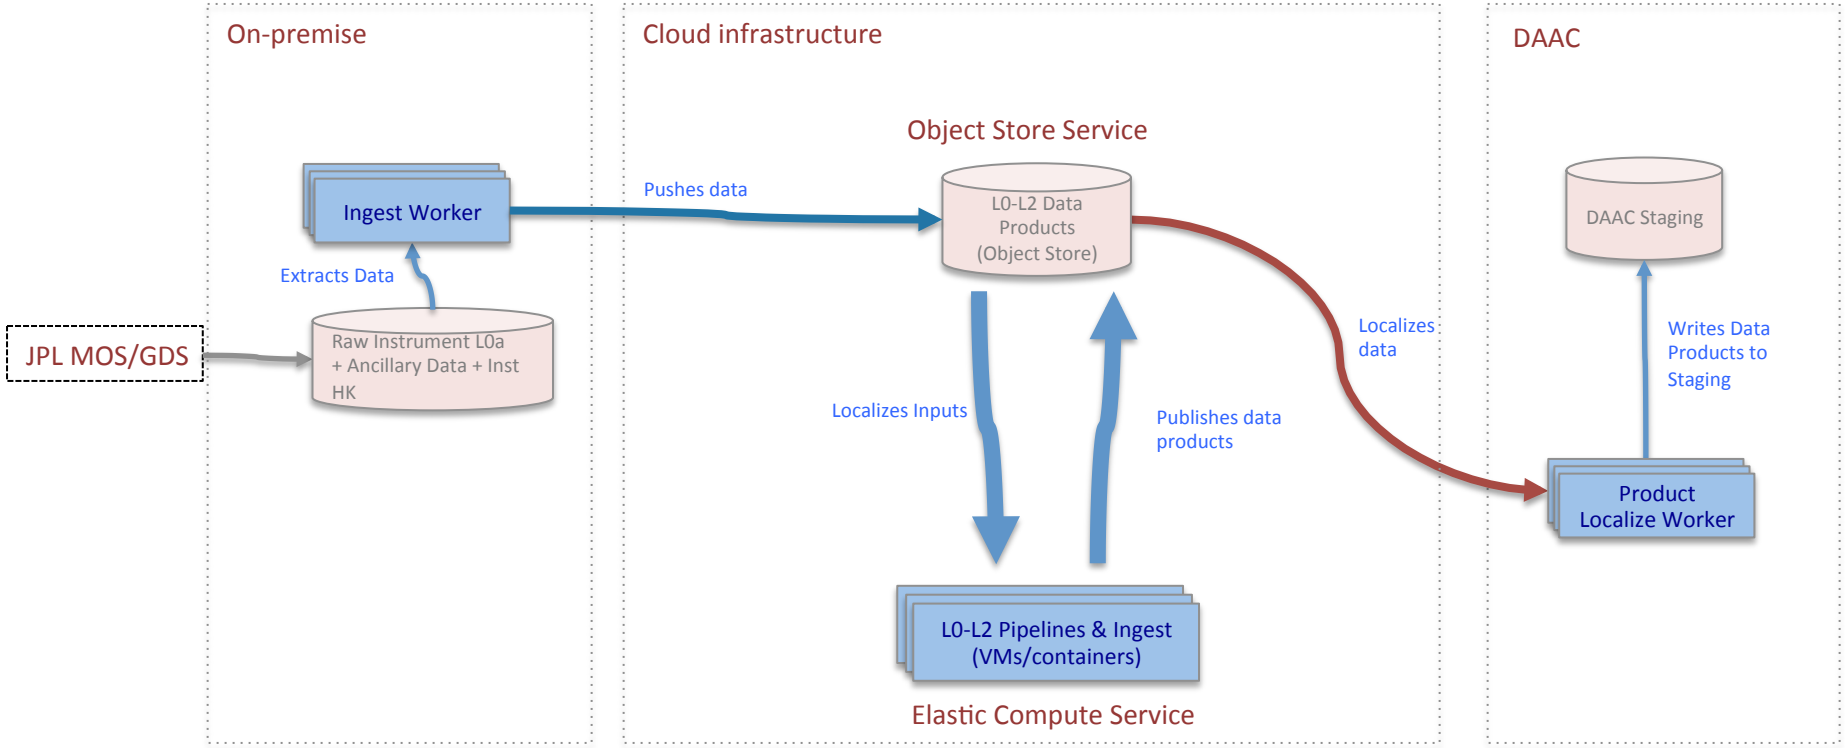

“Data Lake”-extended

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

Reduce redundancy and foster ESDIS-wide services

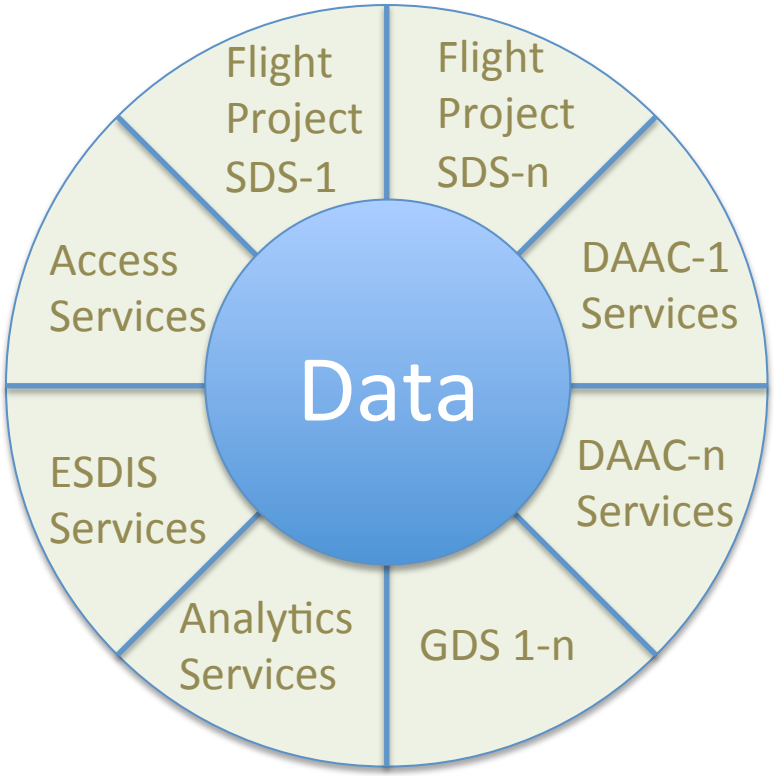

Enabling multi-disciplinary data approach for analysis

"Hot Data Lake"

- Long-term **public cloud storage is expensive** at PB-scales
  - …Unless we can negotiate deals with cloud vendors
- Use object store data lake for ***"hot data"***
  - SDSes generate data into object stores
  - Object stores contain "fresh" / "hot" data as rolling storage
  - Offline moving of older data to DAAC for permanent storage
  - Automate caching of "hot data" back from DAACs