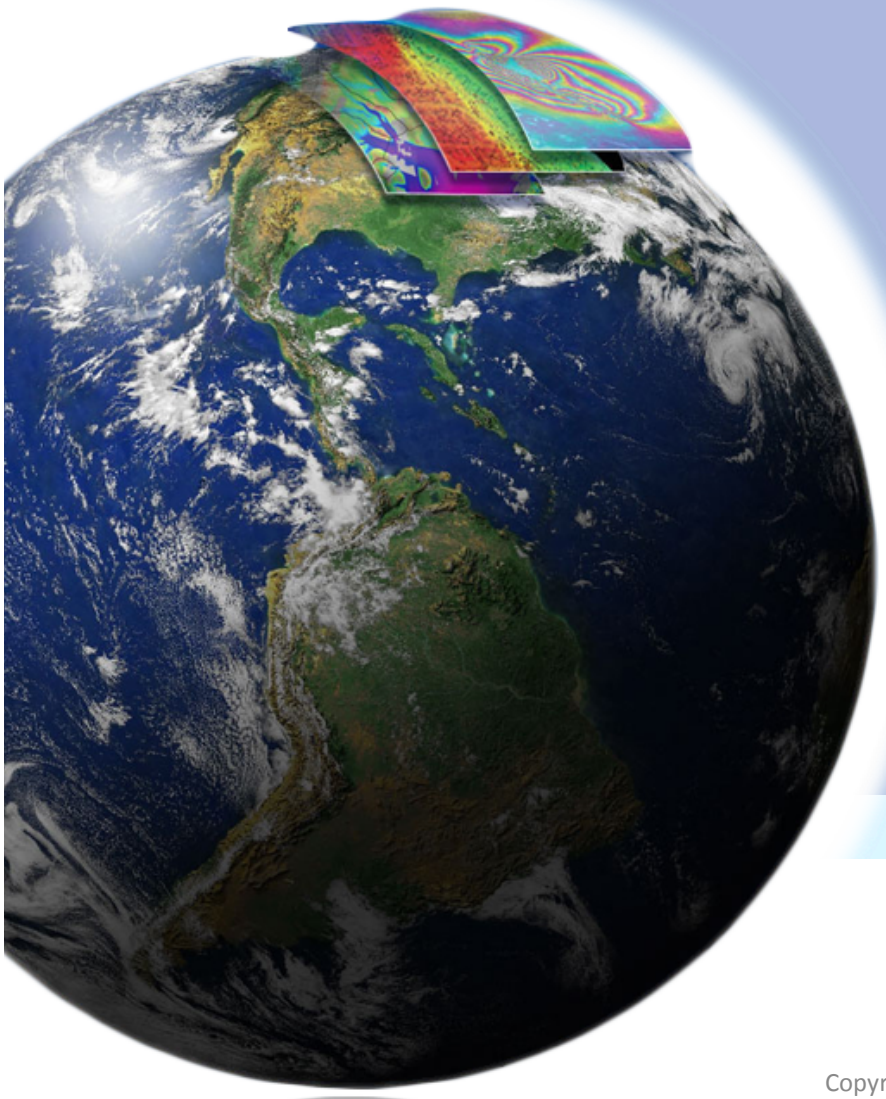

SDS Considerations for SWOT and NISAR

Big Data Task Force of the
NASA Advisory Council Science Committee
Thursday, November 2, 2017

Hook Hua (SDS Architect, SWOT and NISAR)
Jet Propulsion Laboratory, California Institute of Technology.

SWOT Mission Concept

Surface Water and Ocean Topography (SWOT)

Oceanography: Characterize the ocean mesoscale and sub-mesoscale circulation at spatial resolutions of 15 km and greater.

Hydrology: To provide a global inventory of all terrestrial water bodies whose surface area exceeds $(250m)^2$ (lakes, reservoirs, wetlands) and rivers whose width exceeds 100 m (rivers).

- To measure the global storage change in fresh water bodies at sub-monthly, seasonal, and annual time scales.
- To estimate the global change in river discharge at sub-monthly, seasonal, and annual time scales.

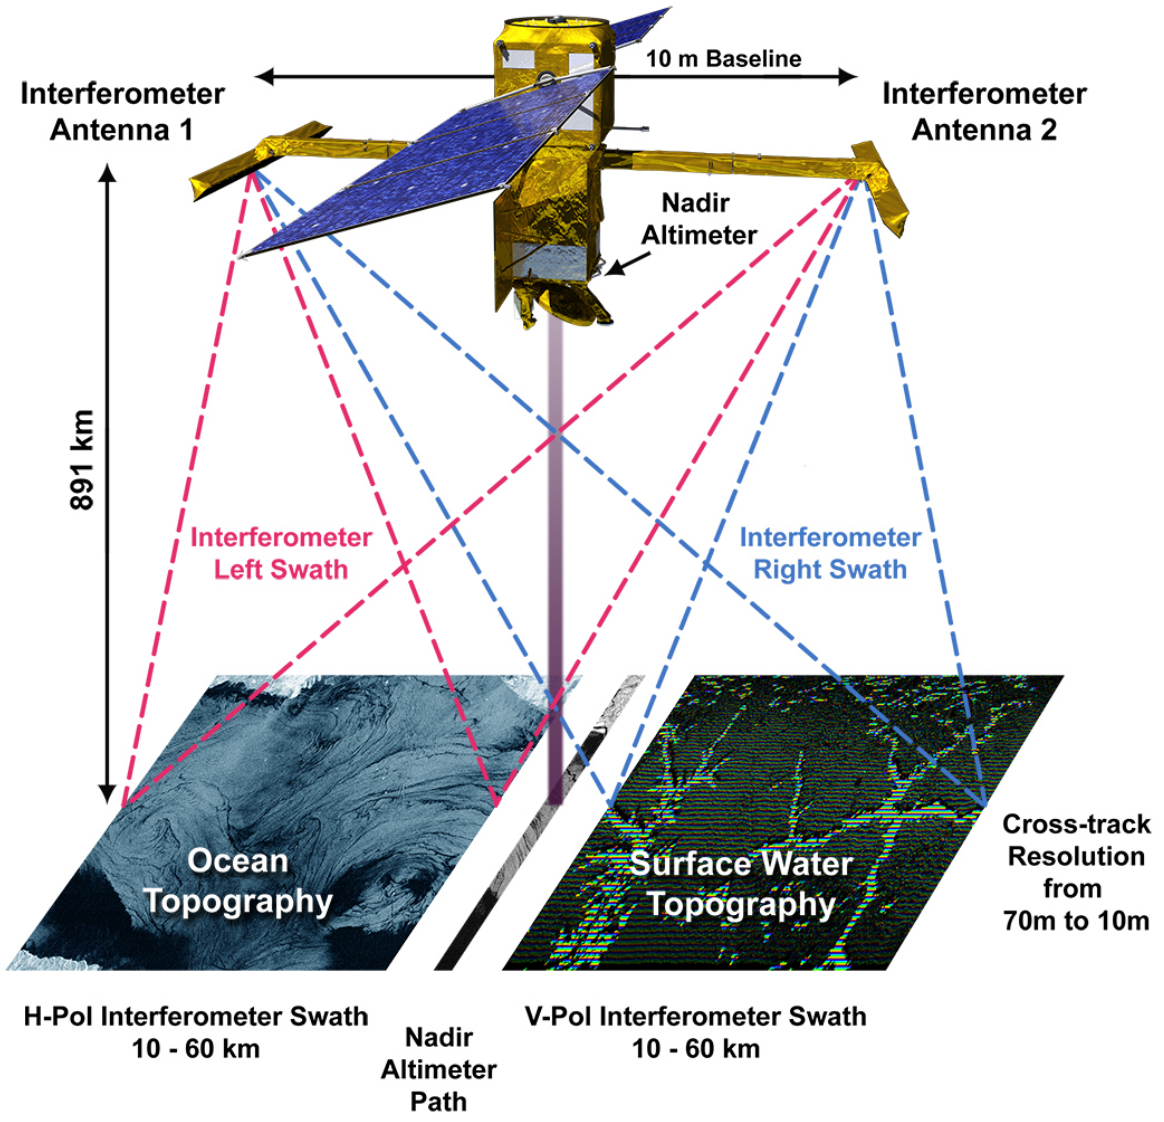

NISAR Mission Concept

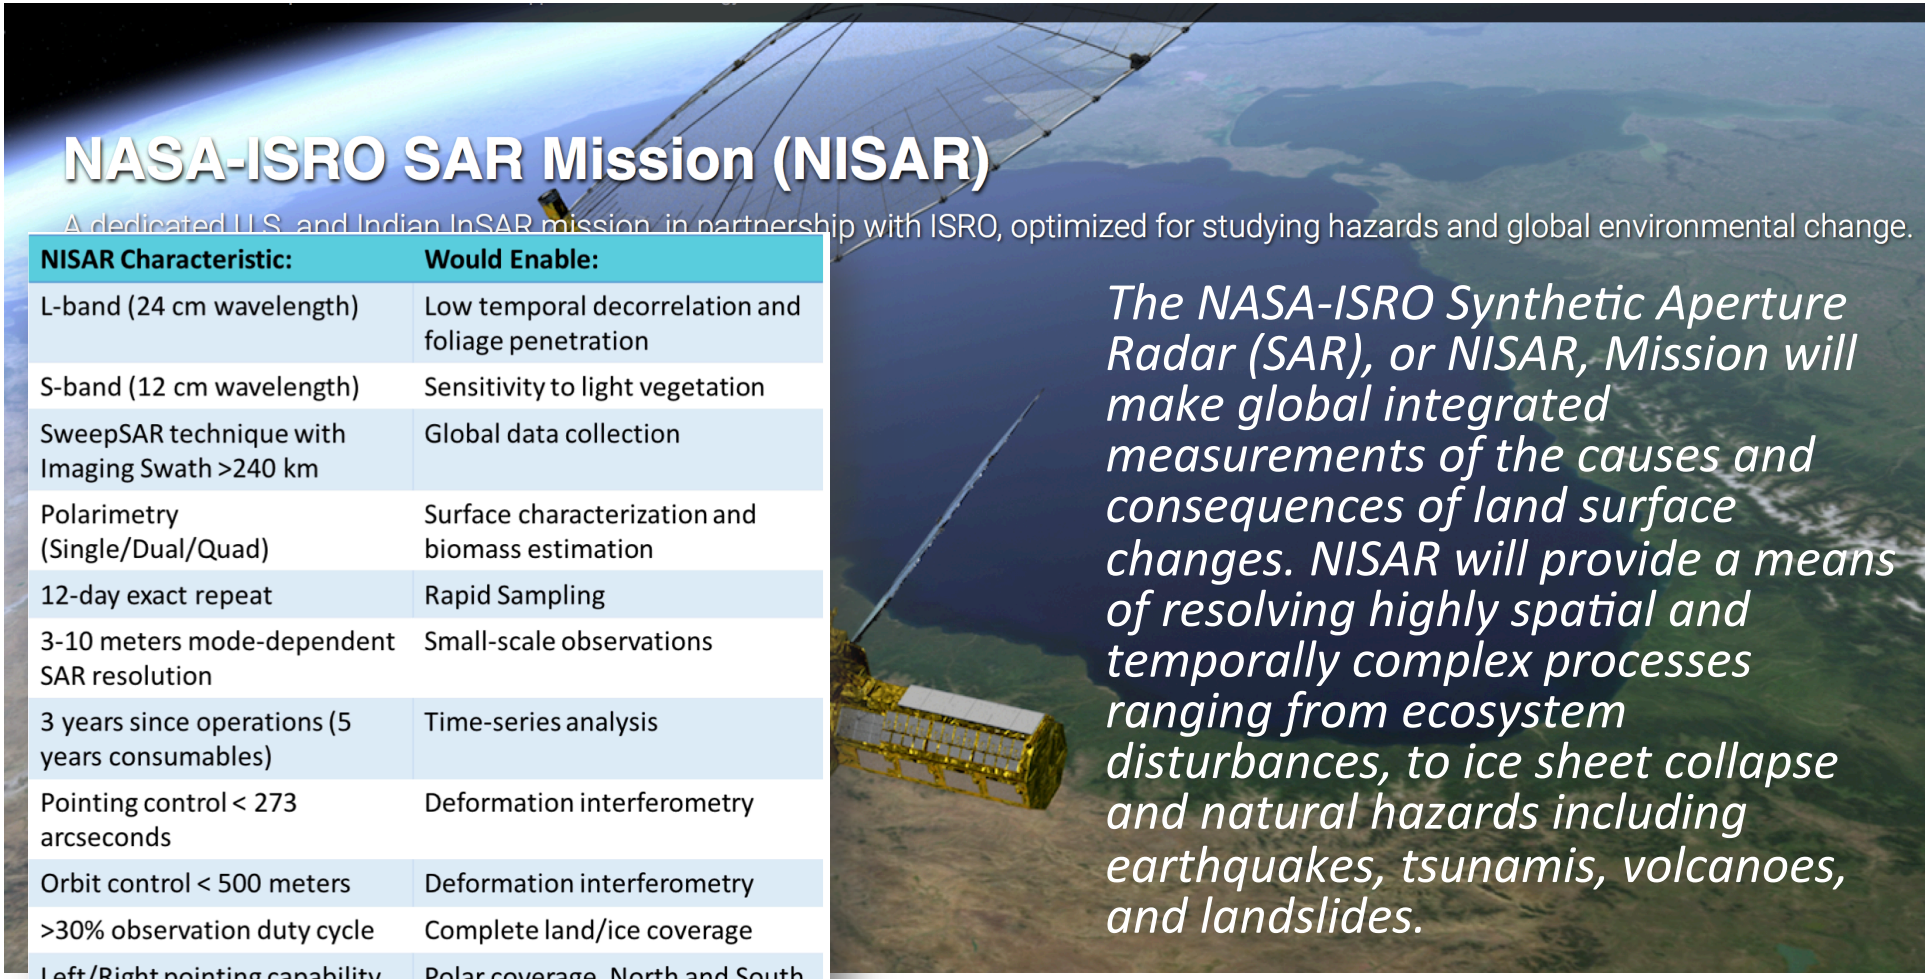

NISAR Characteristic:	Would Enable:
L-band (24 cm wavelength)	Low temporal decorrelation and foliage penetration
S-band (12 cm wavelength)	Sensitivity to light vegetation
SweepSAR technique with Imaging Swath >240 km	Global data collection
Polarimetry (Single/Dual/Quad)	Surface characterization and biomass estimation
12-day exact repeat	Rapid Sampling
3-10 meters mode-dependent SAR resolution	Small-scale observations
3 years since operations (5 years consumables)	Time-series analysis
Pointing control < 273 arcseconds	Deformation interferometry
Orbit control < 500 meters	Deformation interferometry
>30% observation duty cycle	Complete land/ice coverage
Left/Right pointing capability	Polar coverage, North and South
Noise Equivalent Sigma Zero ≤ -23 db	Surface characterization of smooth surfaces

Figure 1 – NISAR radar characteristics, as of Oct. 2015.

Next Generation Earth Science Remote Sensing Missions

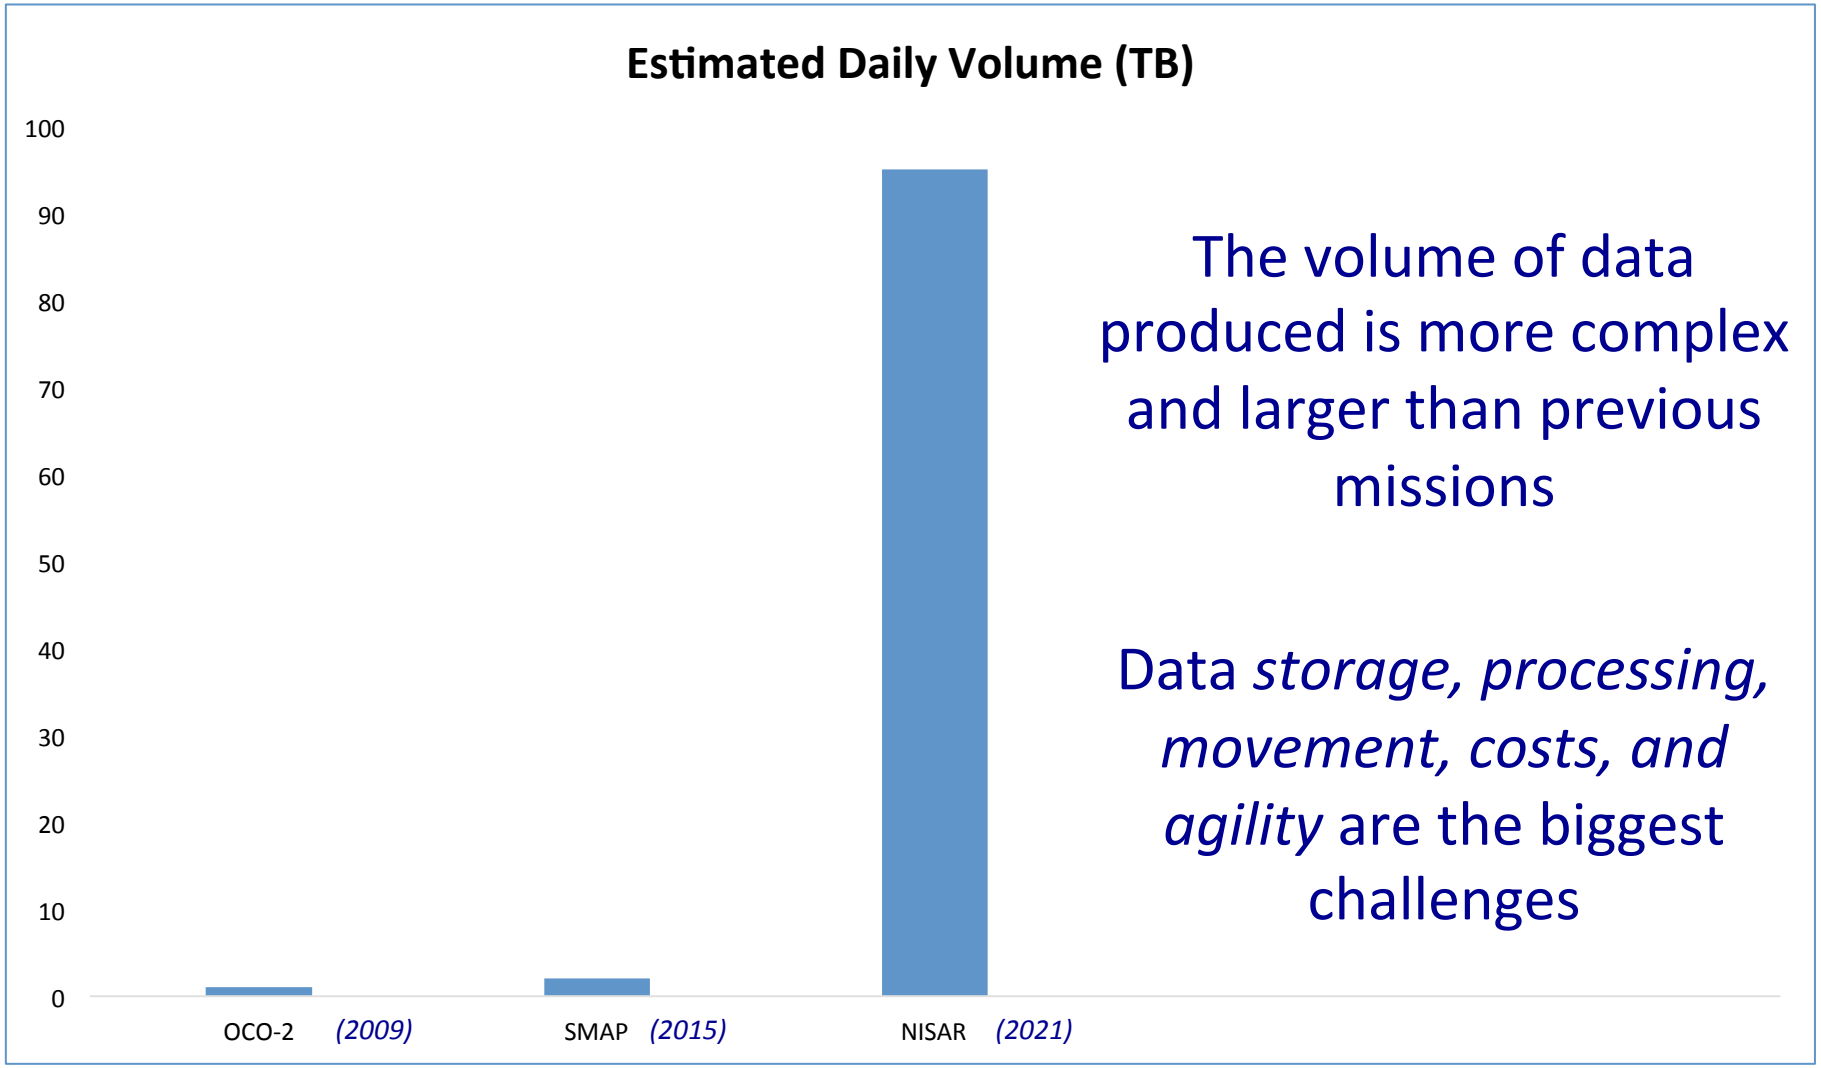

The volume of data produced is more complex and larger than previous missions

Data *storage, processing, movement, costs, and agility* are the biggest challenges

Comparison of Flight Project SDSes

	SWOT SDS	NISAR SDS	SMAP SDS	OCO-2 SDS
Daily Data Acquisition Volume	1 TB	3.25 TB	0.14 TB	0.01 TB
Daily Production Volume	15.5 TB	86 TB	0.44 TB	0.06 TB
L1 Product Latency	29 days	30 days	12 hrs.	2 days
PGE count	> 20	>10	28	17
Ancillary data types	~11	~14	@ Ops:133	@ Ops: 14
Complexity of Ops Workflow	High	Medium	High	Low

Science users will likely need to adopt new strategies for interacting with SWOT and NISAR's high-volume data

Large Data Flow Networks

- GDS handles global downlinks and ground network
- SDS generates science data products
- DAACs provide access, storage, and services to end-users

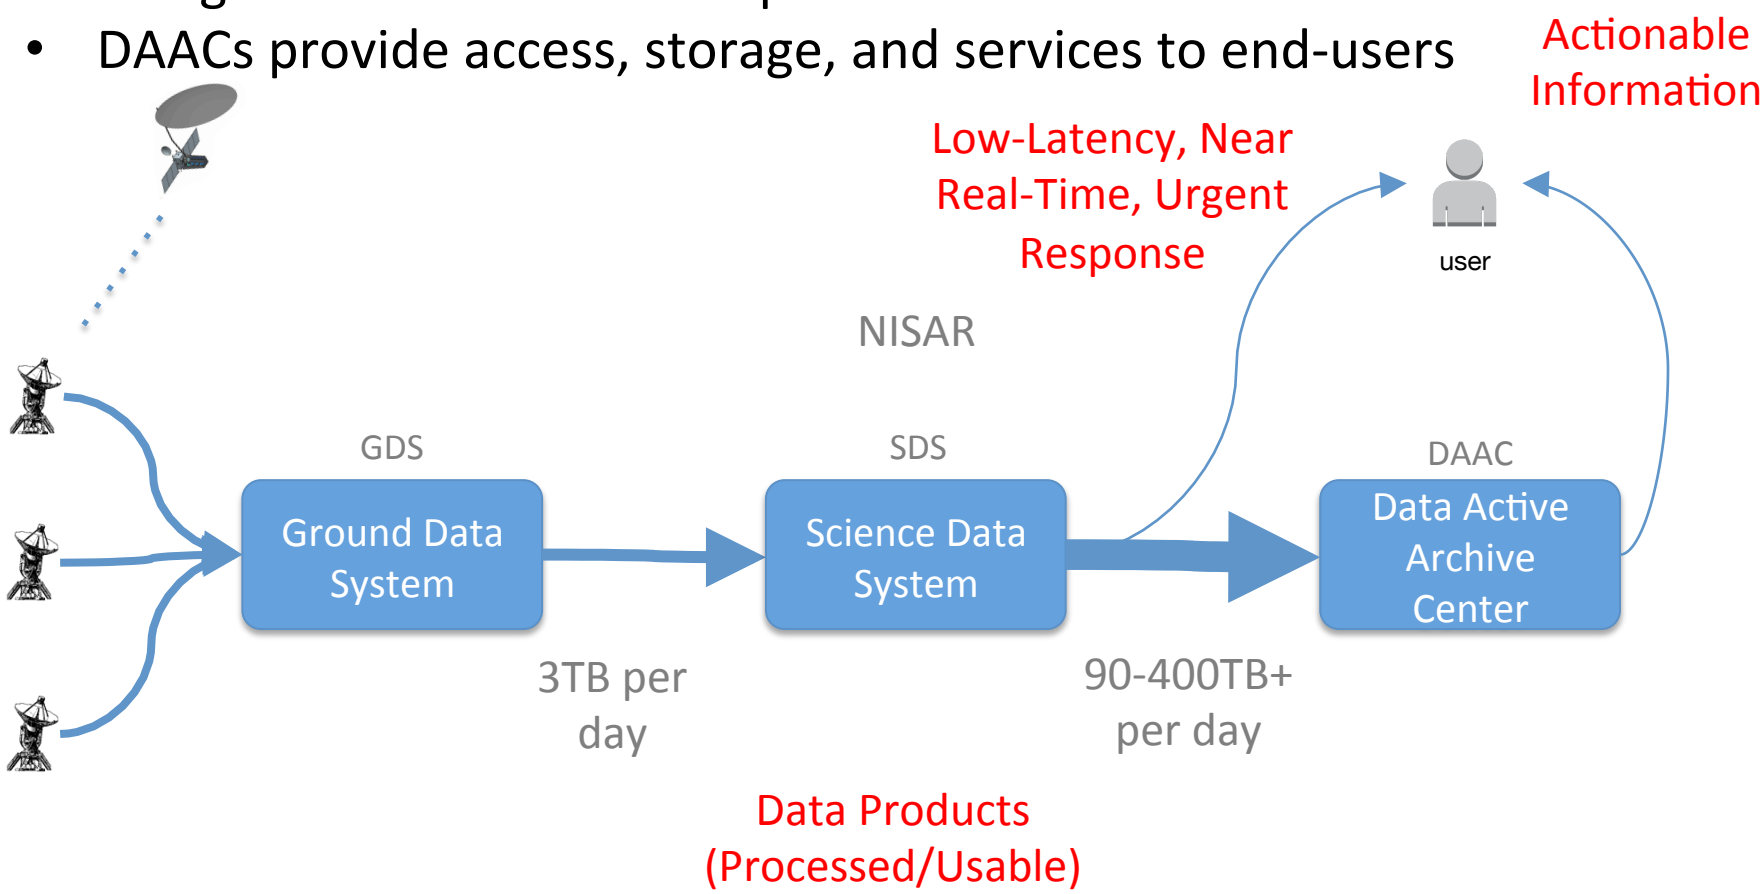

Concurrent Processing Pipelines

- Concurrently keep up with:
  - Forward processing (“keep up”)
  - Bulk (re)processing
  - Urgent response
  - NRT
  - On-demand analysis

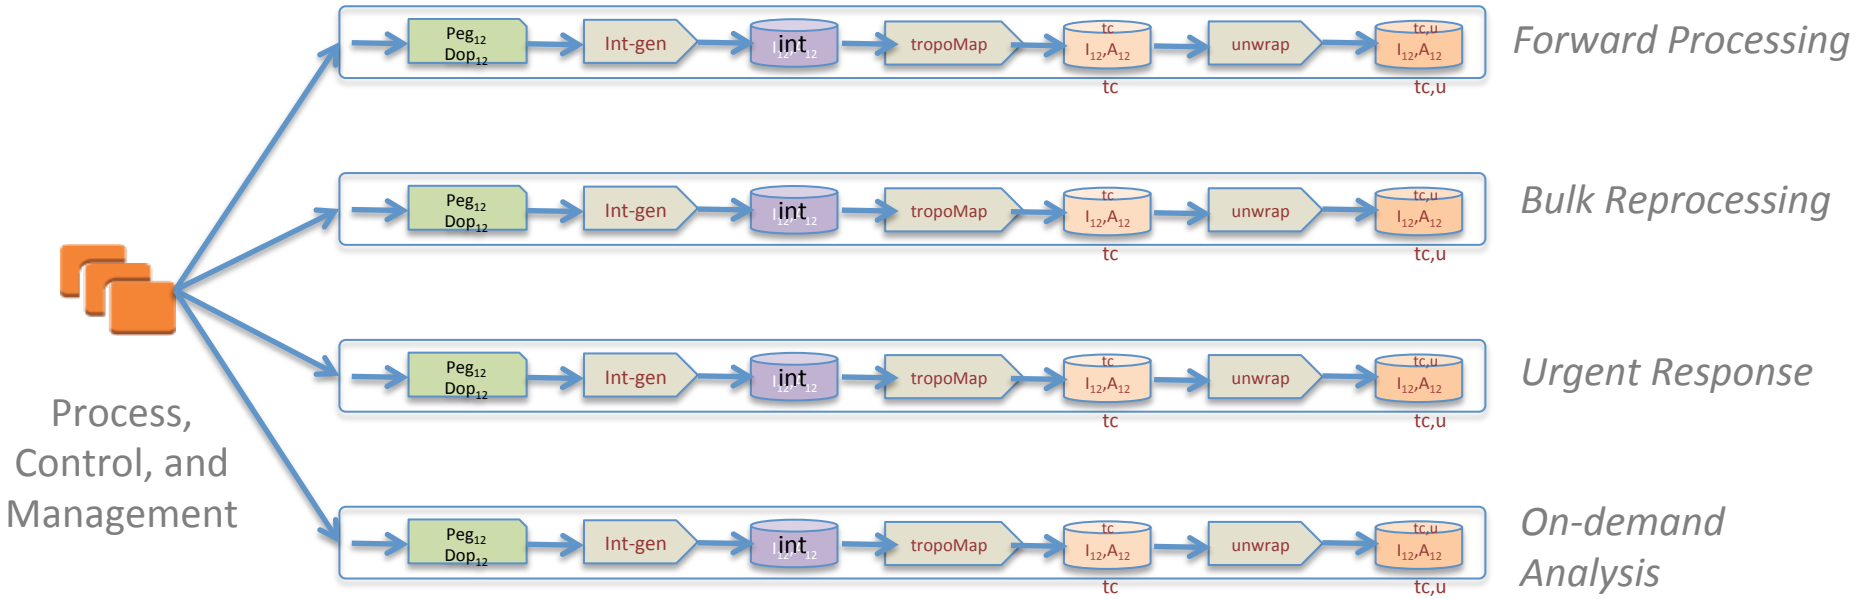

Processing Pipelines

Basic Premise of Cloud-based Science Data Processing

- Science data product into AWS S3 **object storage**
- **Scale up** compute nodes to run in AWS EC2
- **Internal** SDS data throughput needs are scalable via **cloud architecture**
  - Object storage can scale up **data volume** and **aggregate data throughput** by compute instances
- Architectural components can be **collocated**

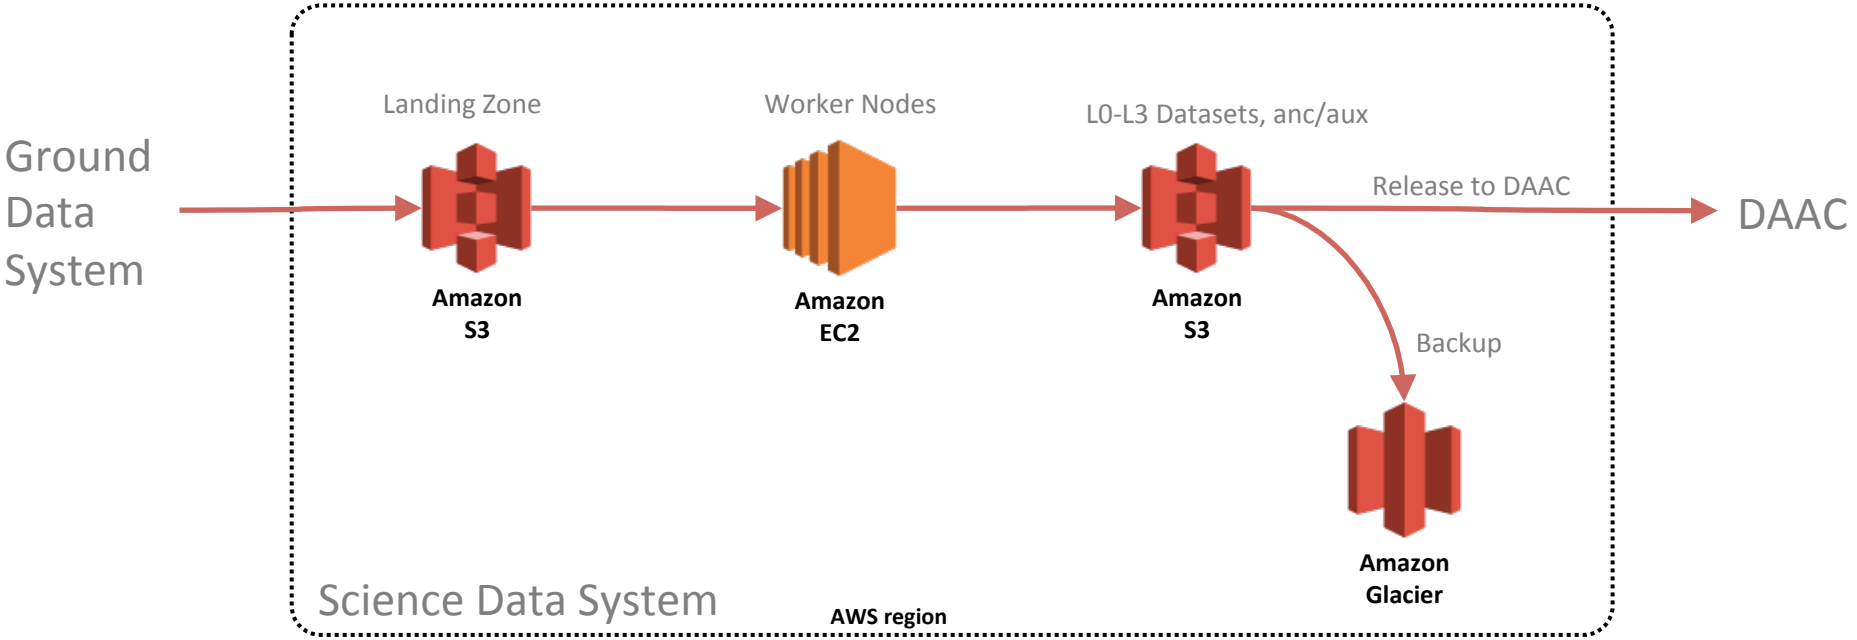

NISAR Science Data Products

Product	To DAAC
L0	
L0A (catalog incoming raw data)	3.25
L0B Radar Signal Data	3.25
L1	
Range-Doppler Single-Look Complex (SLC)	30.33
Multi-Look Detected Browse (MLD)	0.54
L2 (all modes)	
Geocoded Single-Look Complex	30.33
Interferogram (nearest-time pair)	8.67
Unwrapped Interferogram	4.33
Geocoded Unwrapped Interferogram	4.33
L2 Ecosystem (Quad)	
Polarimetric Covariance Matrix	0.43
Geocoded Polarimetric Covariance Matrix	0.43
Total (TB/day, uncompressed)	**85.91**

- Current Baseline Products
- L1 SLCs and L2 geocoded SLCs comprise the **bulk of the data volumes**
- Data **locality** and caching impacts to reprocessing
- **GPU-accelerated** data product generation

Forward Processing in AWS Cloud

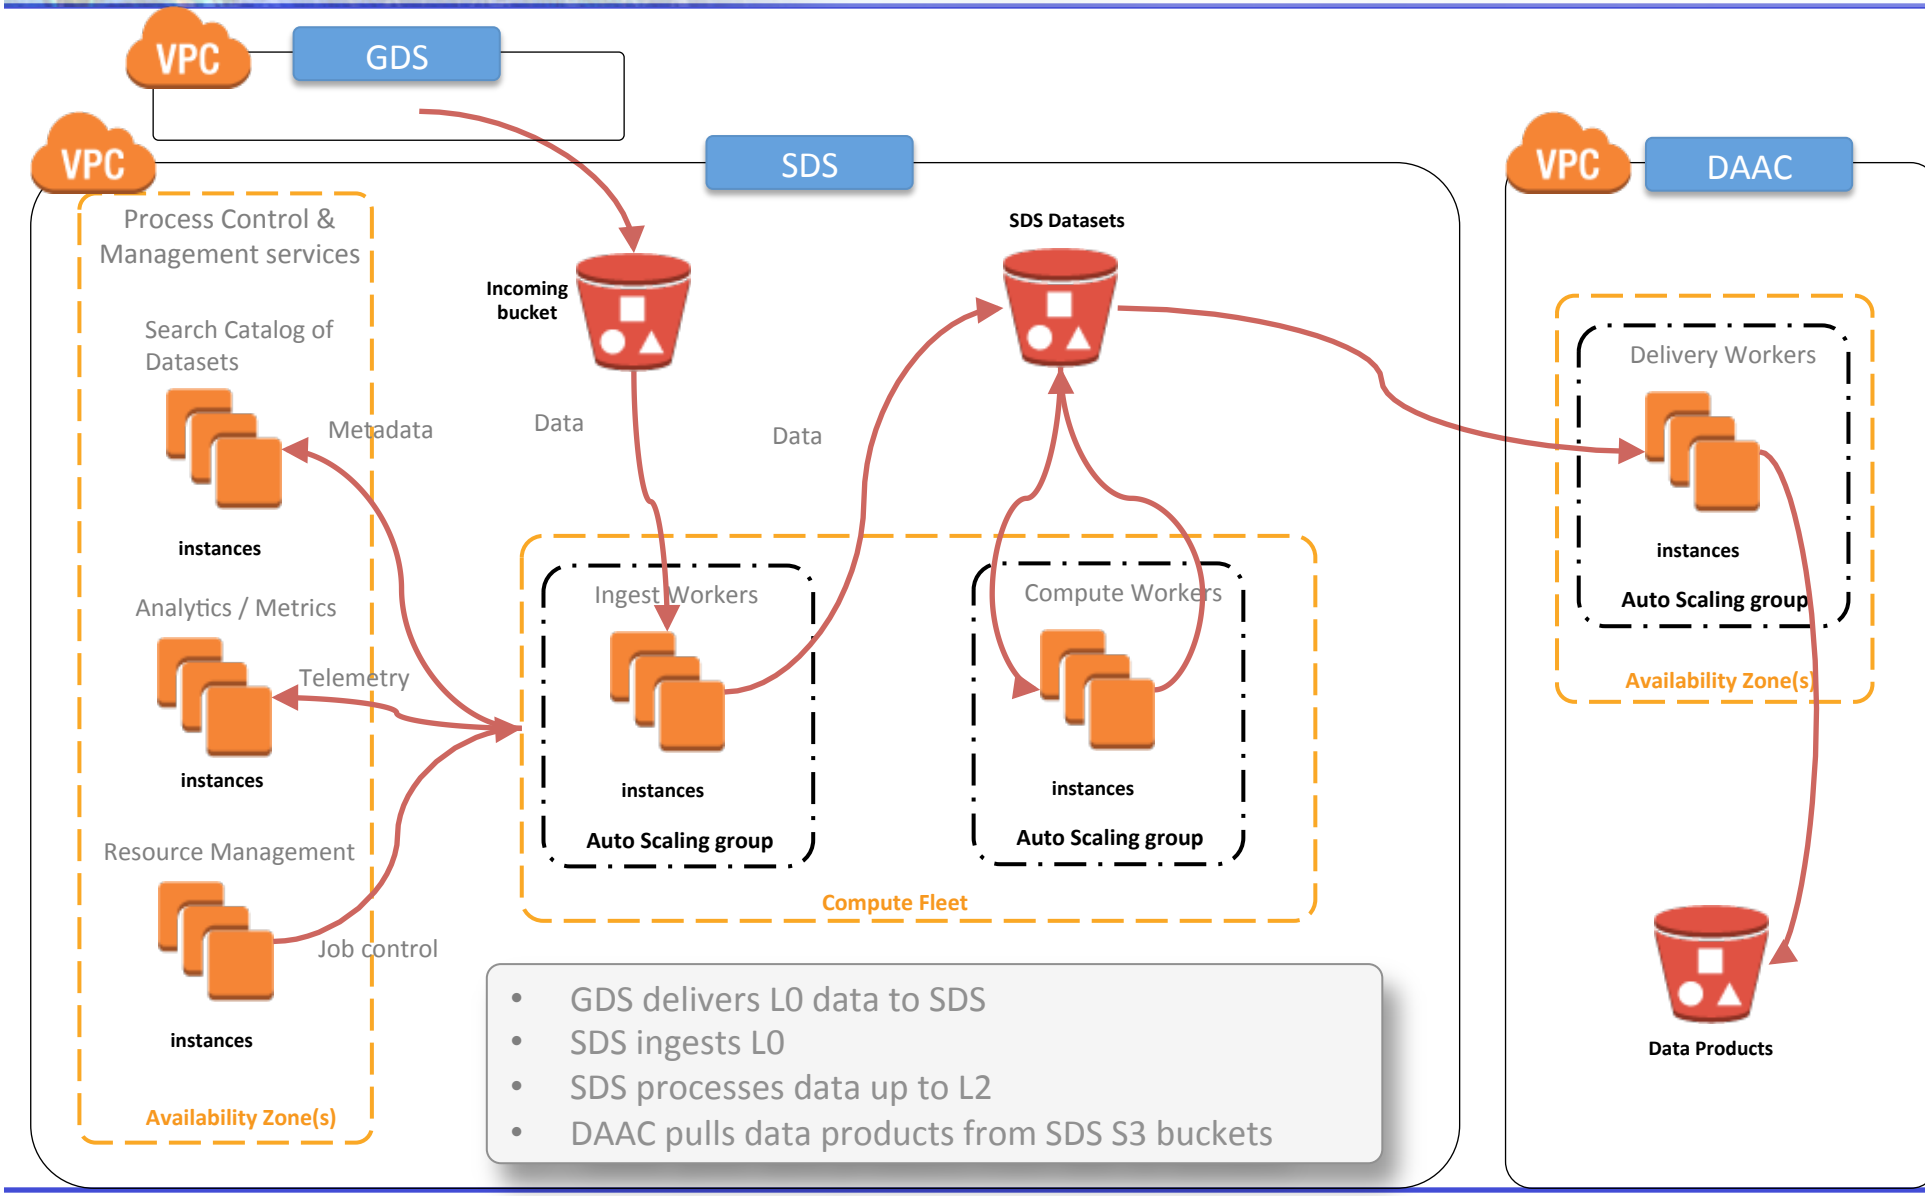

- GDS delivers L0 data to SDS
- SDS ingests L0
- SDS processes data up to L2
- DAAC pulls data products from SDS S3 buckets

"Containerizing" Product Generation Executives (PGEs)

- **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."*

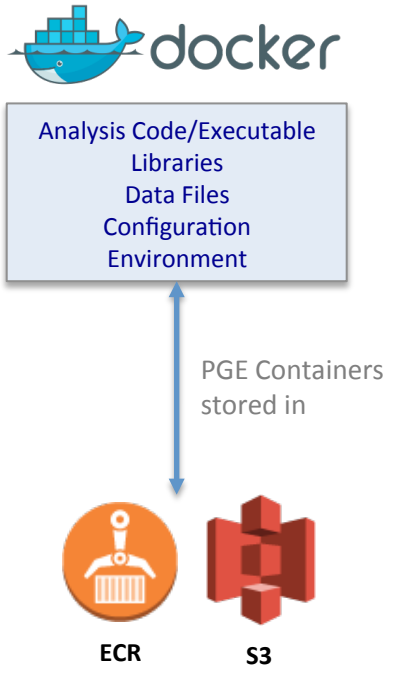

Relevant Cost Components

Component	Service	Unit Costs
Compute	Elastic Compute Cloud (EC2)	• On-demand (~$2/hr for 32-vCPU) • Reserved Instances (~50% of on-demand) • Spot instance (~50% of RI)
Compute	EC2 Container Service (ECS)	• ECS runs on EC2
Storage	S3	• Tiered storage. $0.0275/GB for 5PB+
Storage	S3 Standard - Infrequent Access	• Tiered storage. $0.0125/GB for 5PB+ • Data retrieval: $0.01/GB • PUT, COPY, or POST Request: $0.01 per 1,000 requests • GET and all other Requests: $0.01 per 10,000 requests
AWS Data Movement Out	Egress	• *Discounted* $0.04/GB
On-premise Network	10Gbps to internet (e.g. from SuperNAP)	• ~$2K/month
On-Premise Facilities	rack, cooling, electrical	• ~$1500/month/rack

Data Tiering

- Optimizing cost to storage use

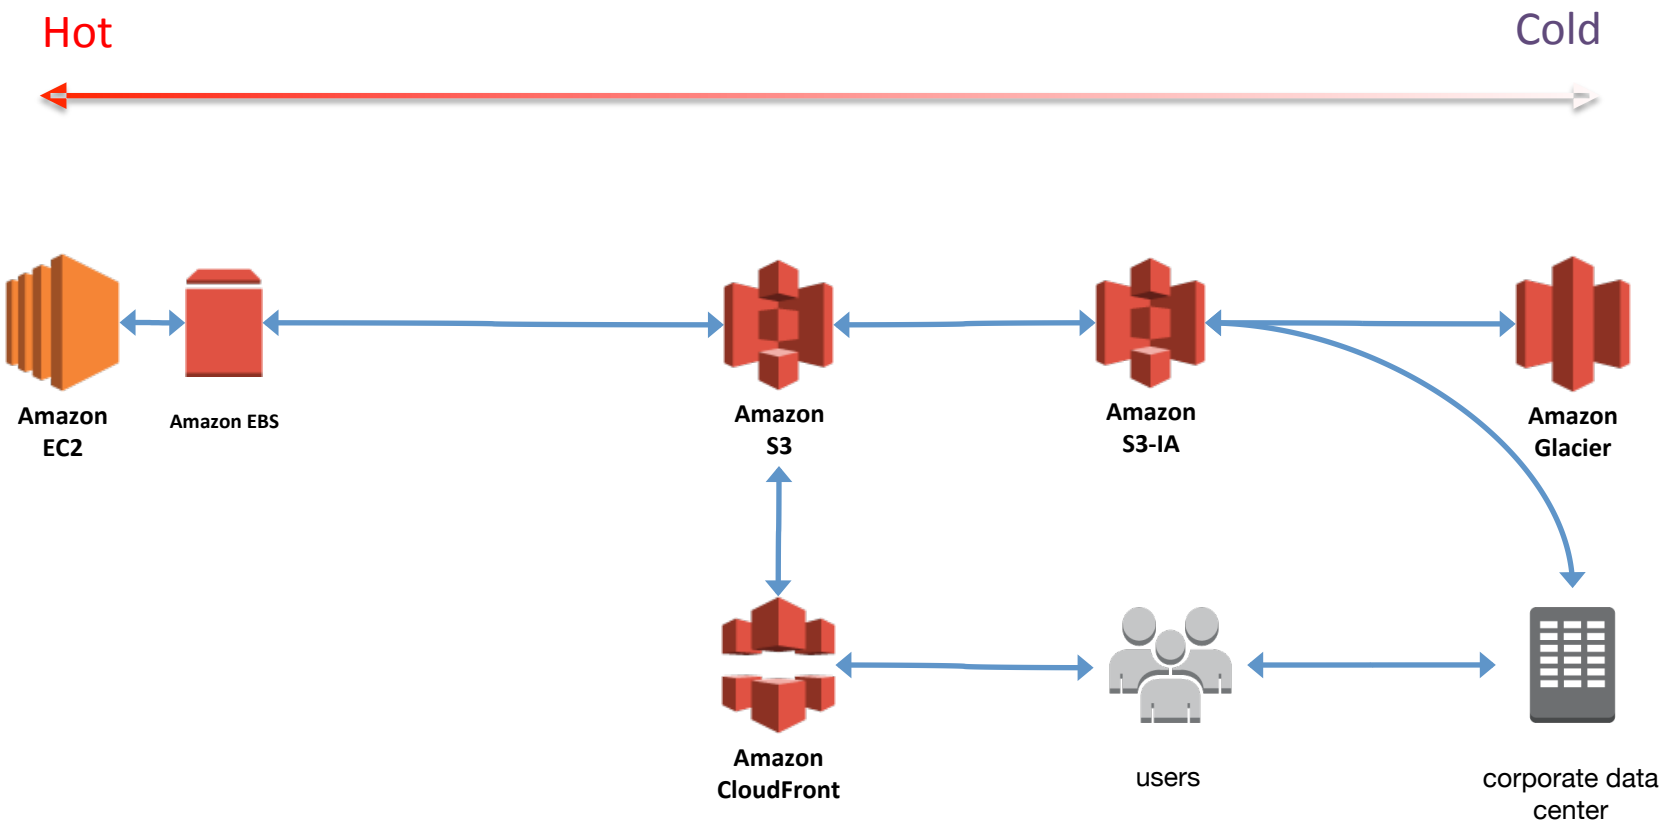

AWS Storage: Monthly Tiered Example

- Storage types
  - EBS
  - S3
  - S3 Infrequent Access
  - Glacier
- Storage cost tiers
- *At public rack-rates… Not cost-effective to store everything in S3*

Monthly S3 Storage Costs (rack rate)

At day 1

Monthly cost			Price/GB	
Tier 1: First 1 TB	1.00 TB	x	0.0300 =	30.00 $
Tier 2: Next 49 TB	49.00 TB	x	0.0295 =	1,445.50 $
Tier 3: Next 450 TB	36.00 TB	x	0.0290 =	1,044.00 **$**
Tier 4: Next 500 TB	0.00 TB	x	0.0285 =	- **$**
Tier 5: Next 4000 TB	0.00 TB	x	0.0280 =	- $
Tier 6: Above 5000 TB	0.00 TB	x	0.0275 =	- $
Total Monthly	86.00 TB			2,519.50 $

At month 1

Monthly cost			Price/GB	
Tier 1: First 1 TB	1.00 TB	x	0.0300 =	30.00 $
Tier 2: Next 49 TB	49.00 TB	x	0.0295 =	1,445.50 $
Tier 3: Next 450 TB	450.00 TB	x	0.0290 =	13,050.00 **$**
Tier 4: Next 500 TB	500.00 TB	x	0.0285 =	14,250.00 **$**
Tier 5: Next 4000 TB	1580.00 TB	x	0.0280 =	44,240.00 $
Tier 6: Above 5000 TB	0.00 TB	x	0.0275 =	- $
Total Monthly	2580.00 TB			73,015.50 $

At year 1

Monthly cost			Price/GB	
Tier 1: First 1 TB	1.00 TB	x	0.0300 =	30.00 $
Tier 2: Next 49 TB	49.00 TB	x	0.0295 =	1,445.50 $
Tier 3: Next 450 TB	450.00 TB	x	0.0290 =	13,050.00 **$**
Tier 4: Next 500 TB	500.00 TB	x	0.0285 =	14,250.00 **$**
Tier 5: Next 4000 TB	4000.00 TB	x	0.0280 =	112,000.00 $
Tier 6: Above 5000 TB	26390.00 TB	x	0.0275 =	725,725.00 $
Total Monthly	31390.00 TB			866,500.50 $

Forward and Bulk Processing Data Volume/ Throughput per Day

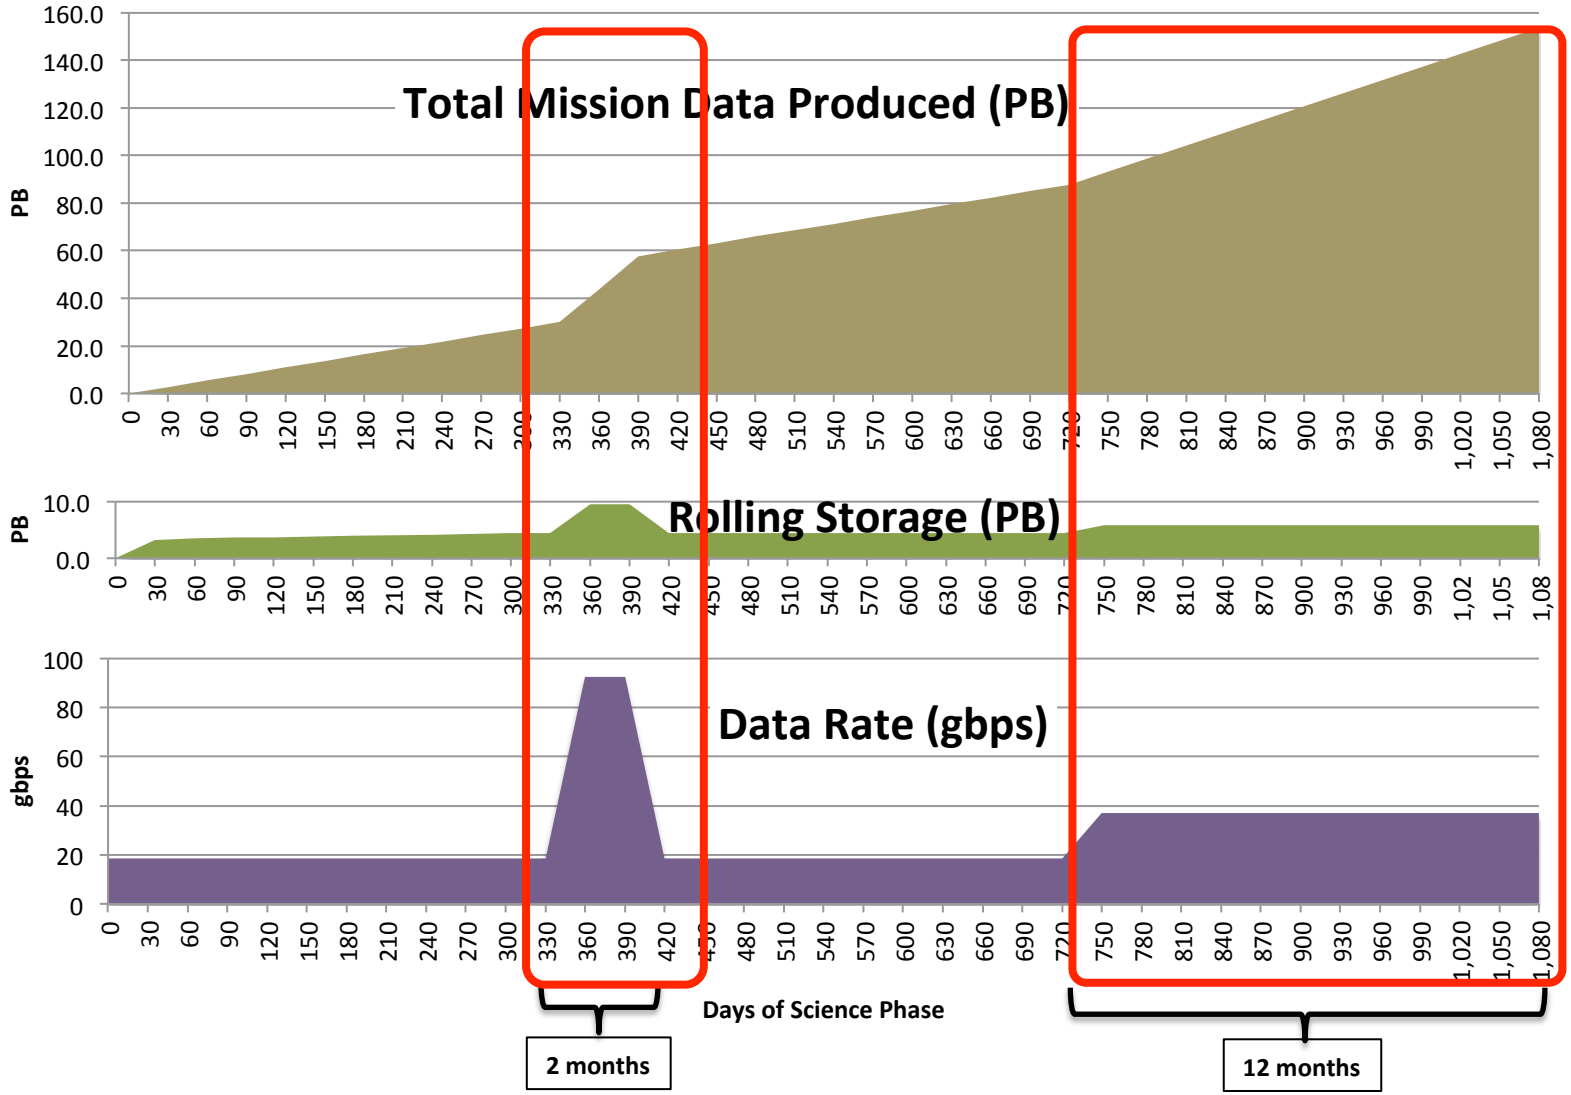

Example Bulk Reprocessing Campaigns

Cloud-based Auto-Scaling of Compute

- *The size of the data system compute nodes can automatically grow/shrink based on demand*
- *Easily scale to over 100,000 vCPUs*
- *"Pay as you go"*

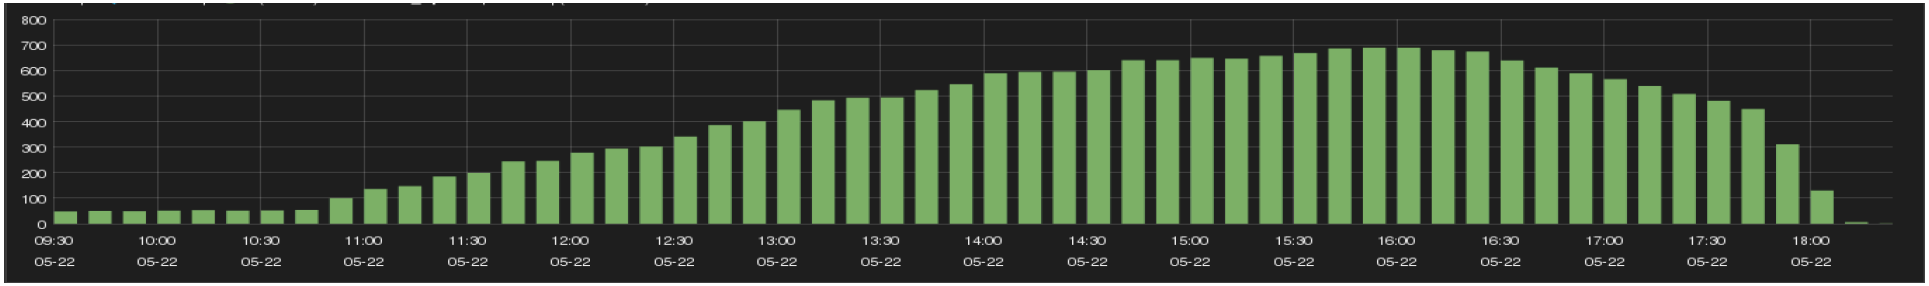

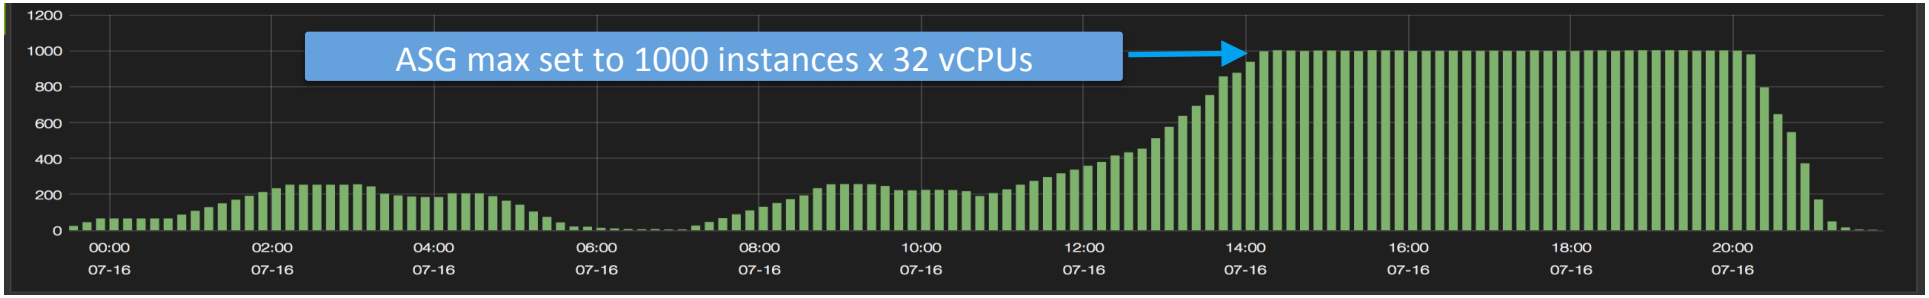

AWS Spot Market

					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 on compute (**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
- *Running in spot market forces data system to be more **resilient** to compute failures*

High-Resiliency at Large-Scales

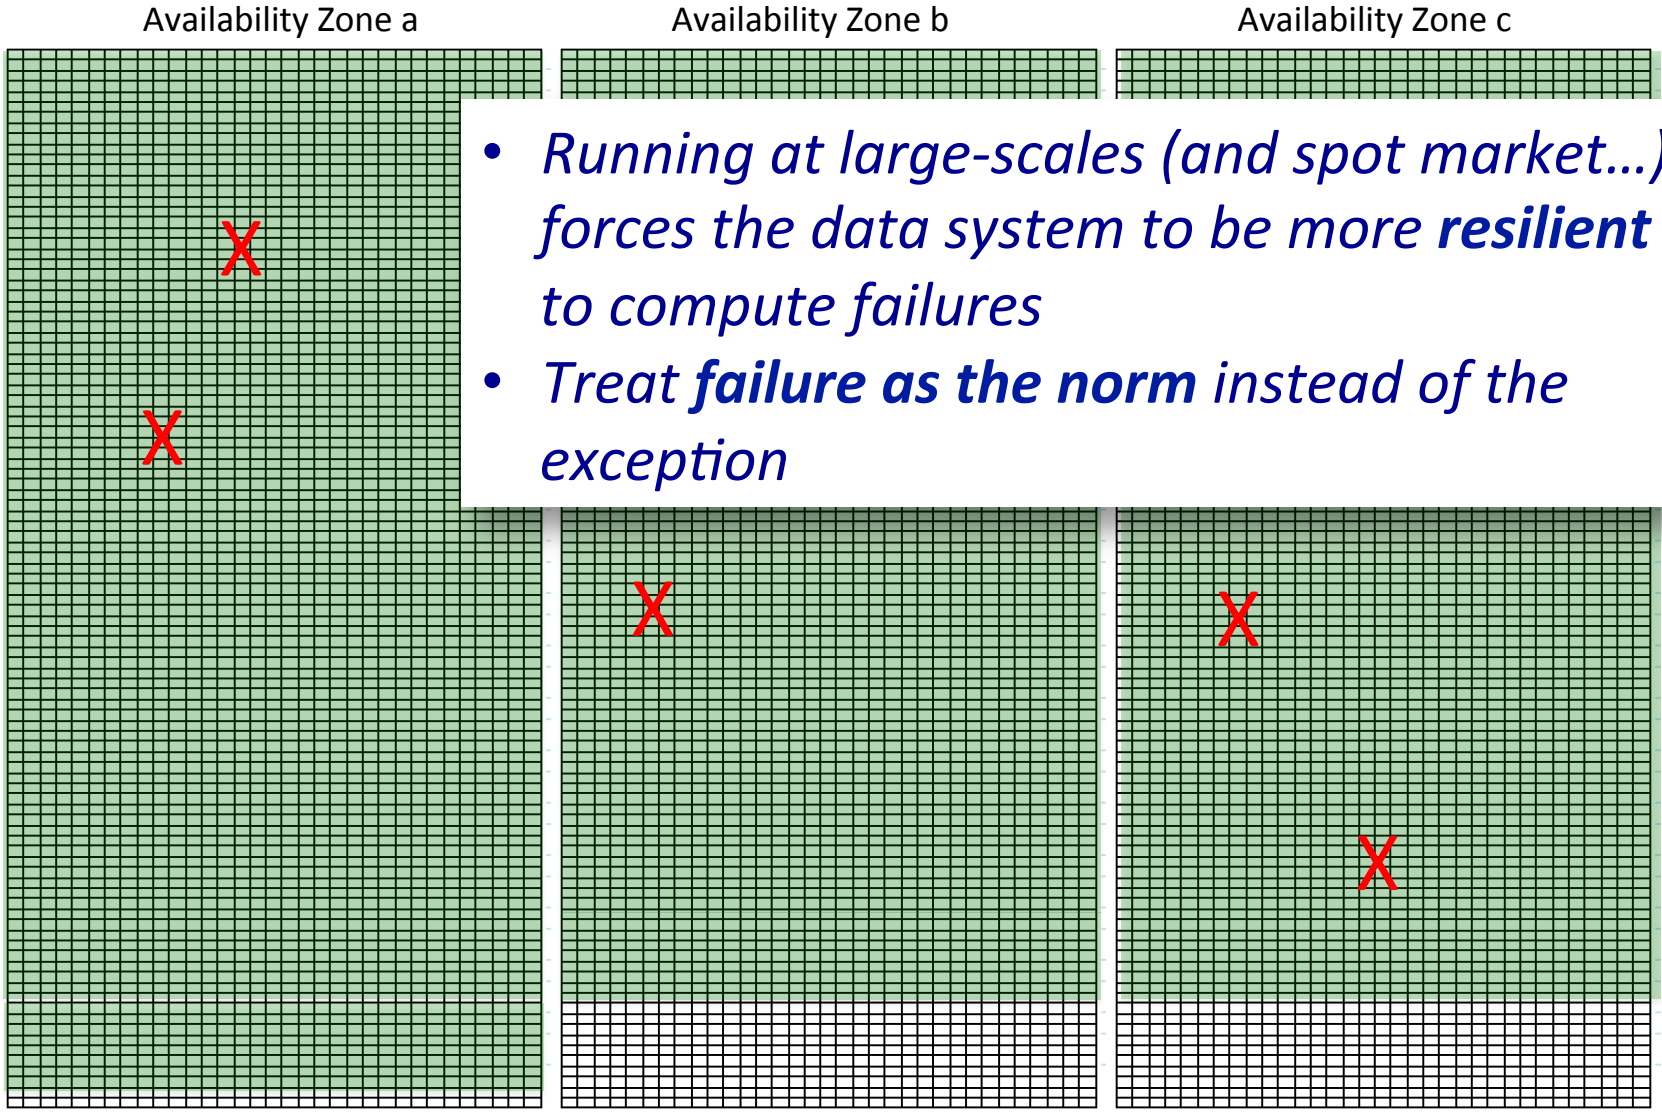

- *Running at large-scales (and spot market…) forces the data system to be more **resilient** to compute failures*
- *Treat **failure as the norm** instead of the exception*

“Market Maker”

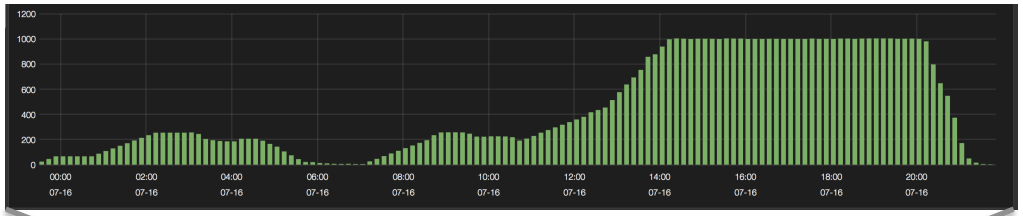

This OCO-2 production run of 1000 x 32vCPUs affected the market prices

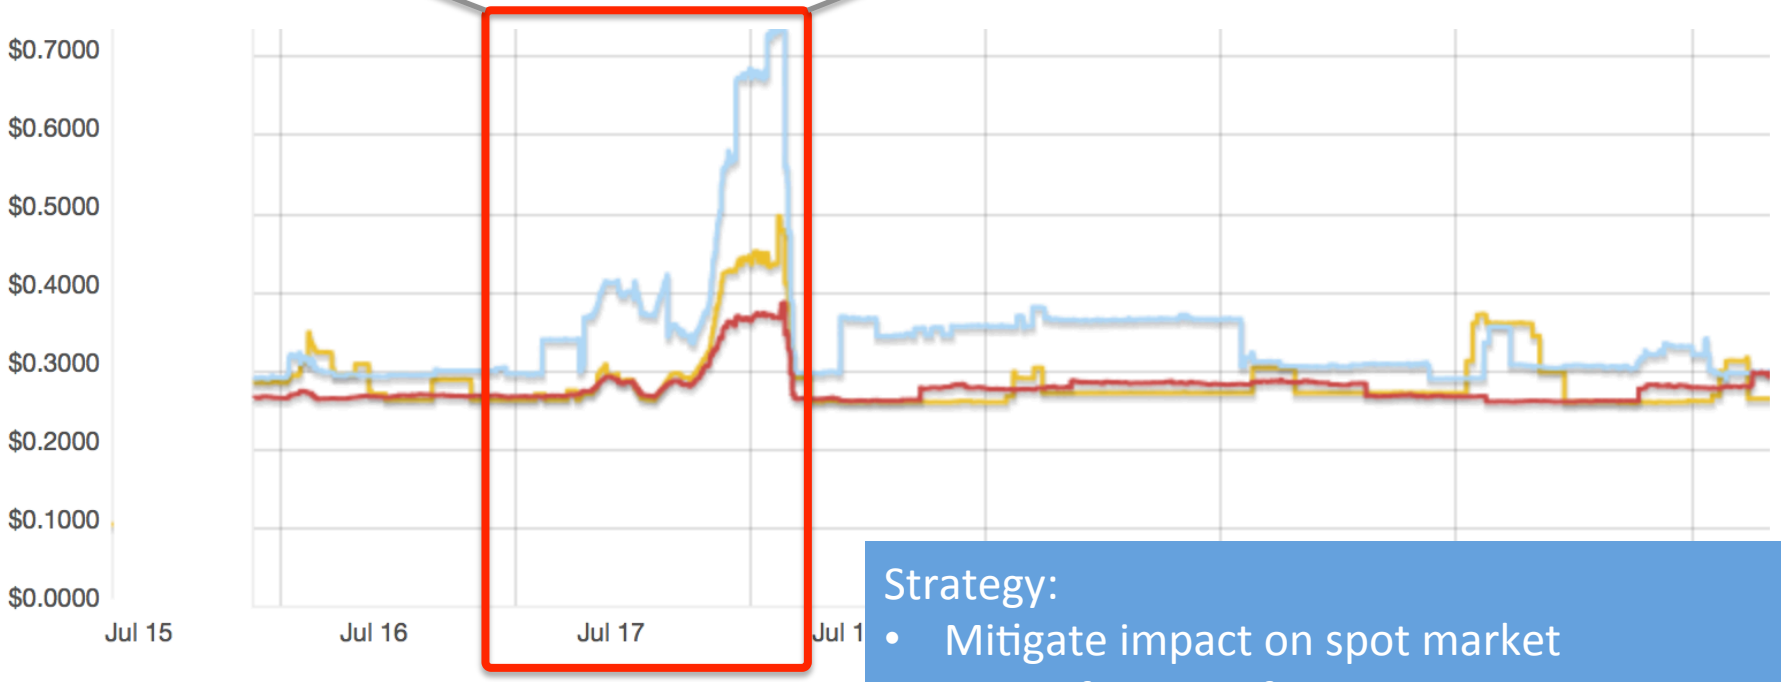

Strategy:

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

Validating Cloud Adaptation Version

- Algorithm team to validate the cloud-native version within acceptable tolerance

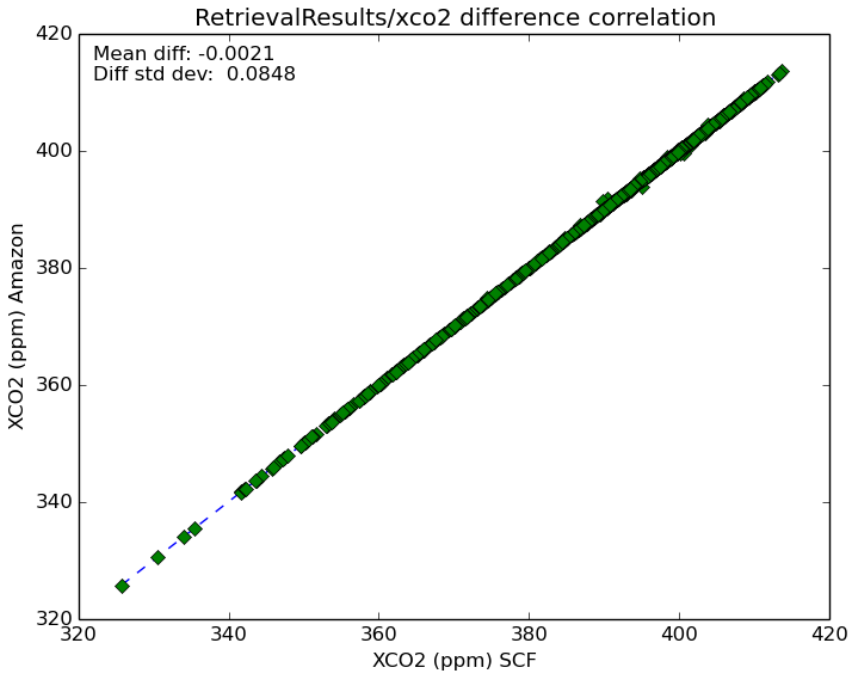

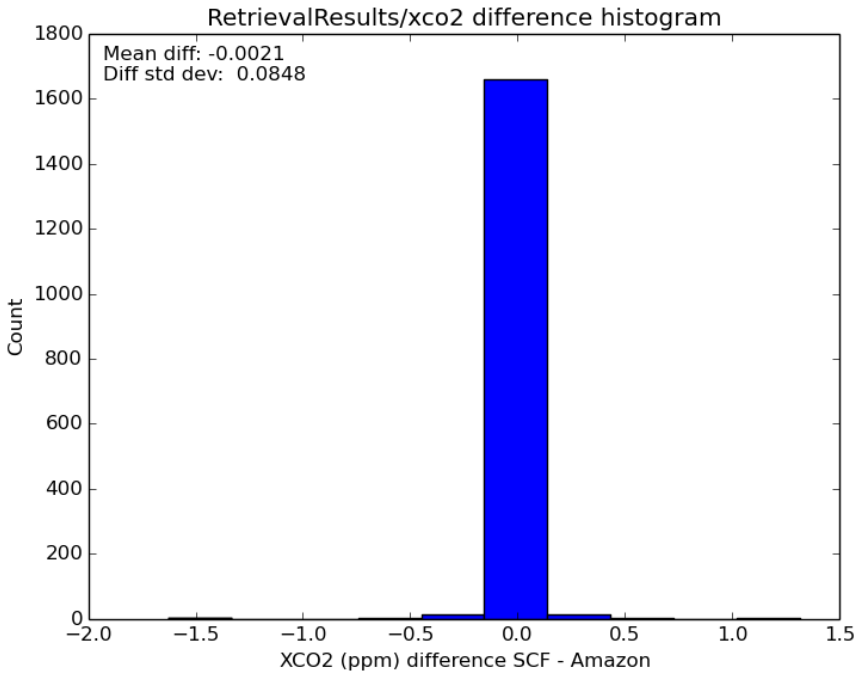

Classic Deployment of SDS and DAAC

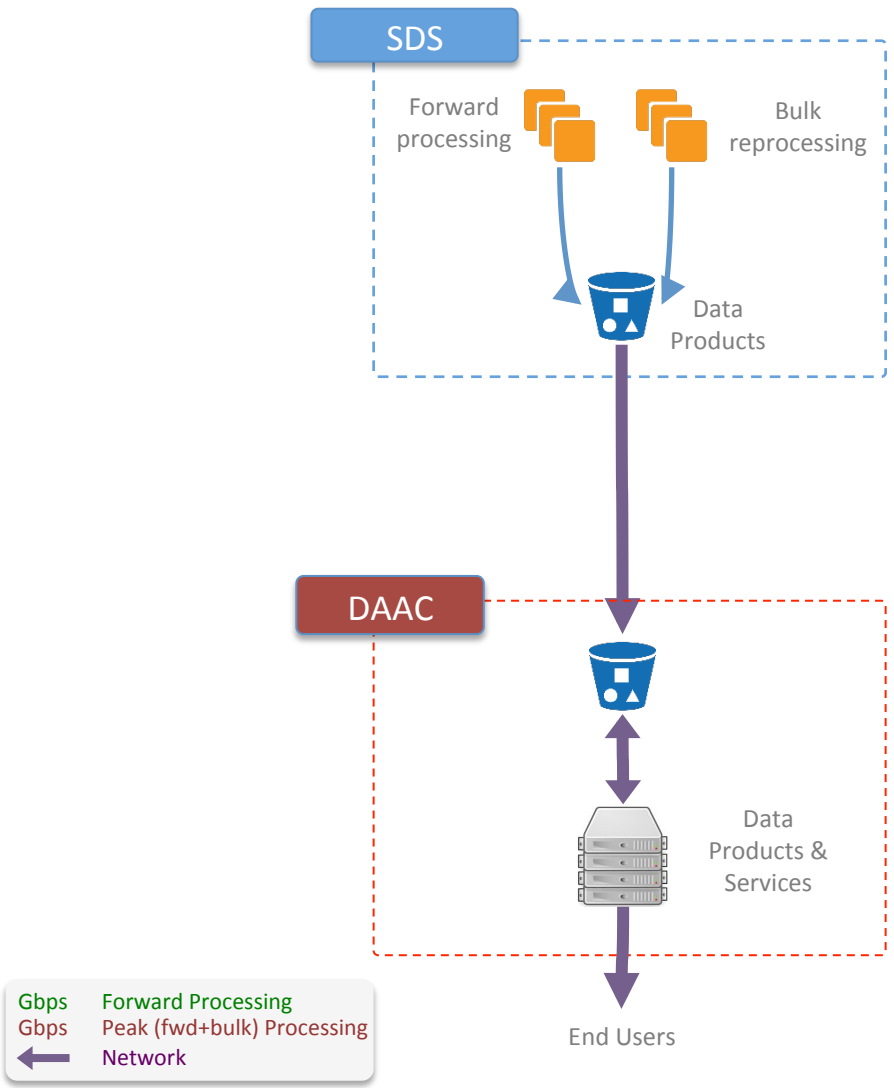

- Science data product generation at SDS
- Science data products moved to DAAC facilities
  - (copying large data volumes)
- End users access from DAAC
- Bottlenecks and cost impact of high network data stream

Collocation of SDS and DAAC in Cloud

- **Collocated data** between SDS and DAAC
- No egress nor external network limitations between SDS and DAAC
- DAAC still incurs end-user egress costs.

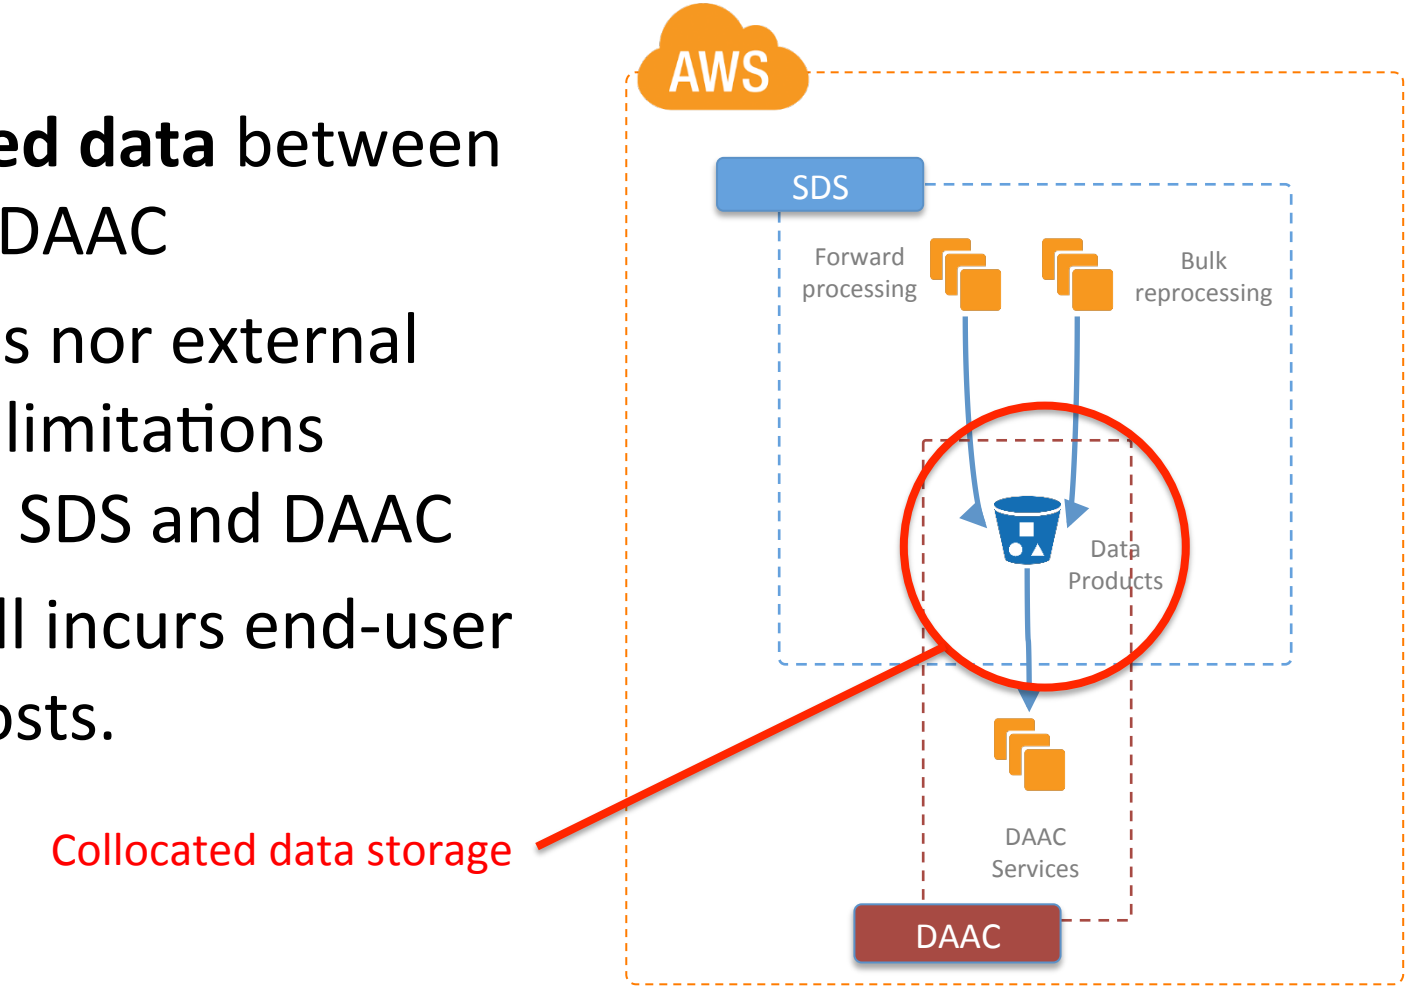

Move the Analysis, not the Big Data

- Factoring in data locality
- Move compute closer to data

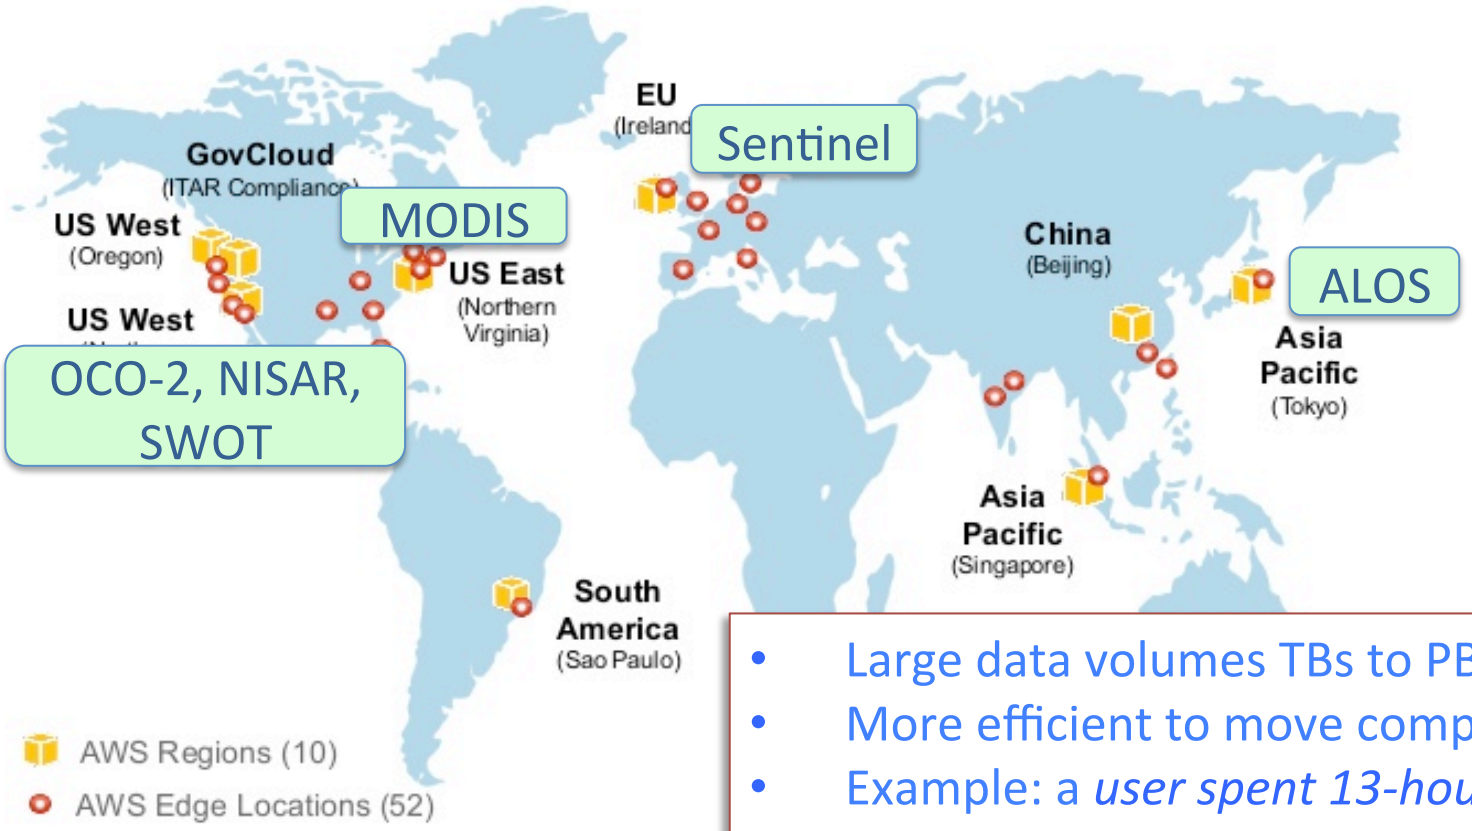

- Large data volumes TBs to PBs
- More efficient to move compute to the data
- Example: a *user spent 13-hours download data + another 3-hours decompressing the data, then started the analysis*

“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

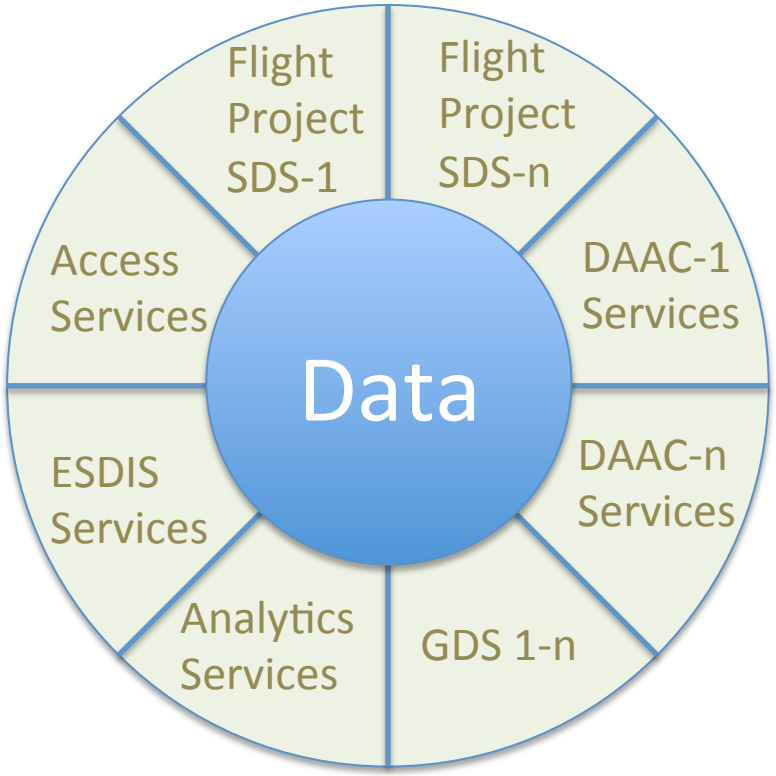

Enabling multi-disciplinary data approach for analysis