

IEEE e-Science 2010
December 9, 2010
Brisbane, QLD, AU

Fault Tolerance and Scaling in e- Science Applications: Observations from the Continuing Development of MODIS Azure

*Keith R. Jackson*¹, Deb Agarwal¹, You-Wei Cheah², Marty
Humphrey³, Catharine van Ingen⁴, Jie Li³, Youngryel Ryu⁵

¹Lawrence Berkeley National Lab

²Indiana University

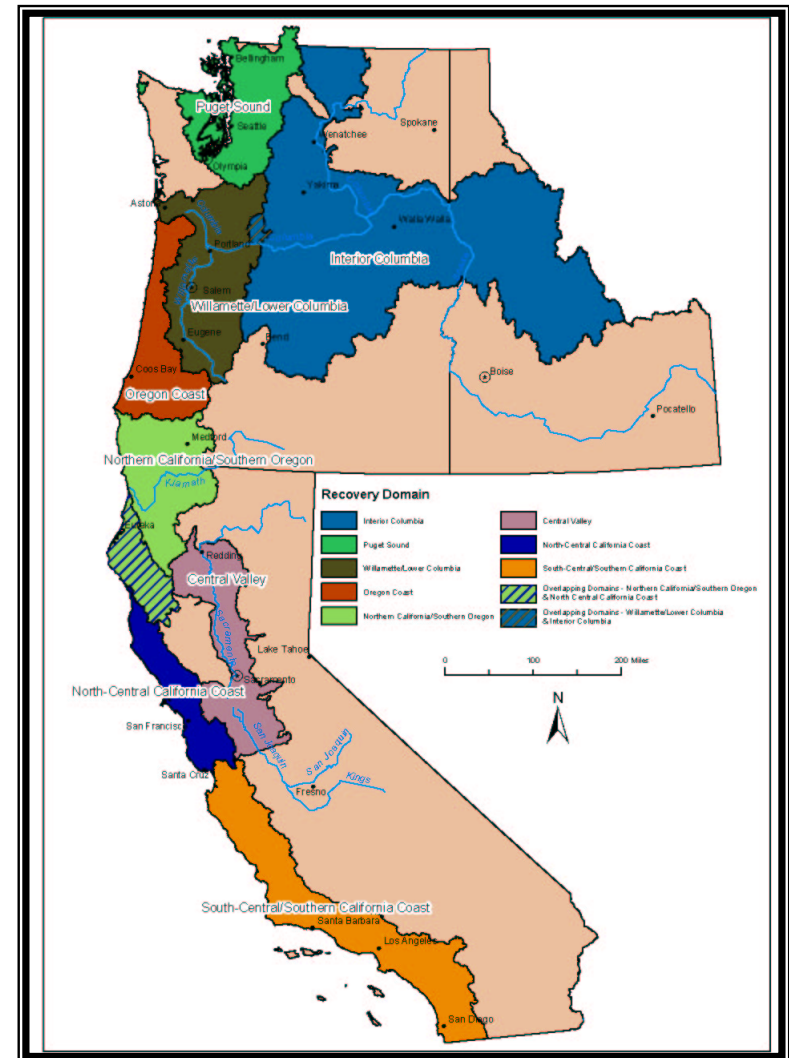
³University of Virginia

⁴Microsoft Research

⁵University of California, Berkeley

Global and Regional Eco-Science

- A major shift is happening in the way eco-science is done.
 - Moving from individual studies of local processes to collaborative studies of regional and global processes. (e.g. studying the impact of climate change)
- Studying global scale environmental processes requires:
 - Integration of local, regional, and global spatial scales.
 - Integration across disciplines, e.g., climatology, hydrology, forestry, etc., and across methodologies (field observations, remote sensing, and modeling).



Global and Regional Eco-Science



Manual Measurement




Automated Measurement



Sample Collection




Typing



Satellite



Aircraft Surveys



Model Output

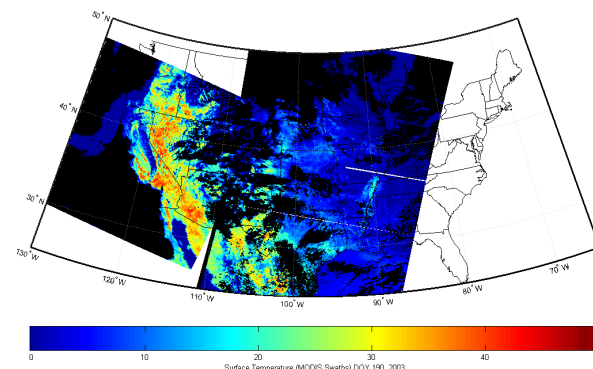


Counting



CS Grunge

- Data download
 - Locate download site for each data product
 - Validate your download
- Reprojection
 - Converts one geo-spatial representation to another.
 - Example: latitude-longitude swaths converted to sinusoidal cells.
- Spatial & Temporal resampling
 - Converts one spatial resolution to another.
 - Example is converting from 1 KM to 5 KM pixels.
- Gap filling
 - Assigns values to pixels without data either due to inherent data issues such as clouds or missing pixels.

Source Data (Swath format)



Reprojected Data (Sinusoidal format – equal land area pixel)



Sinusoidal
US cells

Evapotranspiration

- Evapotranspiration (ET) is the release of water to the atmosphere by evaporation from open water bodies and transpiration, or evaporation through plant membranes, by plants.
- Climate change isn't just about a change in temperature, it's also about a change in the water balance and hence water supply critical to human activity.






Image courtesy of the
[National Oceanic and Atmospheric Administration](#)

Data Inputs

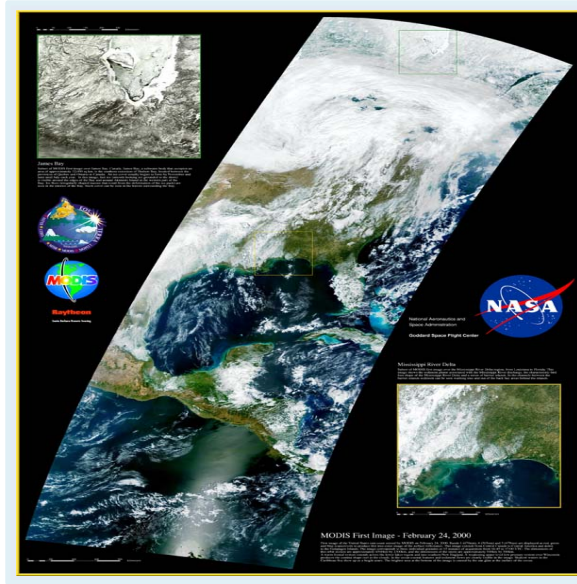
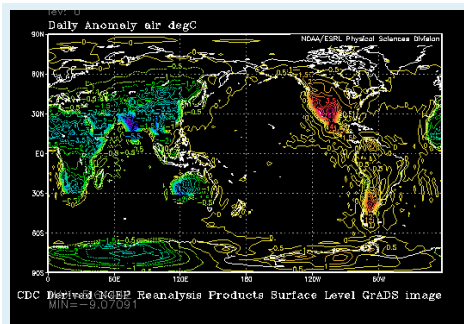


Climate classification

J.M. Chen et al. / Remote Sensing of Environment 97 (2005) 447–457



Vegetative clumping
~5MB (1file)



NASA MODIS imagery archives

**Sizes given are 1 US year
20 US year ~ 1 global land
surface year**



FLUXNET
Curated
sensor
dataset
30GB
(960 files)

FLUXNET
curated field
dataset
2 KB (1 file)

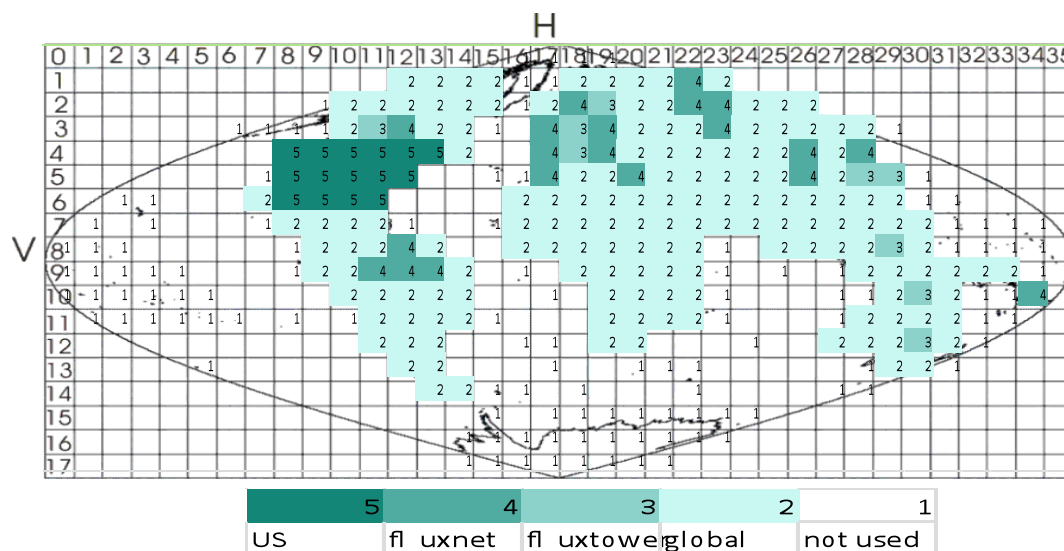


The Resource Barrier

	#Source Files	Source Size	# Result Files	Result Size
USA (18)	21850	238 GB	27375	261 GB
FluxTower (3)	80670	993 GB	58400	210 GB
Global (3)	152670	2414 GB	352225	630 GB

1 Global year


- ▶ ~ 9000 hours or 375 days
 - 1 year takes 1 year !
 - With 200 cpus, that's > 2 days
- ▶ 8 TB download reduced to 2.4 TB by reprojection
 - How to know which are missed (and why) among 350K files?



US: 15 tiles FluxTower: 32 tiles Global: 194 tiles

A Cloud-based Solution

- A MODIS Data Reprojection and Reduction Pipeline in Microsoft Windows Azure cloud computing platform
 - Lower resource entry barriers
 - Hide the complexities in data collection, reprojection and management from domain scientists
 - A generic Reduction Service for scientists to upload arbitrary executables to perform scientific analysis on reprojected data.



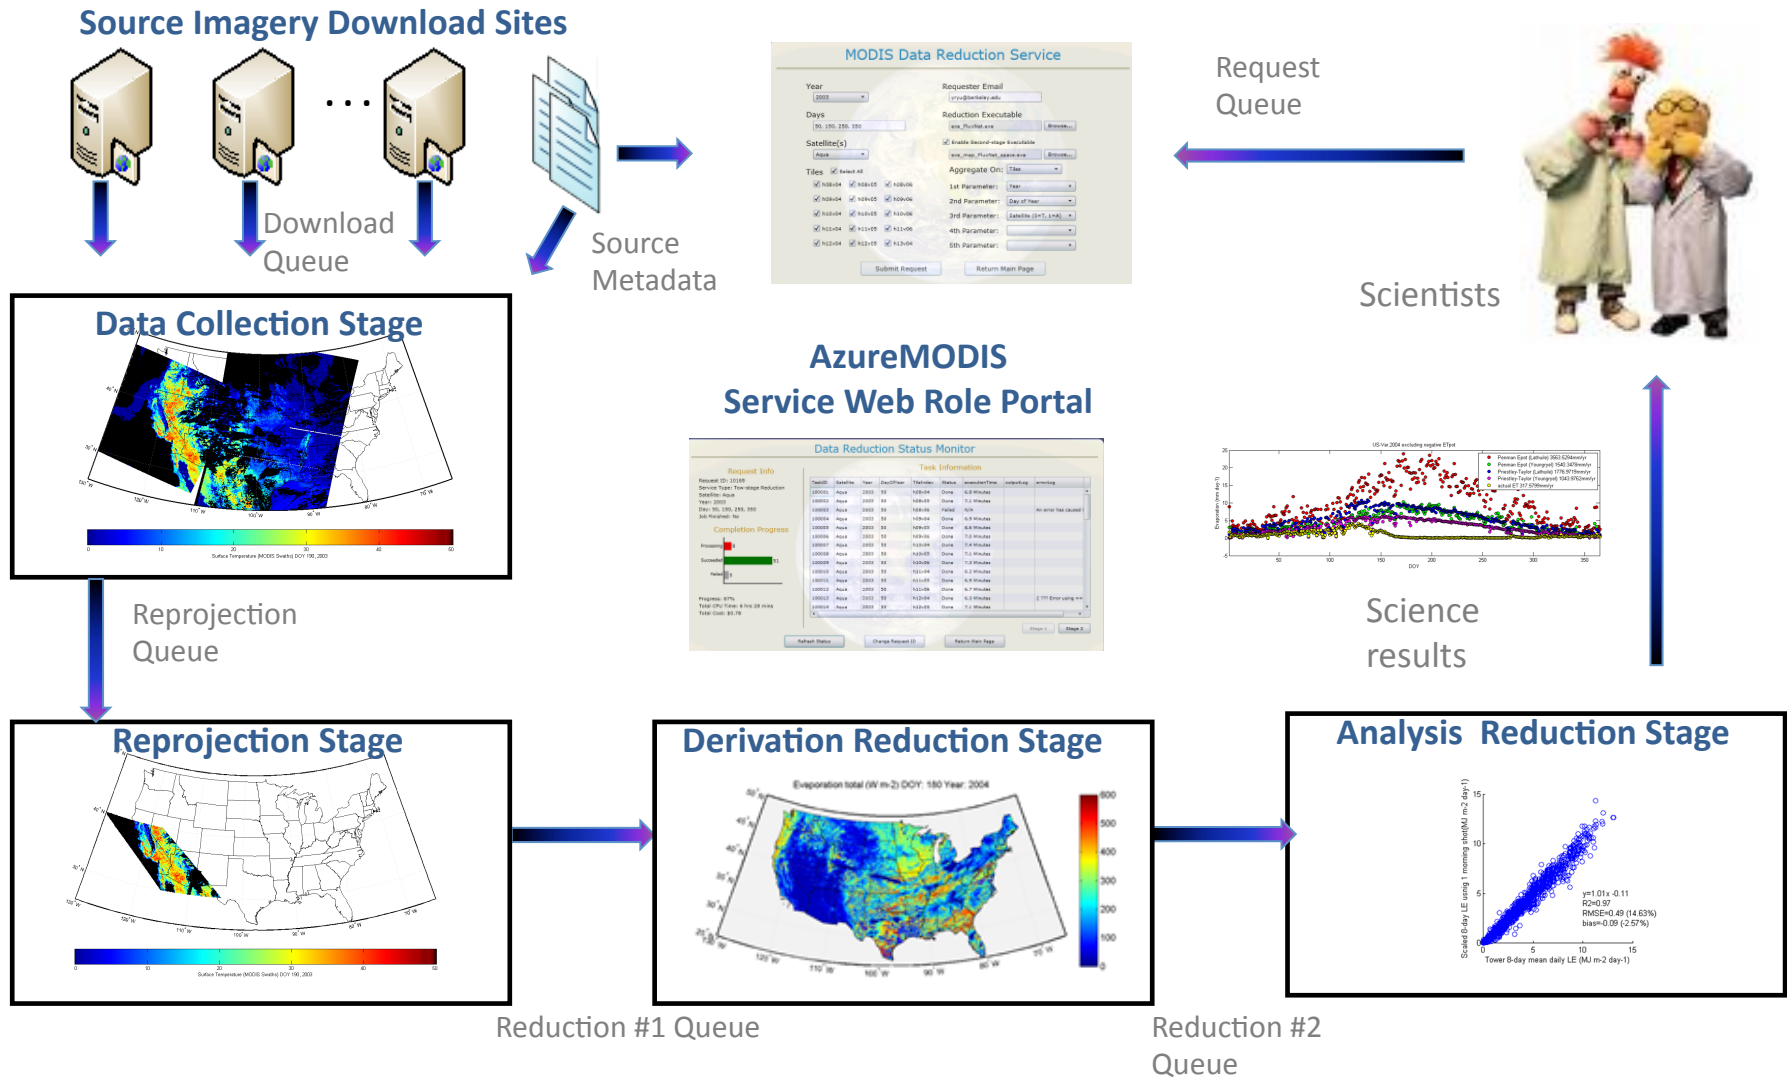
Windows Azure Basics

- Hosted Services
 - Web Role: A Web application accessible via an HTTP and/or an HTTPS endpoint
 - Worker Role: A background processing application
 - Dynamic scalability of role instances
- Storage Services
 - Blob service: Storage for entities such as binary files and data files
 - Queue Service: A reliable, persistent queue model for messaging within and between hosted services
 - Table Service: Structured storage in the form of tables, with simple query support

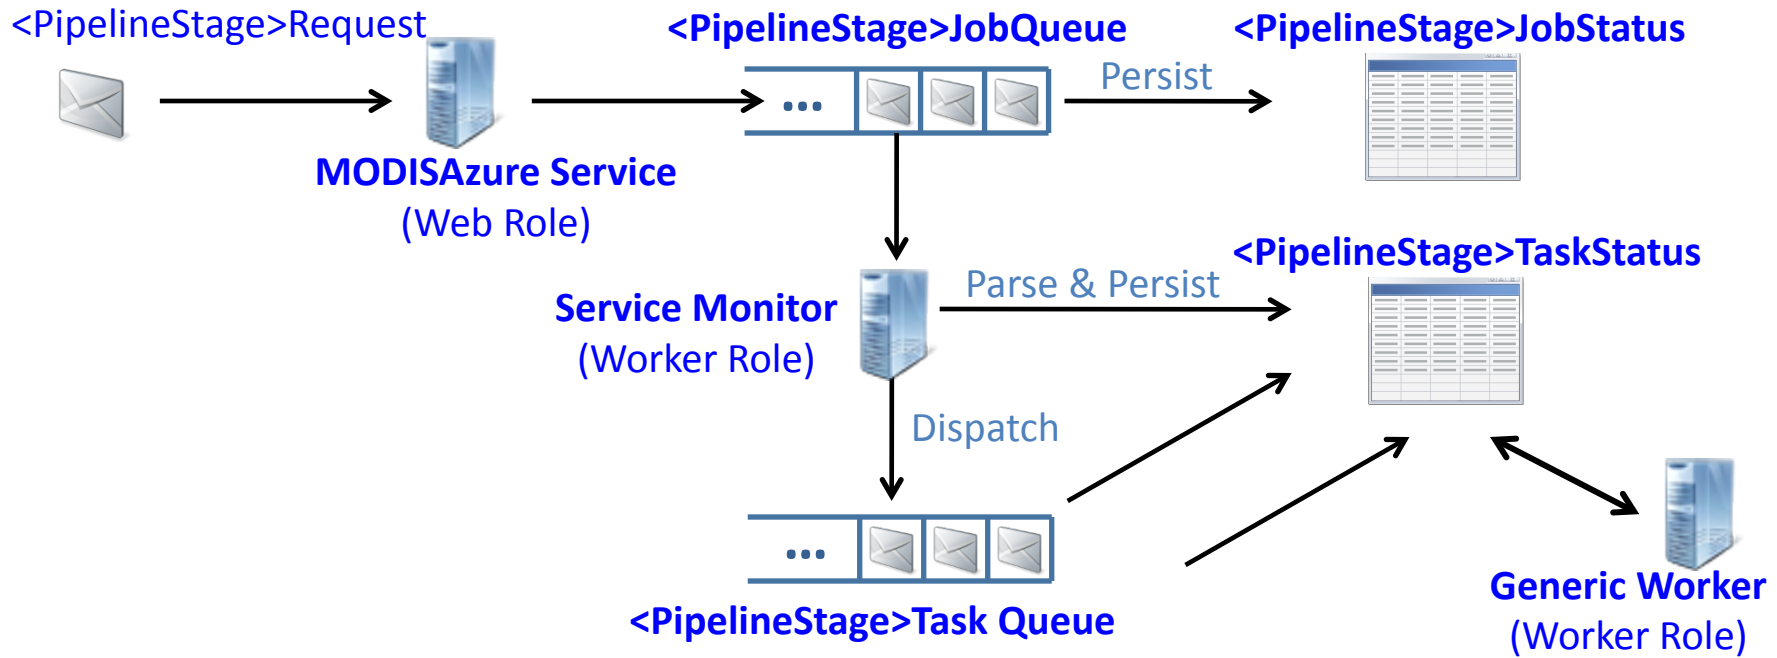


Windows® Azure™

MODIS Azure: Four Stage Image Processing Pipeline



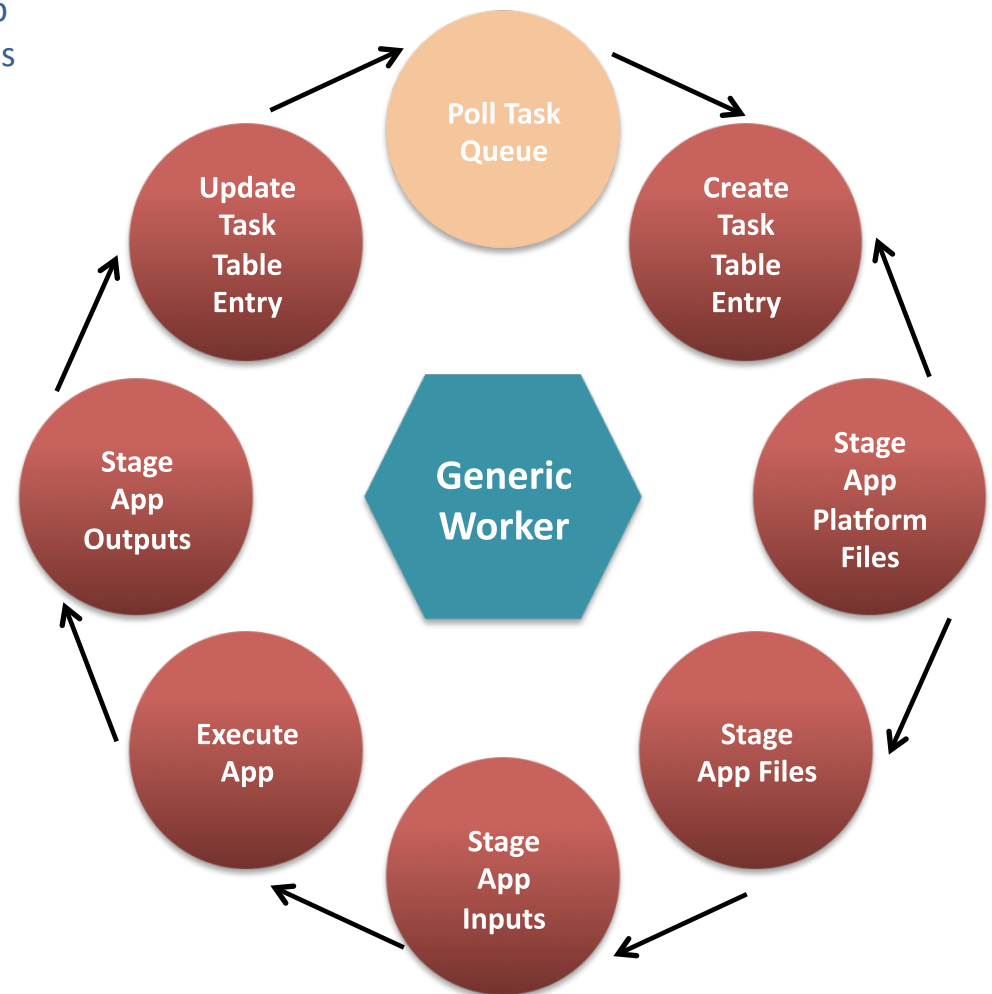
MODIS Azure: Architecture



- **ModisAzure Service** is the Web Role front door
 - Receives all user requests
 - Queues request to appropriate Download, Reprojection, or Reduction Job Queue
- **Service Monitor** is a dedicated Worker Role
 - Parses all job requests into tasks
 - recoverable units of work
 - Execution status of all jobs and tasks persisted in Tables

Generic Worker

- Manages application sandbox
 - Ensures all application binaries such as the MatLab runtime are installed for “known” application types
 - Stages all input blobs from Azure storage to local files
 - Passes any marshalled inputs to uploaded application binary
 - Stages all output blobs to Azure storage from local files
 - Preserves any marshalled outputs to the appropriate Task table
- Manages all task status
 - Dequeues tasks created by the Service Monitor
 - Retries failed tasks 3 times
 - Maintains all task status




Pipeline Scaling & Interactions


- The Web Portal Role, Service Monitor Role and 5 Generic Worker Roles are deployed at most times
 - 5 Generic Workers are sufficient for reduction algorithm testing and development (\$20/day)
 - Early results returned to scientist while deploying up to 93 additional Generic Workers; such a deployment typically takes 45 minutes
 - Deployment taken down when long periods of idle time are known
 - Heuristic for scaling number of Generic Workers up and down
- Download stage runs in the deep background in all deployed generic worker roles
 - IO, not CPU bound so no competition
- Reduction tasks that have available inputs run preferentially to Reprojection tasks
 - Expedites interactive science result generation
 - If no available inputs and a backlog of reprojection tasks, number of Generic Workers scale up naturally until backlog addressed and reduction can continue
 - Second stage reduction runs only after all first stage reductions have completed

Resource Utilization

- Fine scale computation expanded to cover more of the globe: 2x compute requirements and 2x (transient) storage requirements
- Lower resolution global computation added: .5x compute requirements and 2x (transient) storage and higher IOP/cpu reprojection
- Now underway: geo-spatial validation with yearly aggregate: shifts reduction to IO intensive




Performance Variability



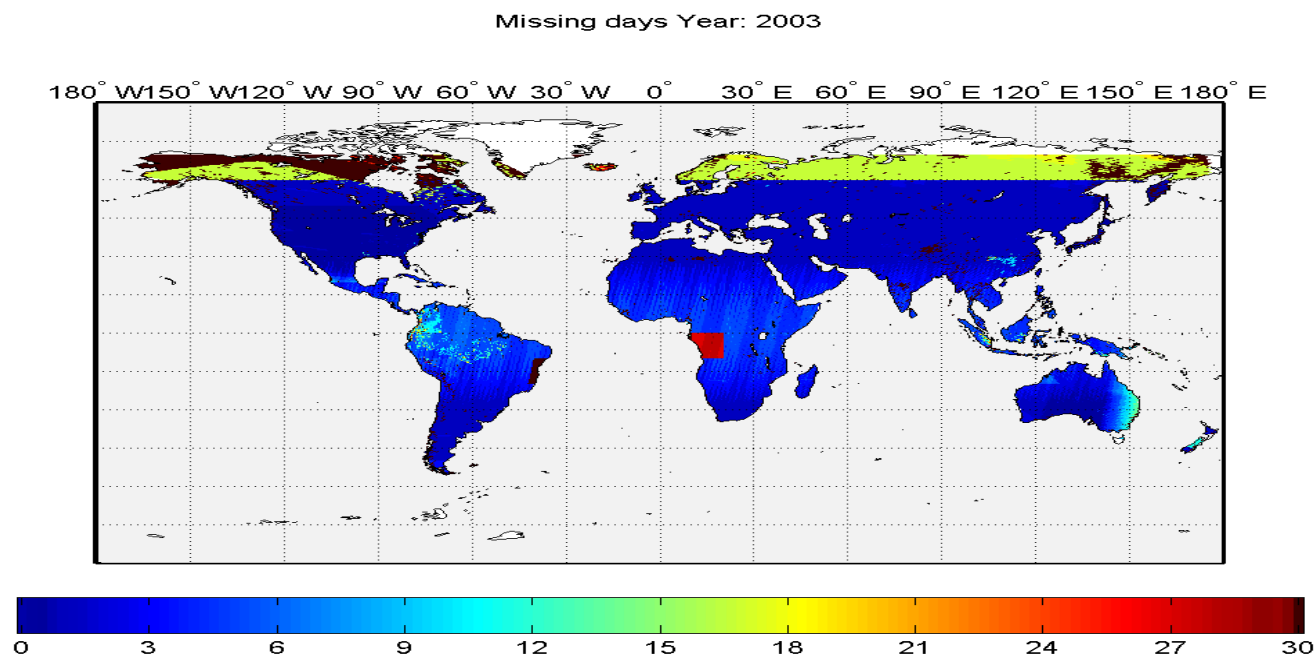
- Run the same 5 reduction tasks twice a day for 15 days
- The difference between the best performance and the worst can be as large as 350%

Infrastructure Reliability



- Instances fail to boot properly
- Temporary storage request failures
- Instances restarted

Data Reliability



- When running hundreds of thousands to millions of tasks, even 99.999% reliability is not enough!
 - 1-2% of MODIS Azure tasks fail, but succeed on retry.

Architectural Lessons

- Assume failures
 - Failures can occur at all levels
 - Log task start and end events
 - Logs must be minable
 - Tasks should be idempotent
 - Much easier failure/restart behavior
 - Simple retry behavior can be very effective
- Decouple components and communicate via reliable message queues

Conclusions

- Clouds like Windows Azure provide a set of building blocks for building scalable, reliable, distributed data-intensive e-Science applications
- Clouds can lower the barriers to doing science
 - Remove CS grunge work
 - Provide access to significant amounts of storage and compute resources
 - Provide Software-as-a-Service web interfaces to scientific data processing

Acknowledgments

- This work was funded in part by Microsoft External Research
- This work was funded in part by the Advanced Scientific Computing Research (ASCR) in the DOE Office of Science under contract number DE-C02-05CH11231



Office of Science
U.S. Department of Energy

