Leveraging On-premise and Public Cloud Computing to Enable Advanced Rapid Imaging & Analysis for Monitoring Hazards

ESIP 2014 Winter Meeting, Washington DC

Cloud Computing Governance

Thursday, January 9, 2014

Hook Hua[1], Susan Owen[1], Sang-Ho Yun[1], Paul Lundgren[1], Angelyn Moore[1], Piyush Agram[1], Gian Franco Sacco[1], Eric Fielding[1], Paul Rosen[1], Frank Webb[1], Mark Simons[2], Brian Wilson[1], Timothy Stough[1], Peter F. Cervelli[4], Michael Poland[3], Jennifer Cruz[1]

[1] Jet Propulsion Laboratory
[2] California Institute of Technology
[3] USGS Hawaiian Volcano Observatory
[4] USGS Menlo Park Science Center

Government sponsorship acknowledged.

ESTO
Earth Science Technology Office

The Challenge of Leveraging Remote Sensing for Disaster Response

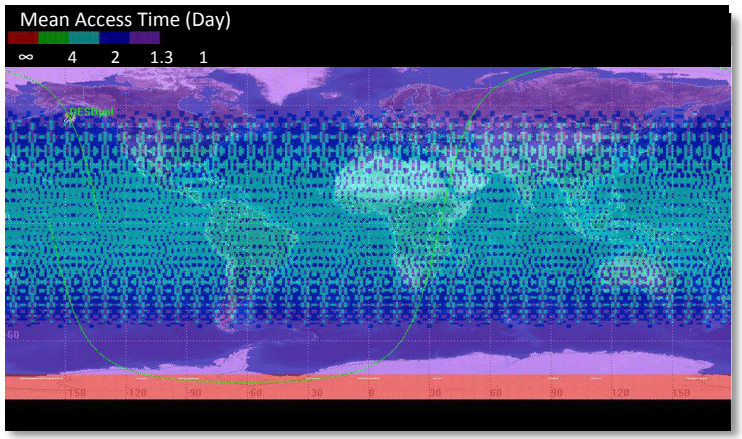

Automated data system are required to analyze large quantities of data from proposed NASA NI-SAR (formerly DESDynI), other satellite missions, and rapidly expanding GPS networks.

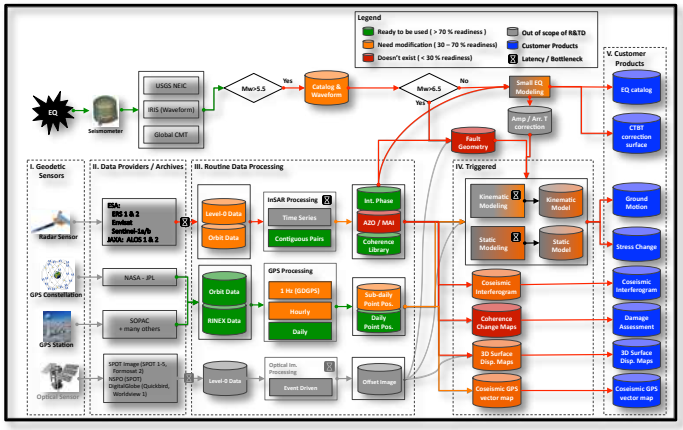

Going from Artisan to Automation: Use system engineering approach to translate specialized data analysis into operational capability.

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

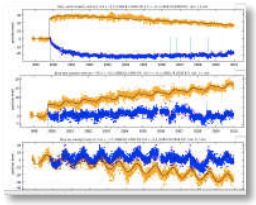

Temporal Records of Ground Deformation

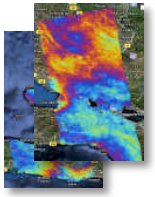

Spatial Maps of Ground Deformation

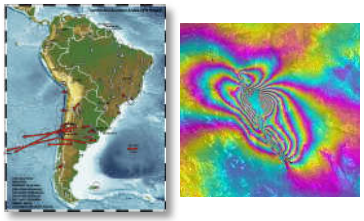

Coseismic Ground Deformation

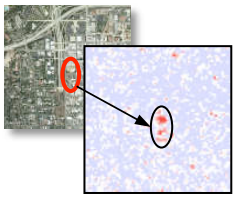

Coseismic Damage

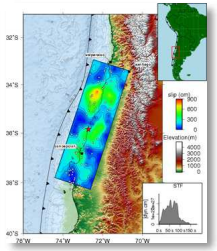

Earthquake Models

Advanced Rapid Imaging & Analysis for Monitoring Hazards (ARIA-MH)

- **High-volume, low-latency**, and automatic generation of NASA Solid Earth science data products (InSAR and GPS) to support hazards monitoring.
- Enabling both **science and decision-support** communities to monitor ground motion in areas of interest with InSAR and GPS data.
- Leverage and geographically optimize **hybrid Cloud**-based processing and data management of geodetic data products
- **Monitoring**
  - Event streams from USGS NEIC
  - Data product streams
- **Conditional** triggering of geodetic data processing

Elastic Processing, Data Management, and Monitoring

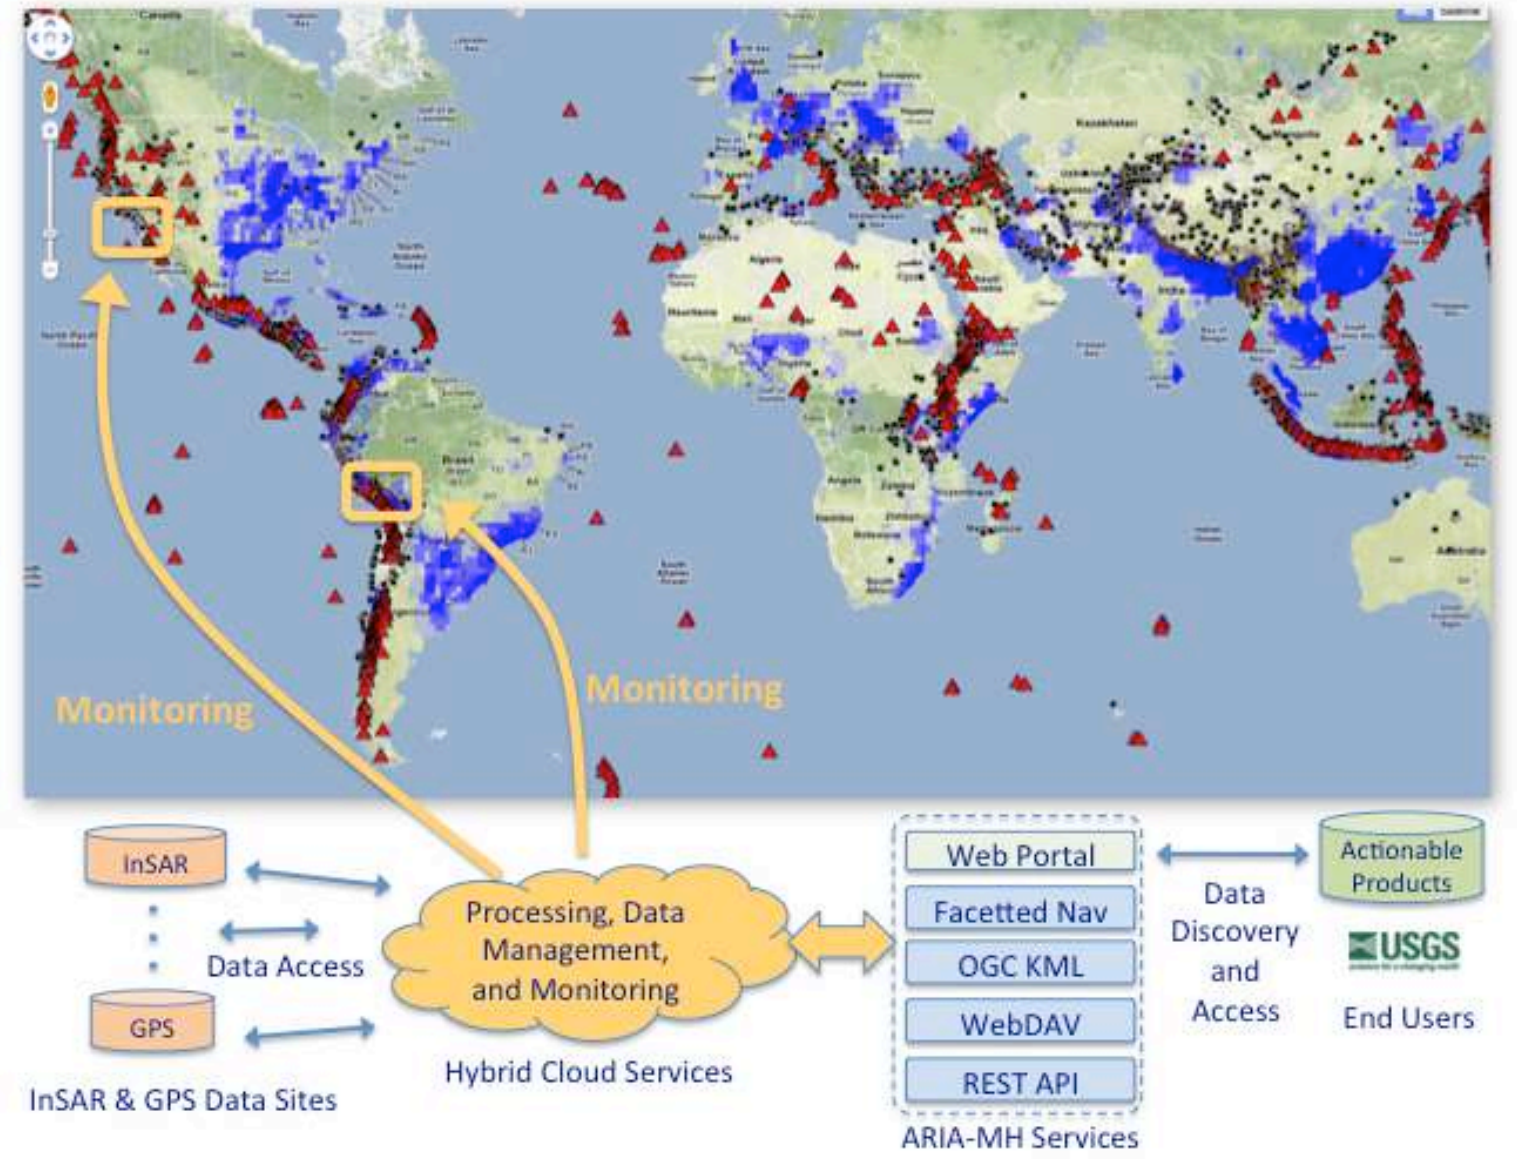

Near Real-Time Data Stream Established

- JPL/Caltech/ASI collaboration effort opens flood of COSMO-SkyMed (CSK) data for select regions
  - Provides access to CSK X-band Level-0 SAR radar data from the Italian Space Agency (ASI)
  - *COSMO-SkyMed (CSK) data provided as part of a **technical collaboration** between JPL-Caltech and the Center for Earth Observations (CIDOT), Italian Space Agency (ASI) (CSK data © 2013 ASI).*
- Near real-time CSK data stream setup into ARIA-MH science data system
  - CSK constellation of 4 satellites has acquisition capacity of 450 frames/day for each satellite
  - Example: San Andreas Fault region of California
    - 580 GB of raw data every 16 days, about 500 frames.
  - ~2TB per month (for CSK on San Andreas Fault region)
  - raw + derived products doubles the data volume

Raw and Interferogram Scene Sizes

Single scene or frame

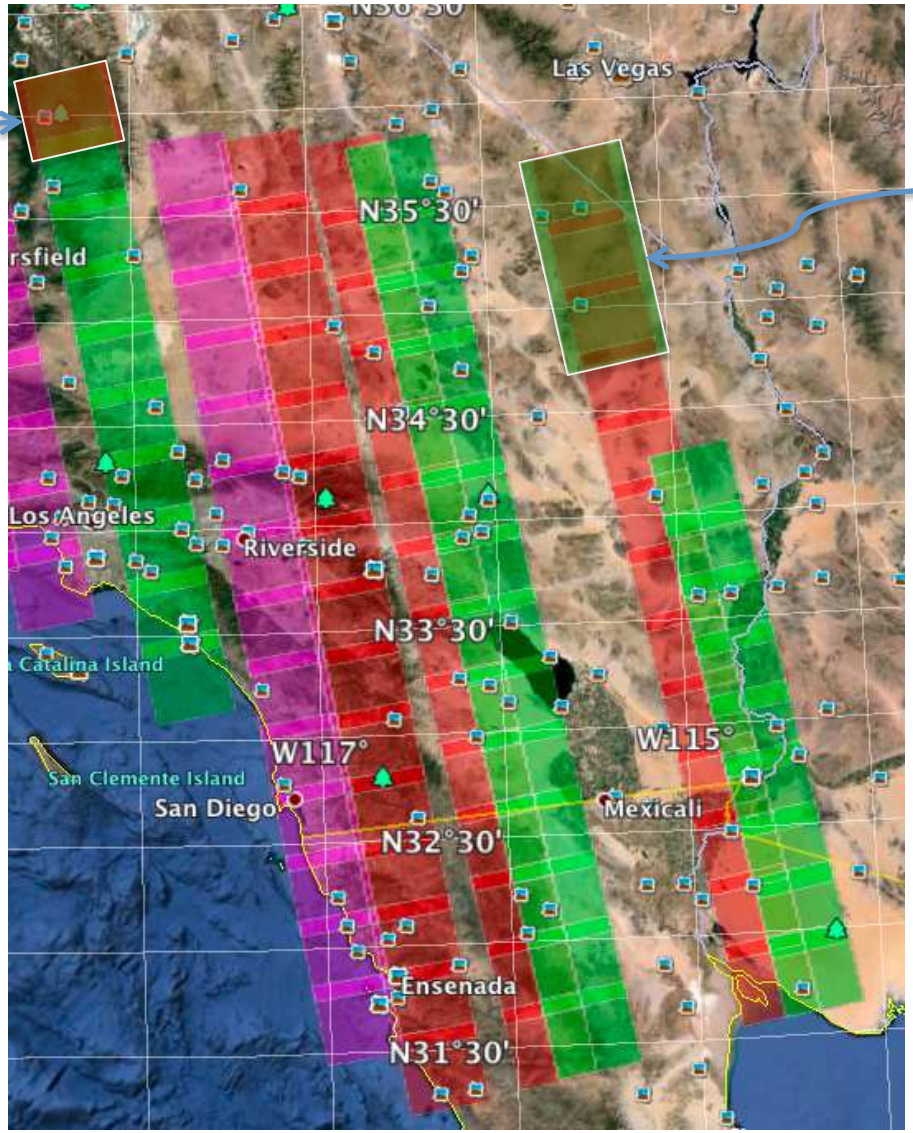

Stitched scene or frame

Big Data Needs

- Raw data
  - 1 frame = 40 x 40 Kms swath, 3 m. res., 1.2 Gb
  - Total = 580 Gb every 16 days of raw data
- Derived Data Volume
  - Minimum 2 frames for each interferogram
  - Up to 8X of baseline pairs from each frame
  - 3X data volume increase in derived data products from each frame
  - **California area only: PB-scale for 36-months of observations**
- Data movement and storage concerns for handling global-scale coverage

CSK Scene Footprints: Rolling 1-day view

Satellite observation of granules from Italian X-band Level-0 SAR radar data

CSKS1
CSKS2
CSKS3
CSKS4

CSK Scene Footprints: Rolling 16-day cycle view

Satellite observation of granules from Italian X-band Level-0 SAR radar data

- CSKS1
- CSKS2
- CSKS3
- CSKS4

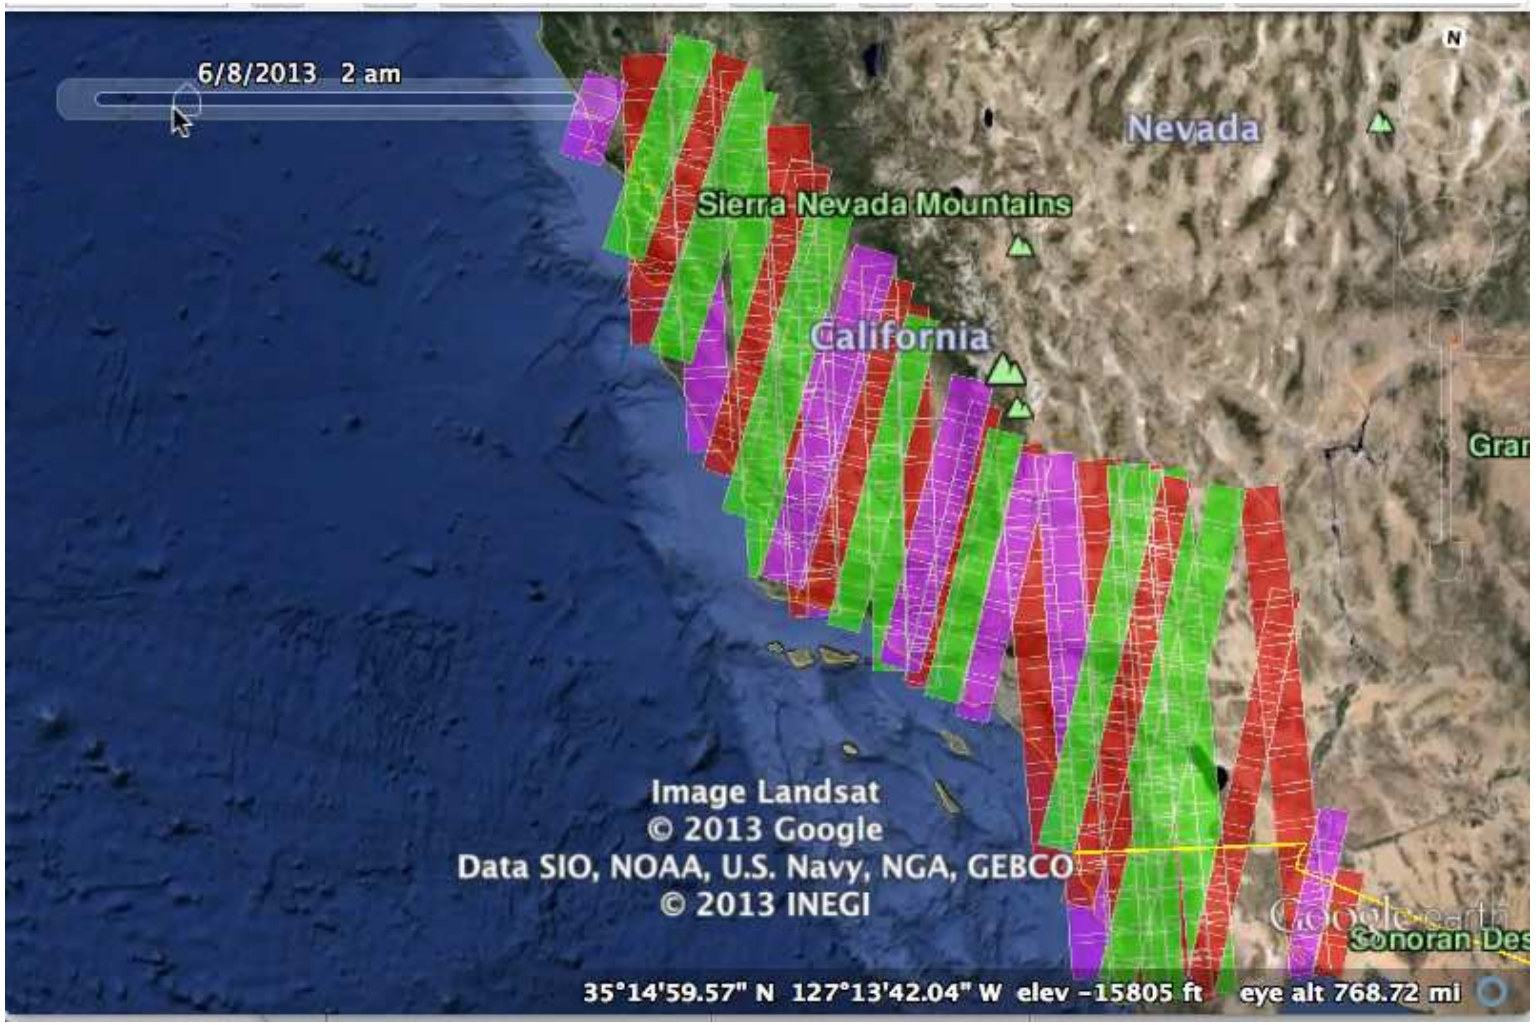

Projects Supported by ARIA-MH

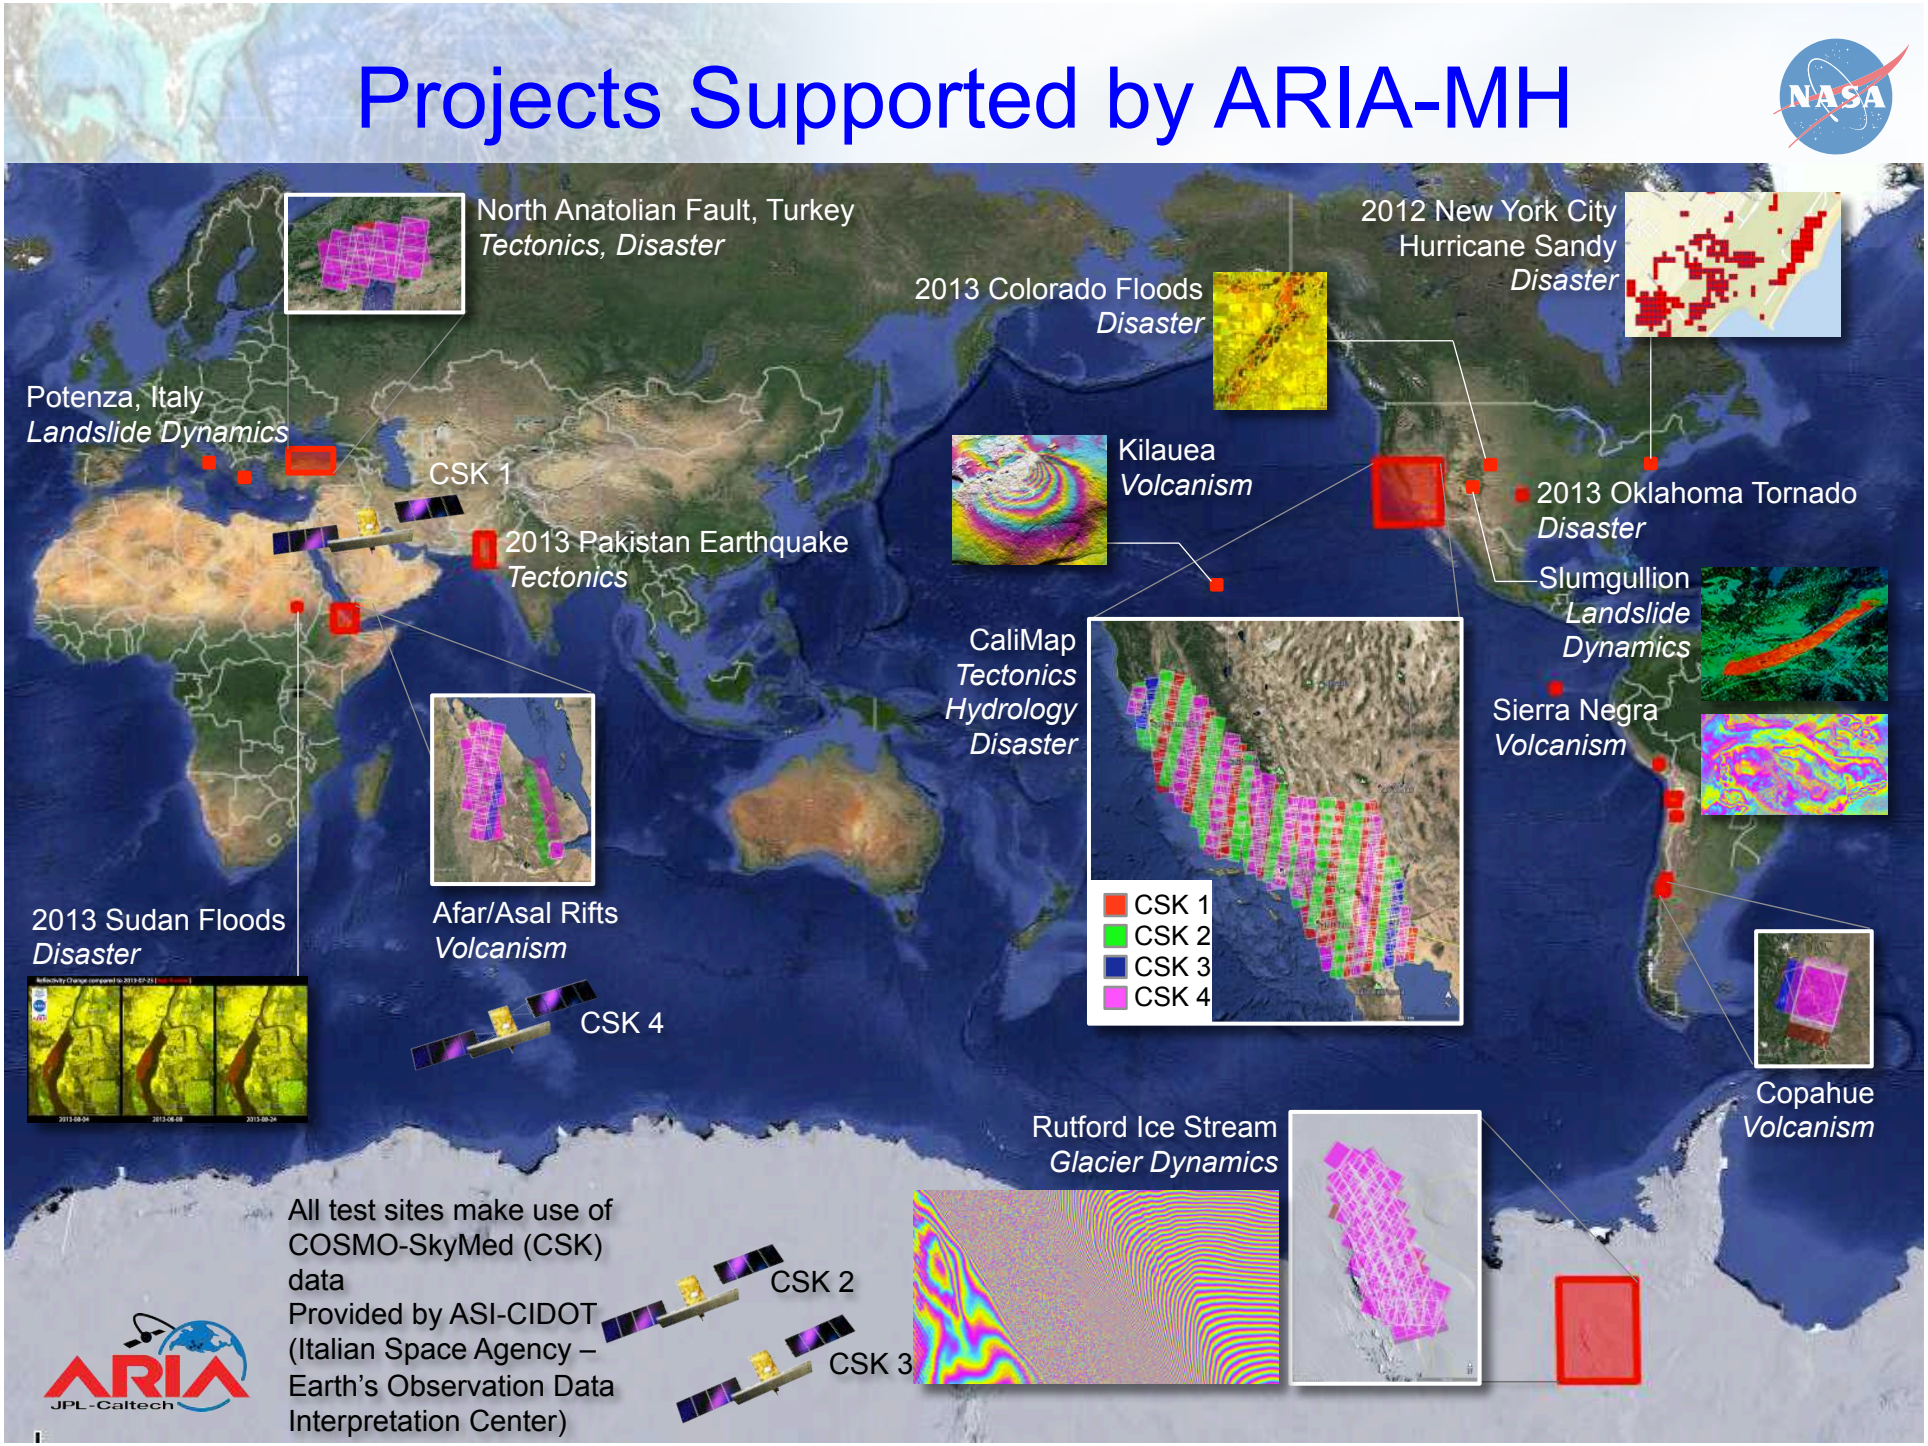

ARIA's Haiyan Response Timeline

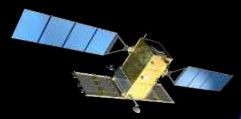

COSMO-SkyMed
acquires data 8/19

Day -81
09:00

Day 0
00:00

Haiyan hits Philippines
2013/11/08 00:00 UTC

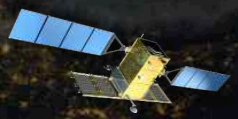

COSMO-SkyMed
acquires data 11/11

Day 3
09:00

Day 2
02:00

GeoEye-1
acquires data 11/10

ftp://

ARIA accesses
the data

Day 3
14:00

Day 3

Copernicus EMS
produces damage
grading maps of
selected areas
2013/11/11

ARIA produces
DPMs 2013/11/11
(11 hours after
data acquisition)

Day 3
20:00

time

Source: Sang-Ho Yun (JPL)

Processing Architecture (1 of 2)

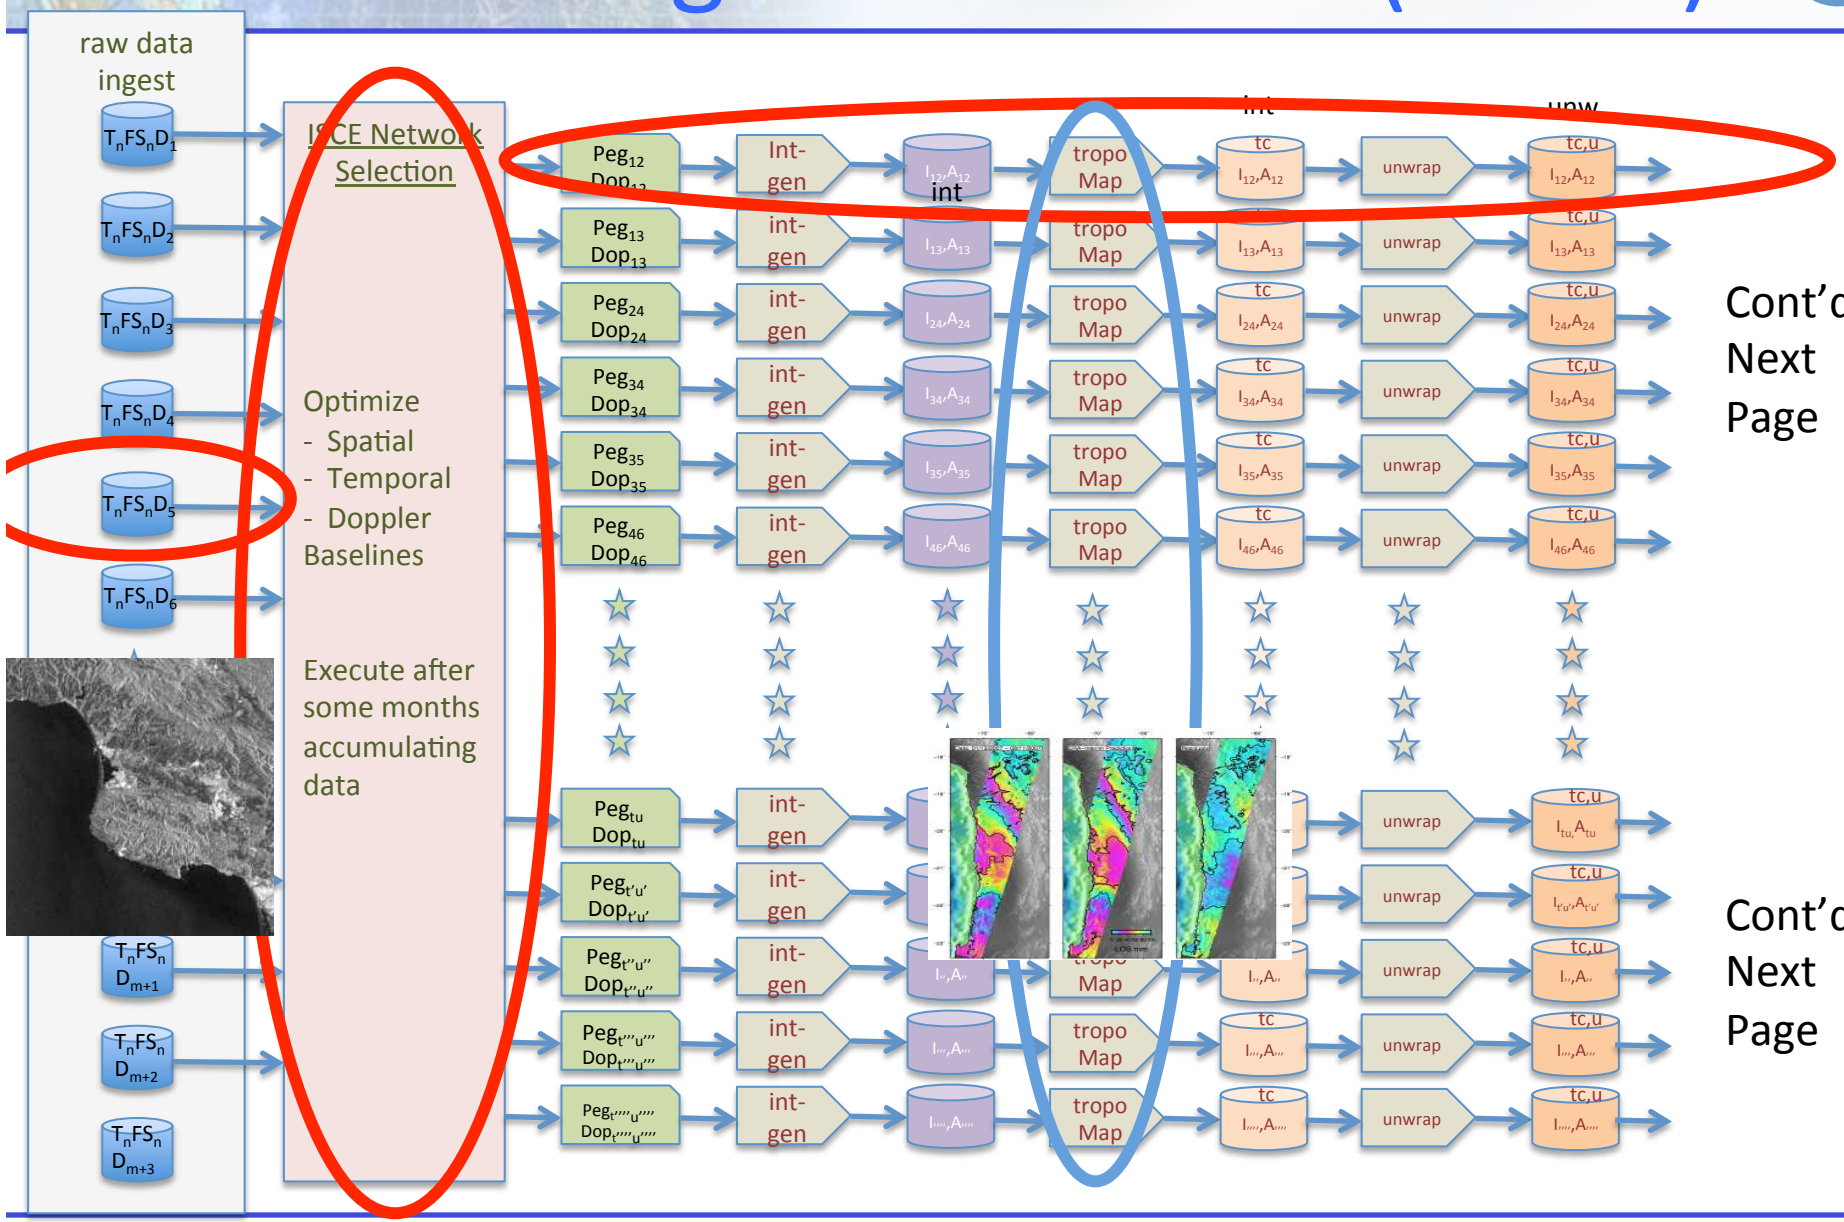

Processing Architecture (2 of 2)

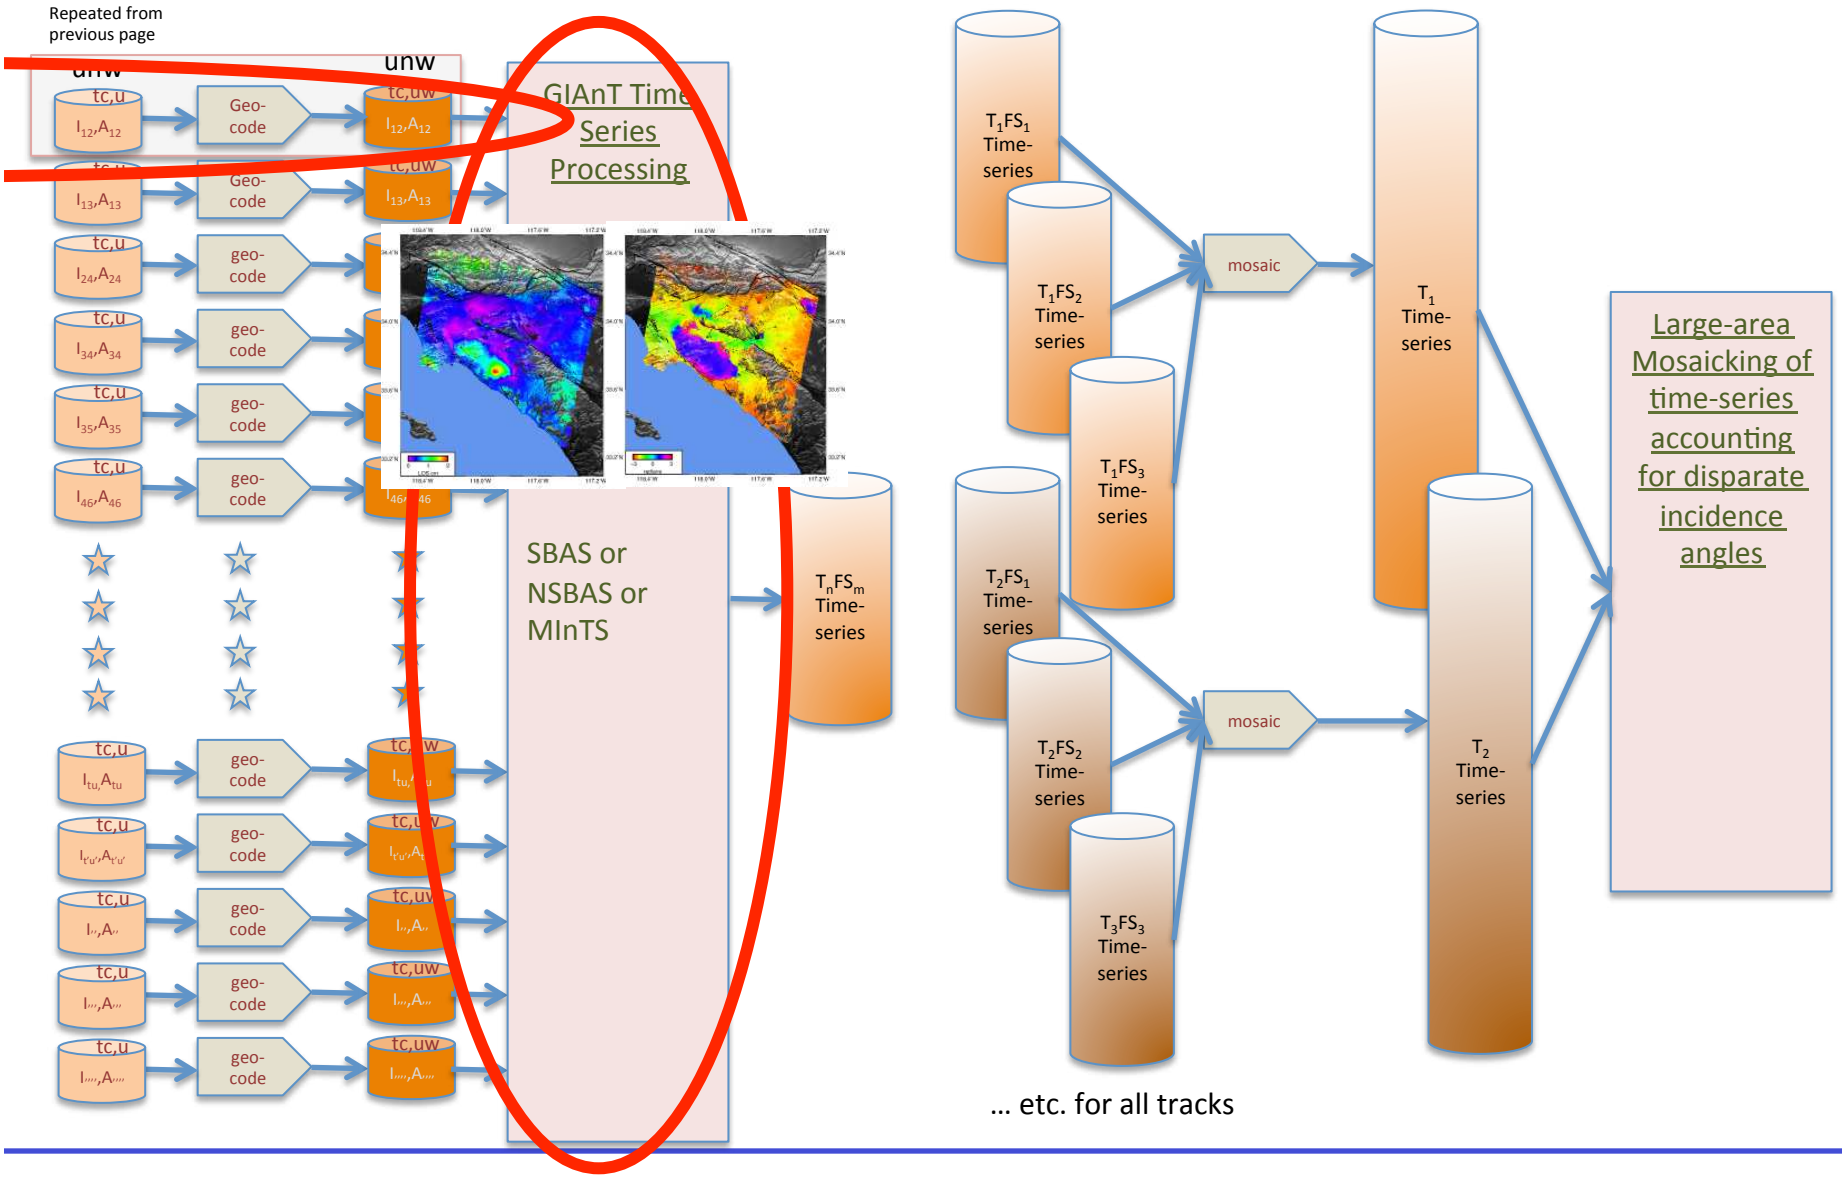

DPM in ARIA-MH Project 2012.11.24

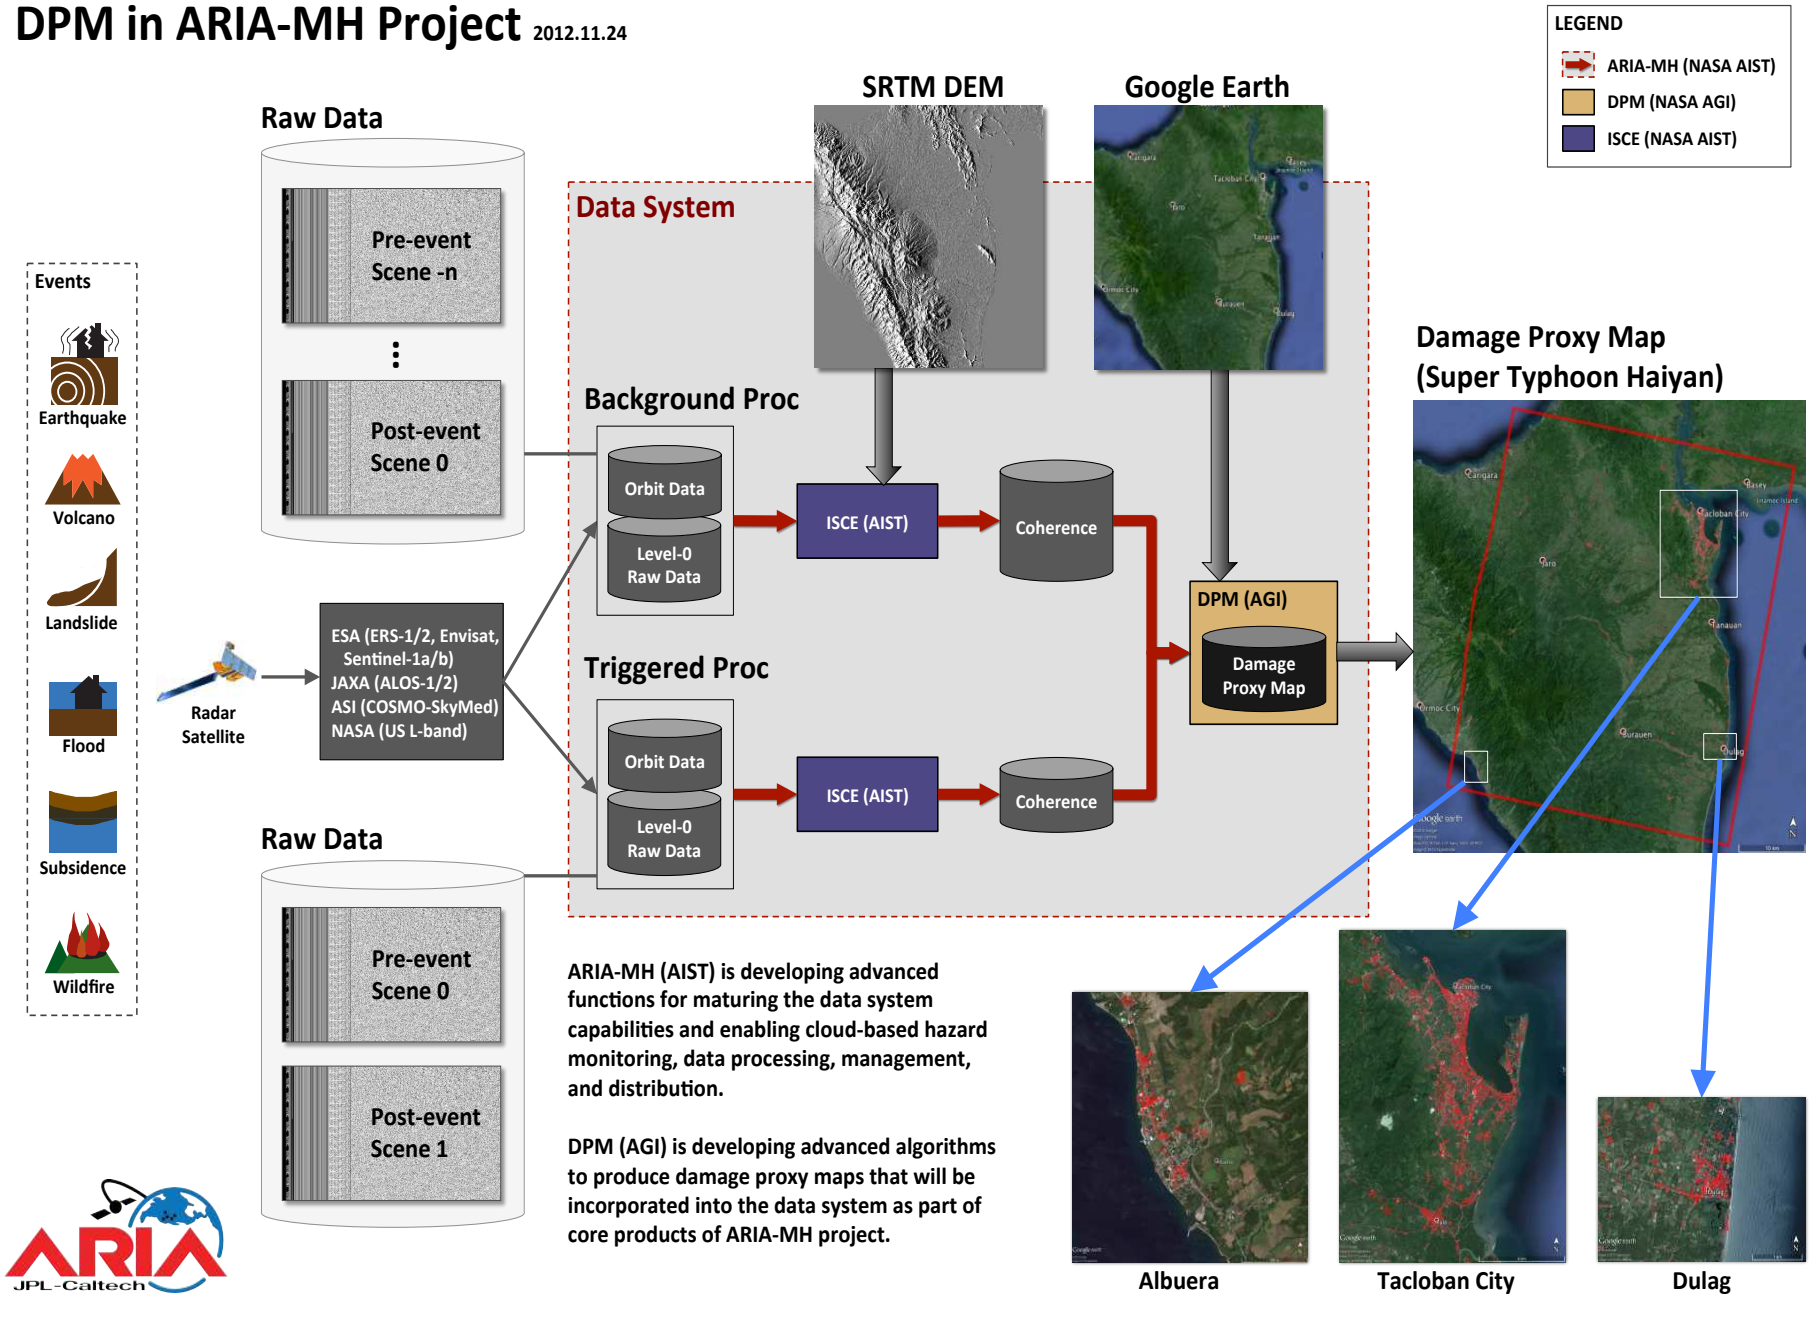

Motivation for Cloud Computing

- Incoming flood of data volume
  - Nominally 1.2GB/scene
  - 100Ks of scenes
  - order of 10GBs-100GBs temp storage per data product processing
  - PBs total storage of products
  - Example InSAR satellites
    - COSMO-SkyMed (CSK) and CSK second generation data from ASI
    - Sentinel 1A/1B
    - ALOS-2
    - Decadal Survey: proposed NI-SAR mission (US L-band SAR)
- "Embarrassingly parallel" data product generation
- Monitoring
  - User definable bounding box regions of interest for nominal background monitoring/processing.
- Elasticity of computing when responding to events
- Process migration to geographically disperse data centers
  - ESDIS DAACs (e.g. ASF)
  - UNAVCO SAR Archive
  - ASI for CSK
  - DLR for TerraSAR-X
  - JAXA
  - Various GEO Supersites

Compute, Data, & Cost Estimates

- Notional analysis comparing local hardware purchase versus AWS GovCloud usage
  - Process 16-days of data in at most 8-days.
    - 26K compute hours on 8-core nodes (3 years wall-clock processing)
  - EC2 instances with
    - Persistent EBS for cached data
    - Ephemeral local VM disk for scratch disk
  - Use AWS Glacier for cheaper long term storage (with lower data access latency)
  - Break-even point before 1-year mark
  - AWS market prices frequently changes *(estimate already outdated)*
  - Costs based on today's dollars. Does not account for inflation. Does not account for future AWS price drops.
  - On-premise costs do not consider overhead costs of cooling and electrical

Cumulative months

Local Cluster (Cumulative)		3	6	9	12	15	18	21	24	27	30	33	36
Long-Term Storage	$460.1	$2.7	$11.9	$27.5	$49.7	$78.3	$113.4	$155.0	$203.0	$257.6	$318.6	$386.1	$460.1
Transfer Fee for LTS data	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0
Transfer Fee for results	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0
CPU	$135.0	$11.3	$22.5	$33.8	$45.0	$56.3	$67.5	$78.8	$90.0	$101.3	$112.5	$123.8	$135.0
Total		$14.0	$34.4	$61.3	$94.7	$134.6	$180.9	$233.7	$293.0	$358.8	$431.1	$509.9	$595.1
Amazon GovCloud (Cumulative)													
Long-Term Storage		$0.21	$1.1	$3.2	$7.0	$13.0	$21.6	$33.4	$48.9	$68.5	$92.8	$122.2	$157.3
Transfer Fee for LTS data		$0.97	$5.2	$15.1	$32.9	$61.0	$101.7	$157.3	$230.2	$322.6	$436.9	$575.4	$740.5
Transfer Fee for results		$0.00	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0	$0.0
CPU		$6.22	$18.7	$37.3	$62.2	$93.3	$130.6	$174.2	$223.9	$279.9	$342.1	$410.6	$485.2
Total		$7.40	$25.01	$55.65	$102.14	$167.31	$253.96	$364.92	$503.01	$671.05	$871.85	$1,108.23	$1,383.01

Hybrid Cloud Computing

- Utilizes both *on-premise* and *off-site* infrastructure
- "Sunken costs": leverage *existing* infrastructure investment
- Hybrid data system architecture
  - Do "*keep up*" processing on-premise
  - "*Burst out*" to public cloud when demand exceeds on-premise resources
  - AWS-compatible Eucalyptus cloud stack
- Vendor lock in: decrease dependency on commercial vendors
- Recognize that not all resources need to be in the public cloud
  - ITAR / EAR99 software
  - AWS: GovCloud US
  - Currently high cloud storage costs

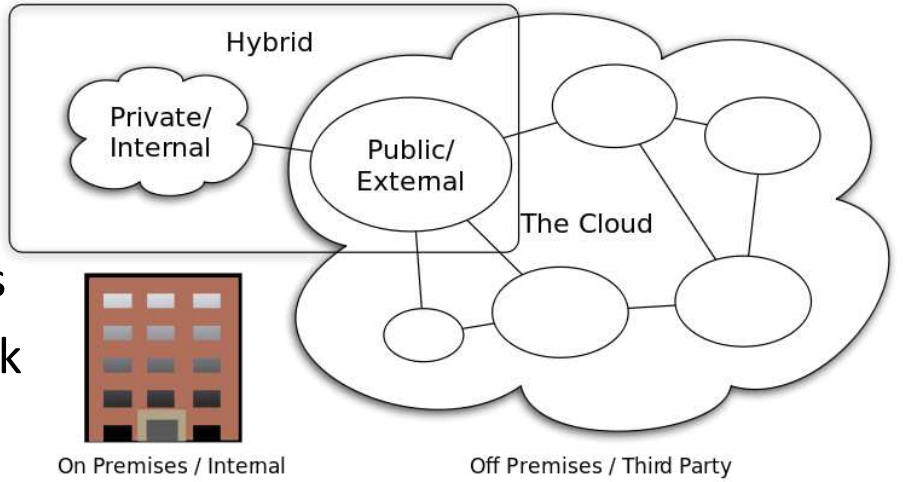

Cloud Computing Types

AWS and Eucalyptus Compatibility

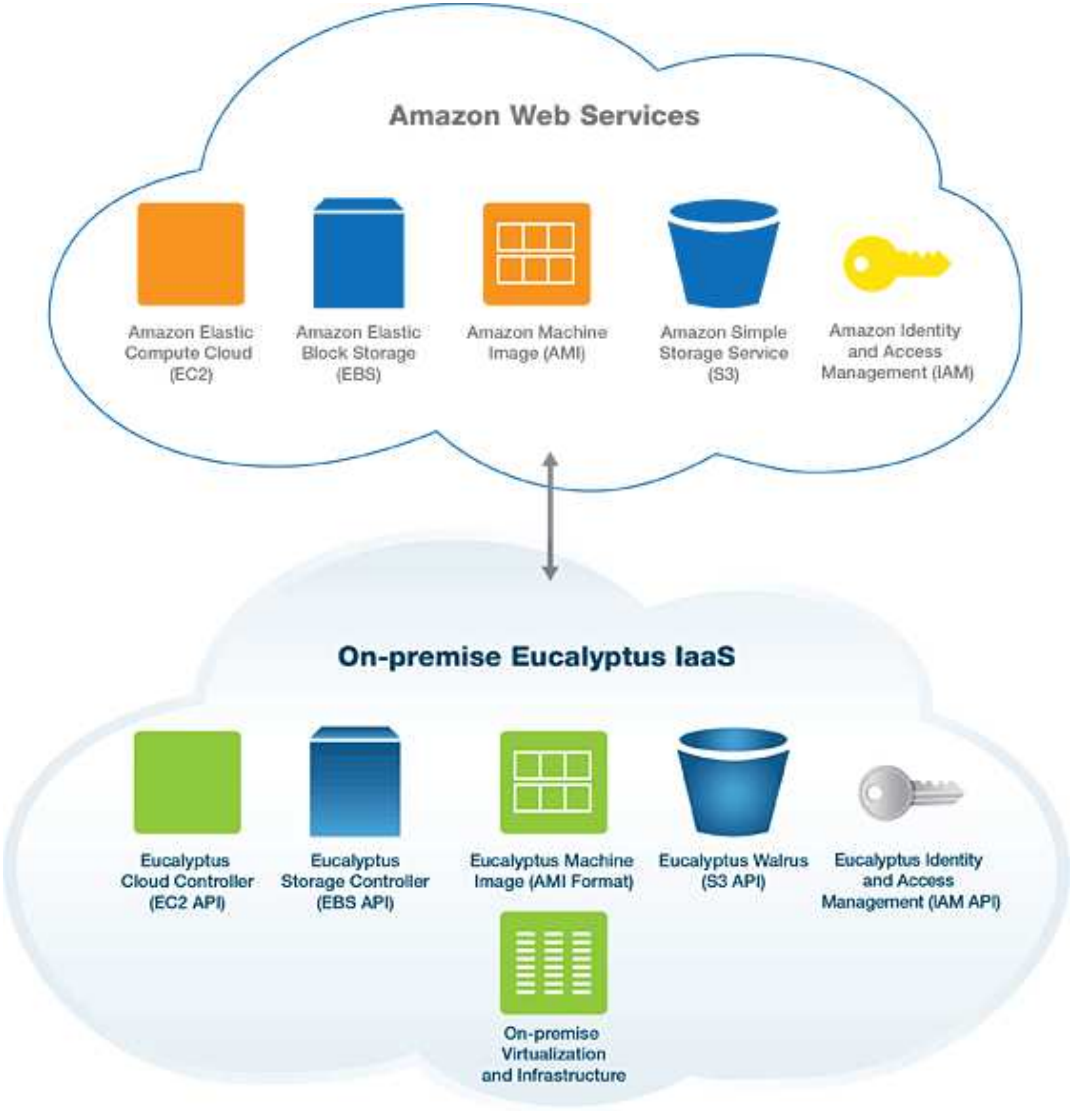

Hybrid Cloud Computing Strategy

- Utilizes both **on-premise** and **off-site** infrastructure
  - Leverage existing infrastructure investment
  - PB-scale processing and storage in public cloud currently too expensive
- Hybrid Cloud data system architecture
  - **Burst out** to public cloud when demand exceeds on-premise resources
  - Deploy AWS-compatible Eucalyptus cloud stack **on-premise**
- **Heterogeneous** computing nodes
- Resource management and data discovery can run anywhere
- Deploy **localized data repositories** closer to processing VMs
- Leverage **Amazon GovCloud US** to address export control and firewall security issues

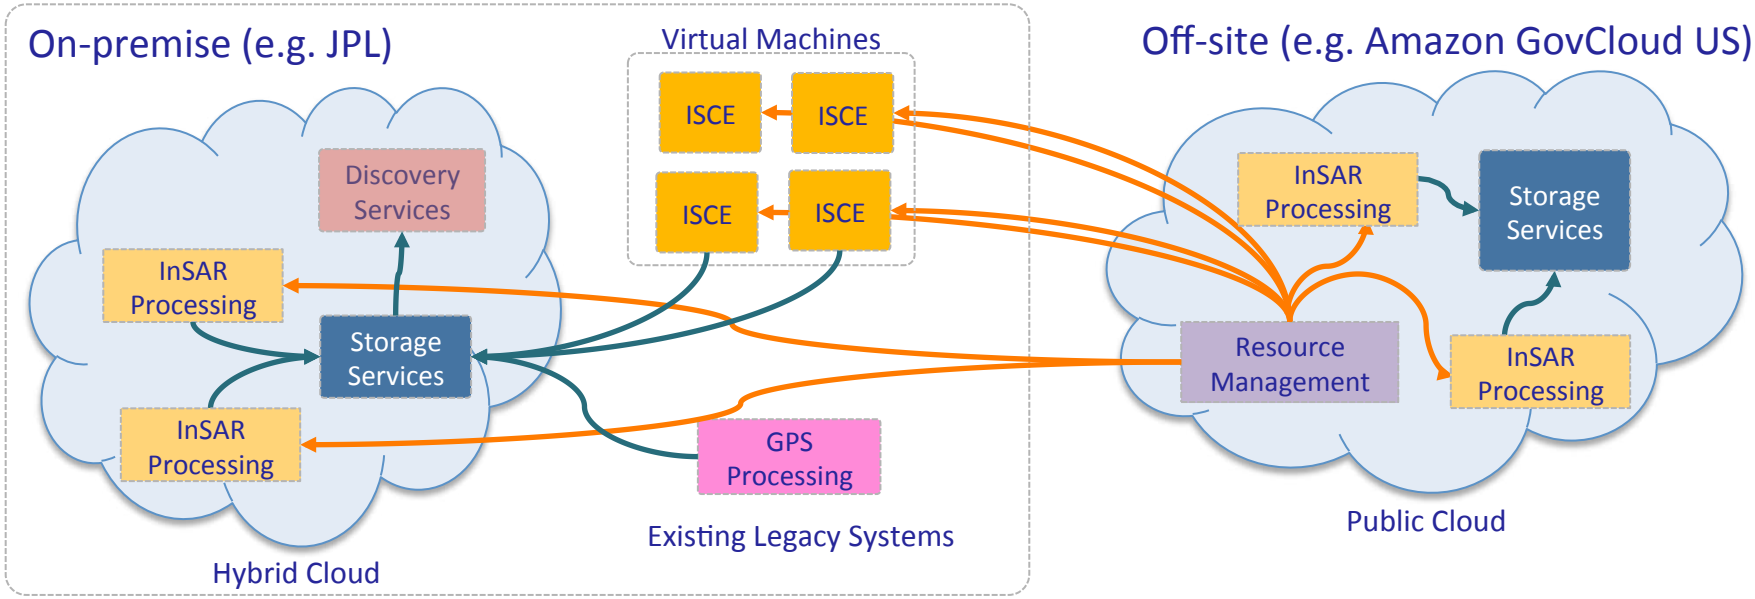

Hybrid Cloud Auto-Scaling

- *Auto-scaling*: burst out to other cloud zones

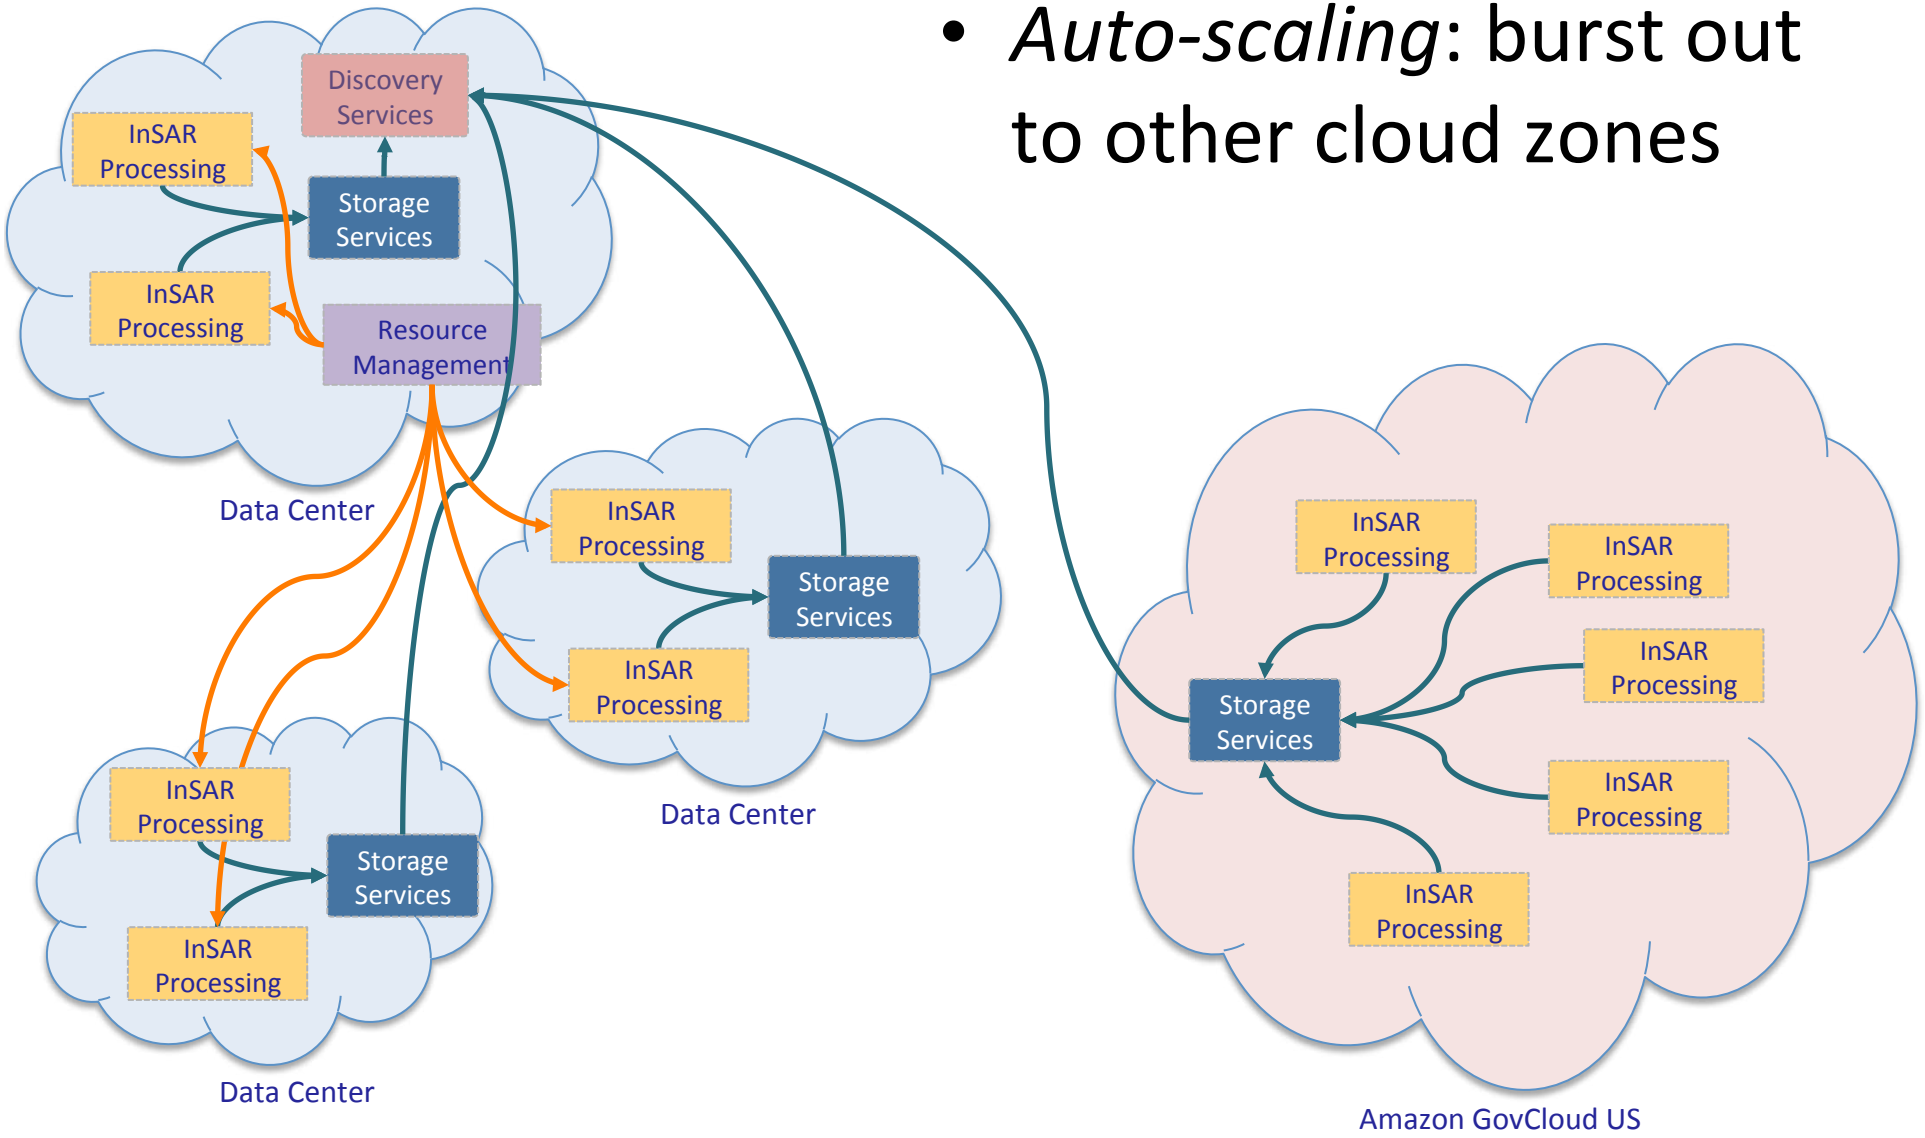

Leveraging Spatial Locality of Data

- **Want to minimize data movement time and costs**
- Leverage data locality on each compute node
- Run entire workflow pipeline on same compute node
  - Minimizes data movement for job staging in each processing step
- Can map and *cache* data on compute nodes
- Run storage repository close to compute nodes

- *Everything runs in parallel.*
- *Run pipeline on same node.*

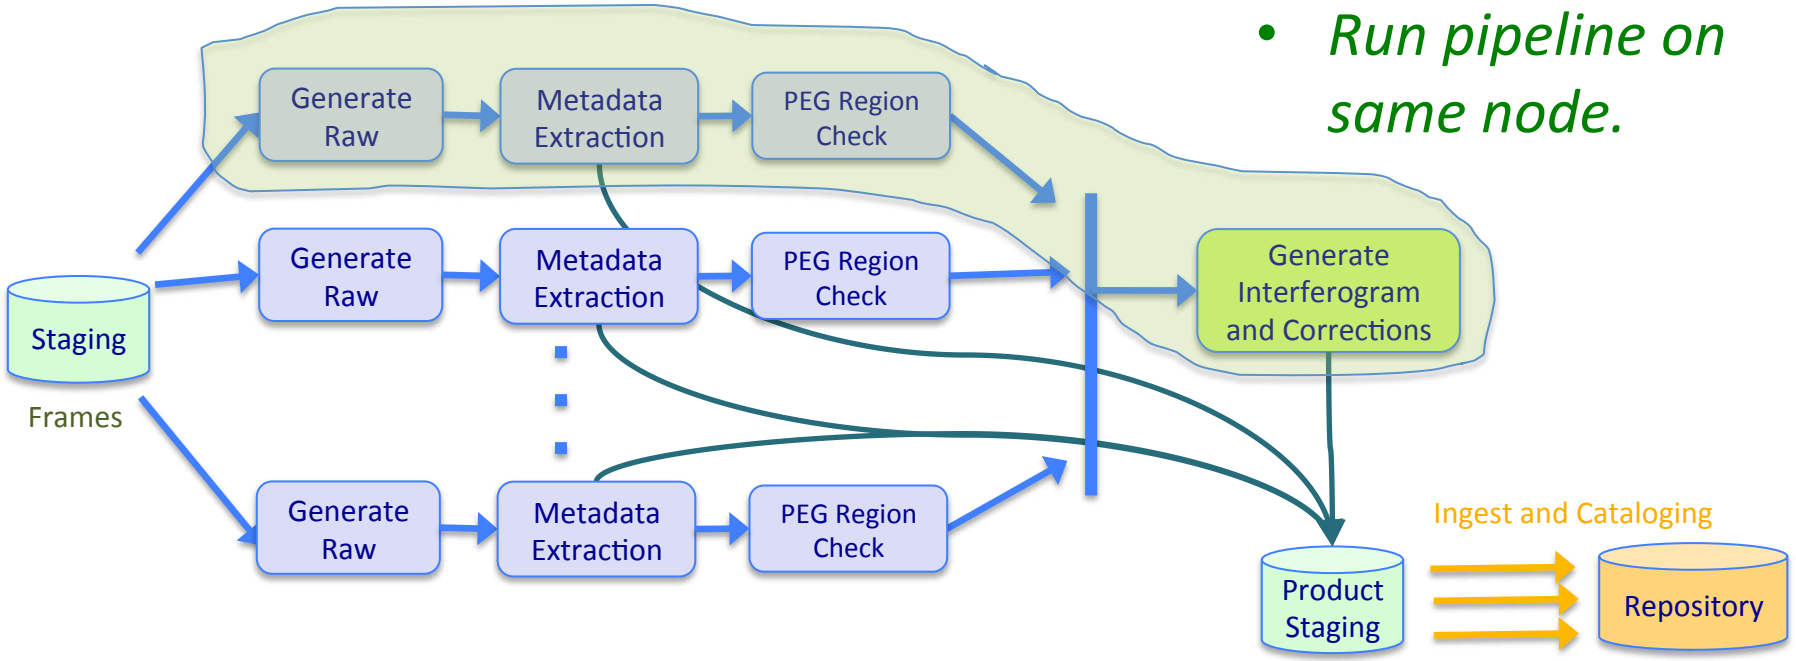

Faceted Resource Management

- Facetted view of distributed cloud compute jobs
- Distribution of job
  - Status
  - Types
  - Compute nodes
  - Hardware resource utilization
- Live DAV view

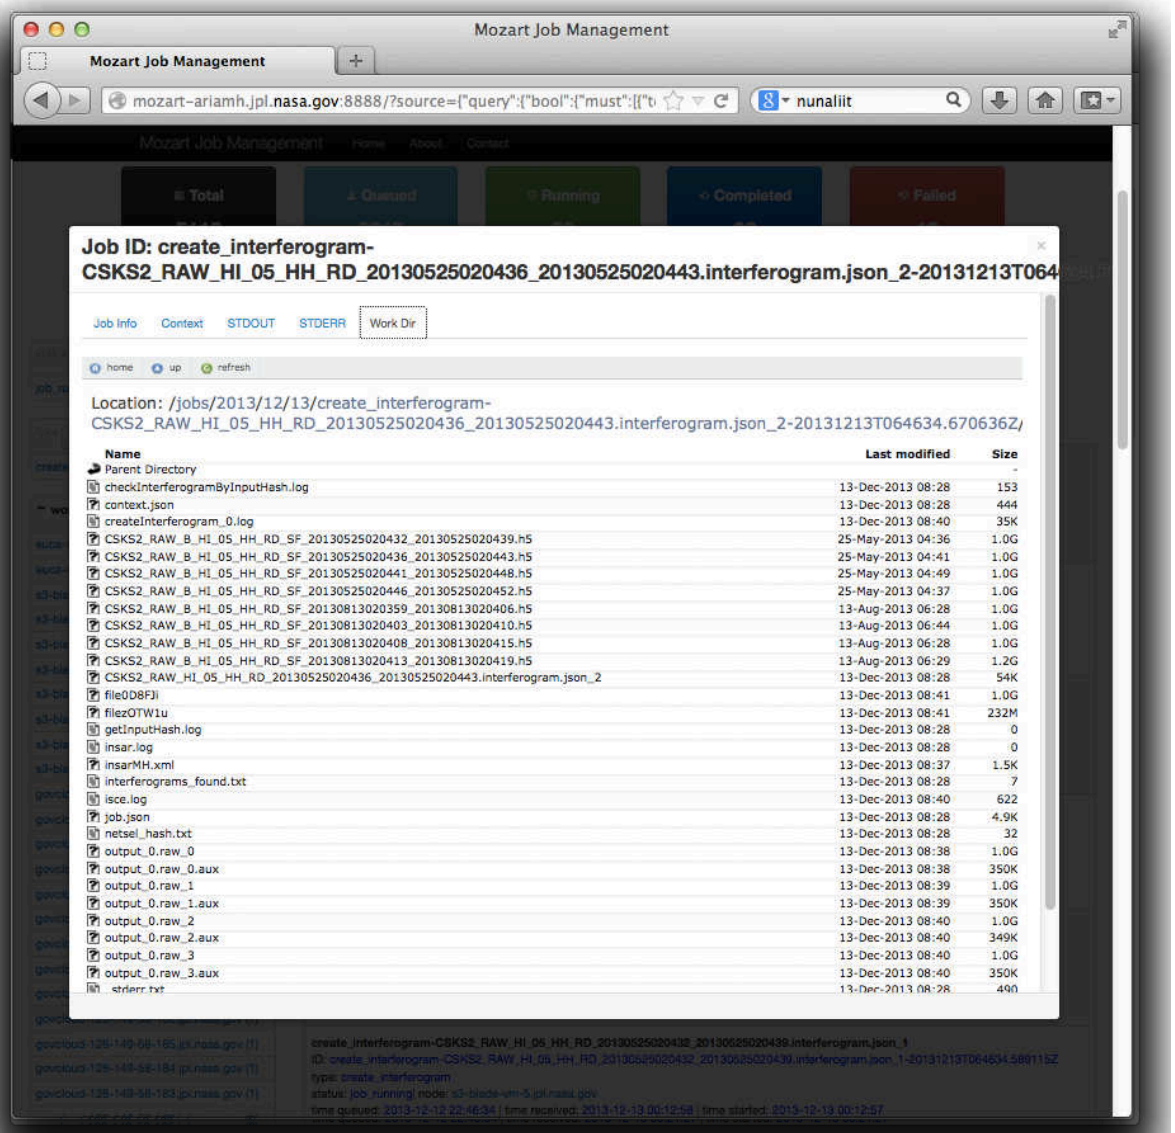

Faceted Resource Management

- Facetted view of distributed cloud compute jobs used **in additional to AWS and Eucalyptus console**

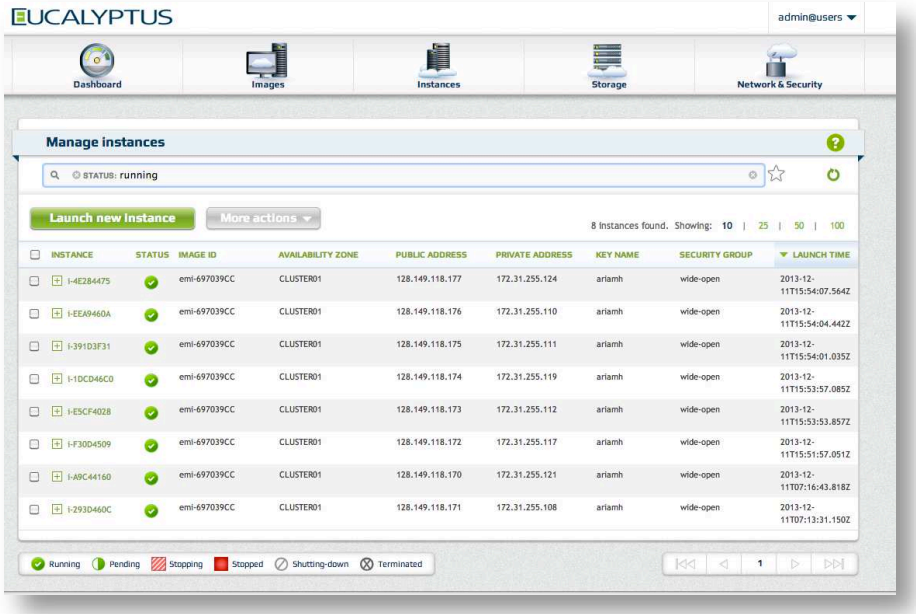

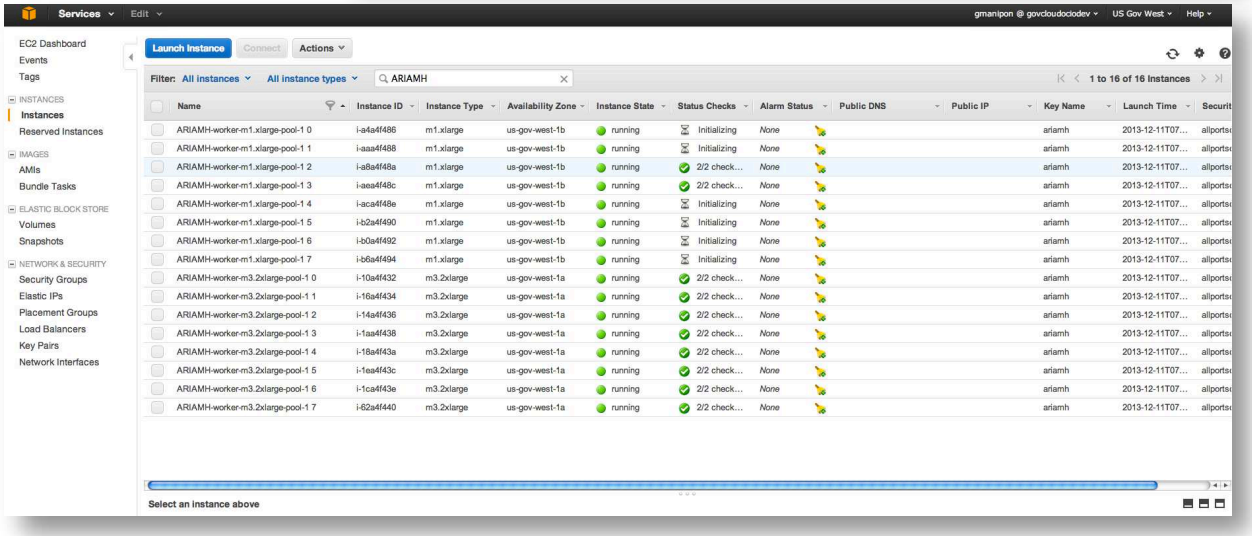

Facetted Data Product Repository

- Facetted view of data products
  - Enable users to "drill down" into multi-dimensional facets of data
  - Sequentially applies constraints
- Facets for *reverse geocoding* for all data products
- Interactive *leaflets*

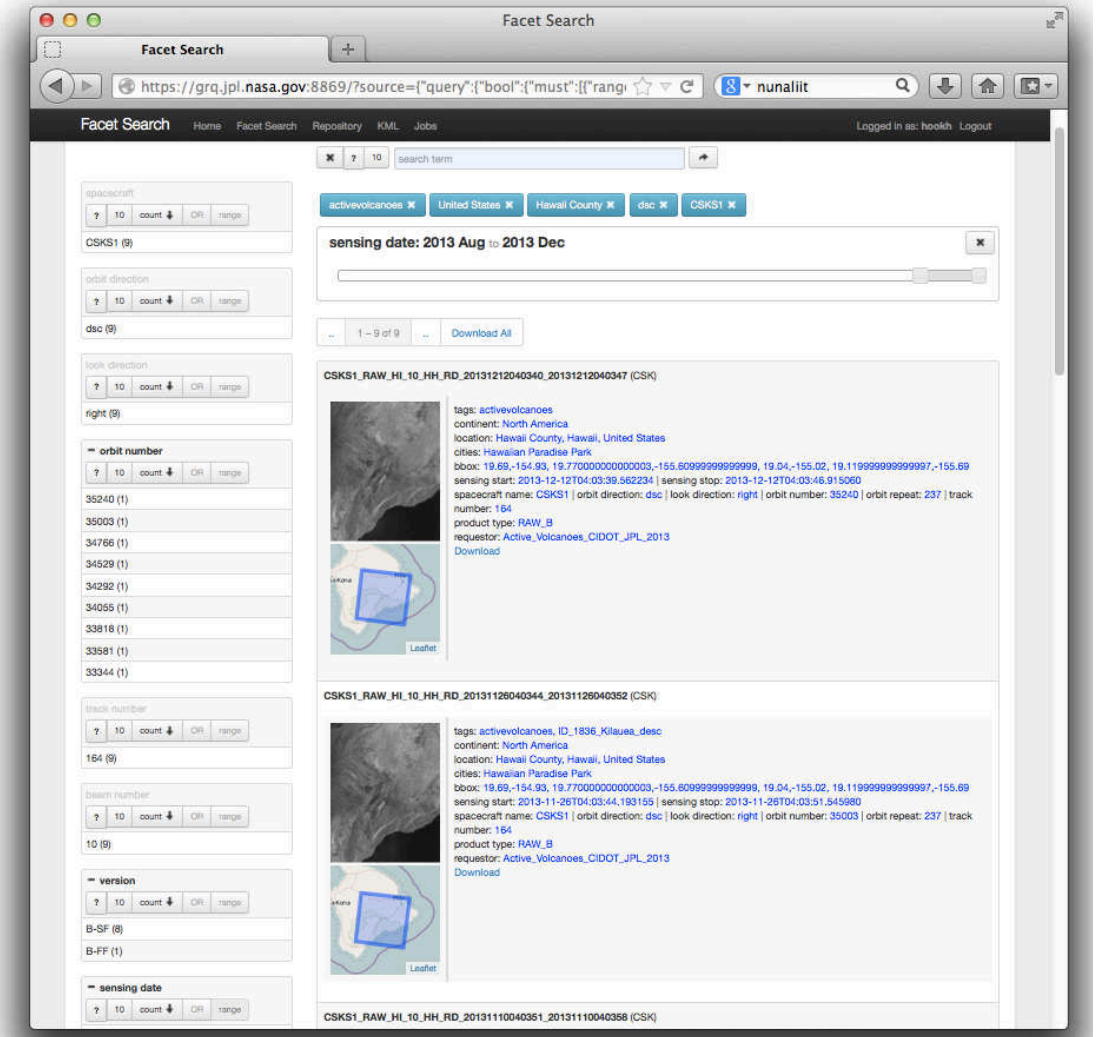

Data Access

- Leverages interoperable WebDAV specification
- Extended from prior ACCESS-2011 work
- Enable access to data products via various interfaces

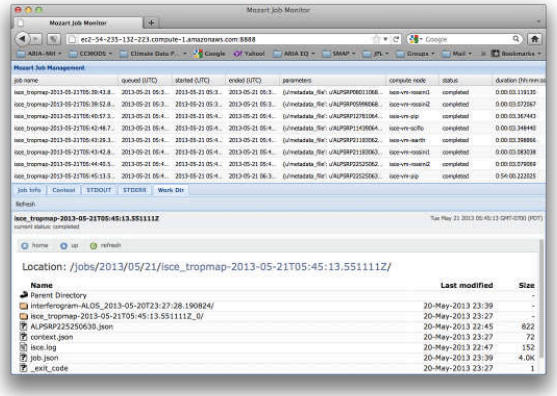

Job Management: browse remote jobs

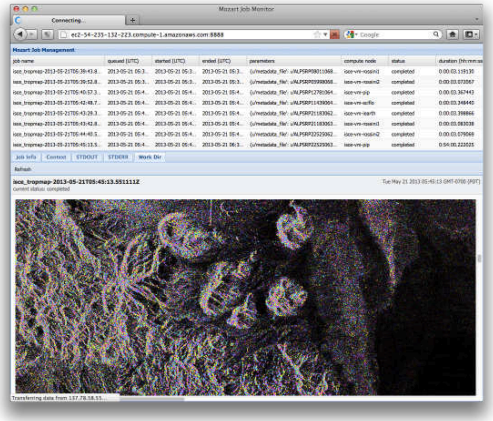

Job Management: browse “first looks”

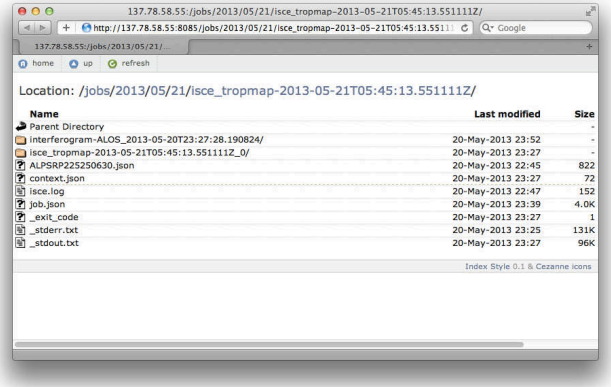

Web Interface to Repository of Data Products

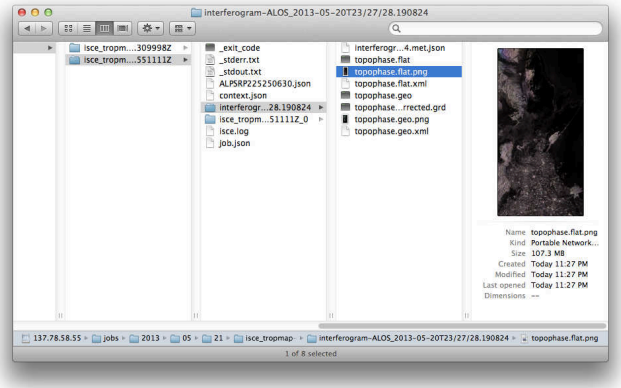

MACOSX Finder: drag-and-drop to local file system mount

Command-line interface

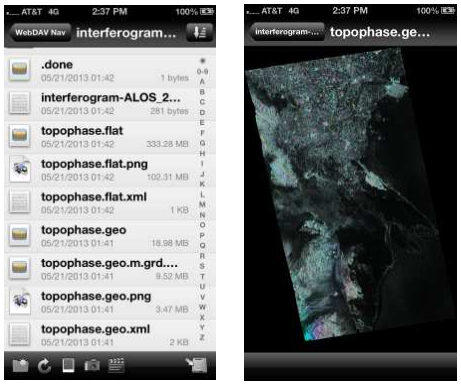

Mobile phone interface

Conclusions

- The global coverage offered by satellite-based SAR missions, and rapidly expanding GPS networks can provide orders of magnitude more observations and improve hazard response
  - …if we have a data system that can efficiently monitor and analyze the voluminous data, and provide users the tools to access data products.
- Hybrid cloud may be more effective for these needs
  - Do "keep up" processing **on-premise**
  - **Off-site** elasticity (bursting)
- Hybrid cloud computing may be a more effective approach to addressing lower latency and high data volume processing needs
- PB-scale data volumes in cloud computing and storage significant enough to **affect architectural design of data system**
- Lowest common denominator of services
  - Compute, Block storage, Bucket Storage, VM, and Identity
- Package management (e.g. Puppet/Chef) key to ease of deployment across cloud stacks
- Monitoring of events for automatic processing
- Monitoring of data products for custom actions