Re-organizing Earth Observation Data Storage to Support Temporal Analysis of Big Data

Christopher Lynnes*

NASA Goddard Space Flight Center

*NASA Civil Servant

Earth Observing System Data and Information System (EOSDIS)

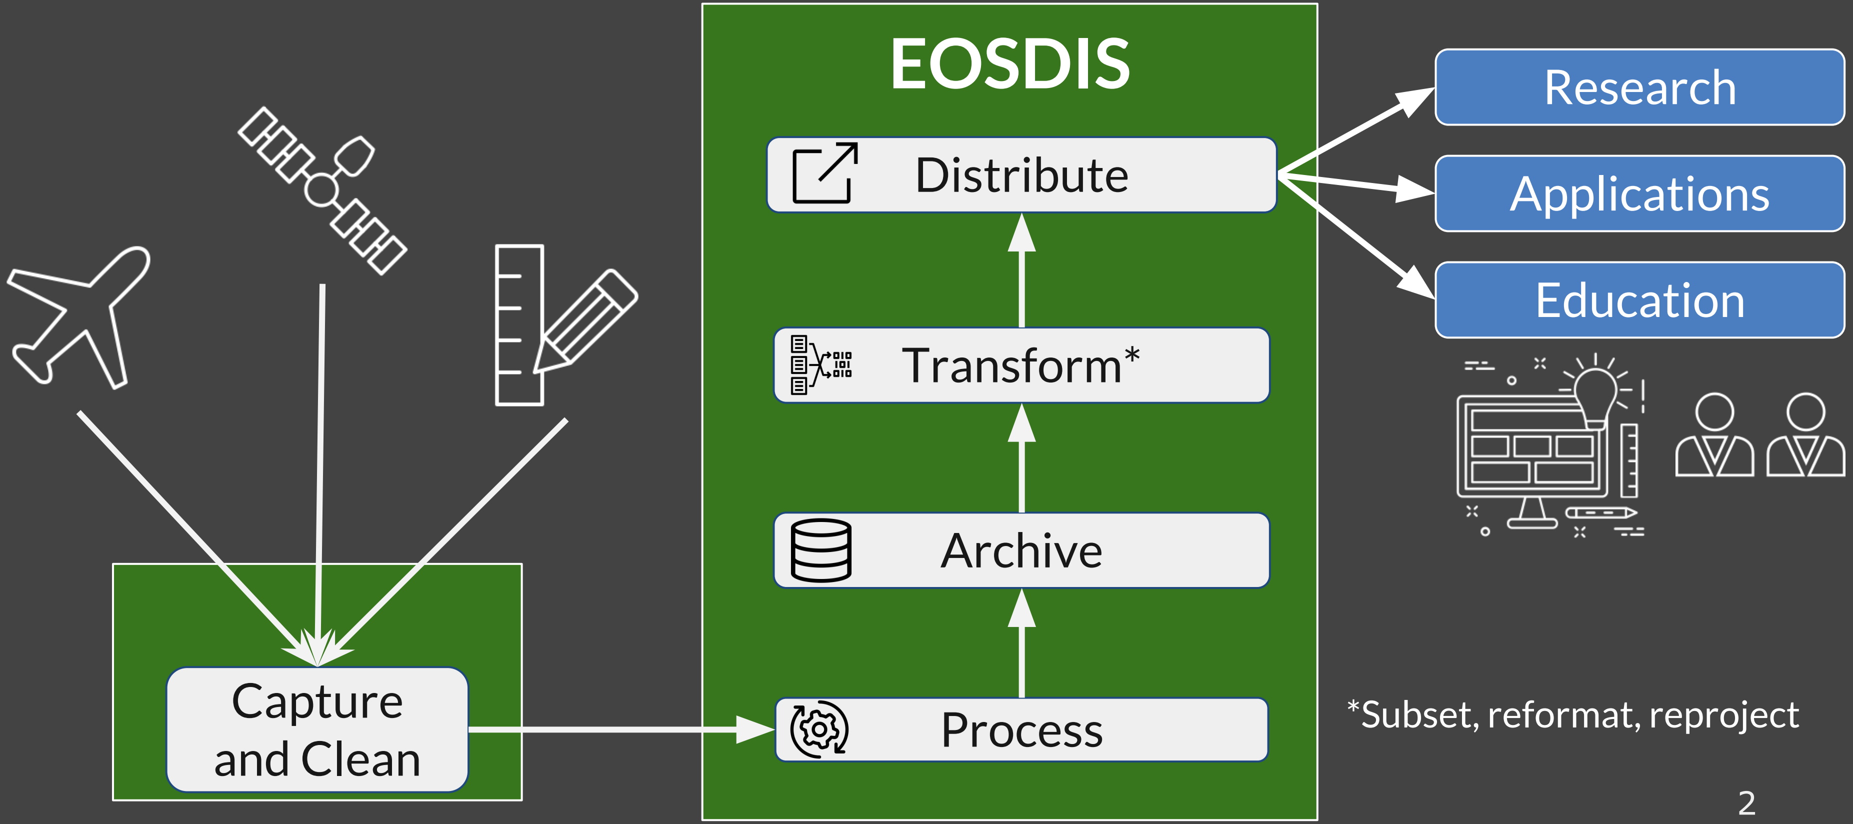

*Subset, reformat, reproject

Earth Observing System Data and Information System (EOSDIS)

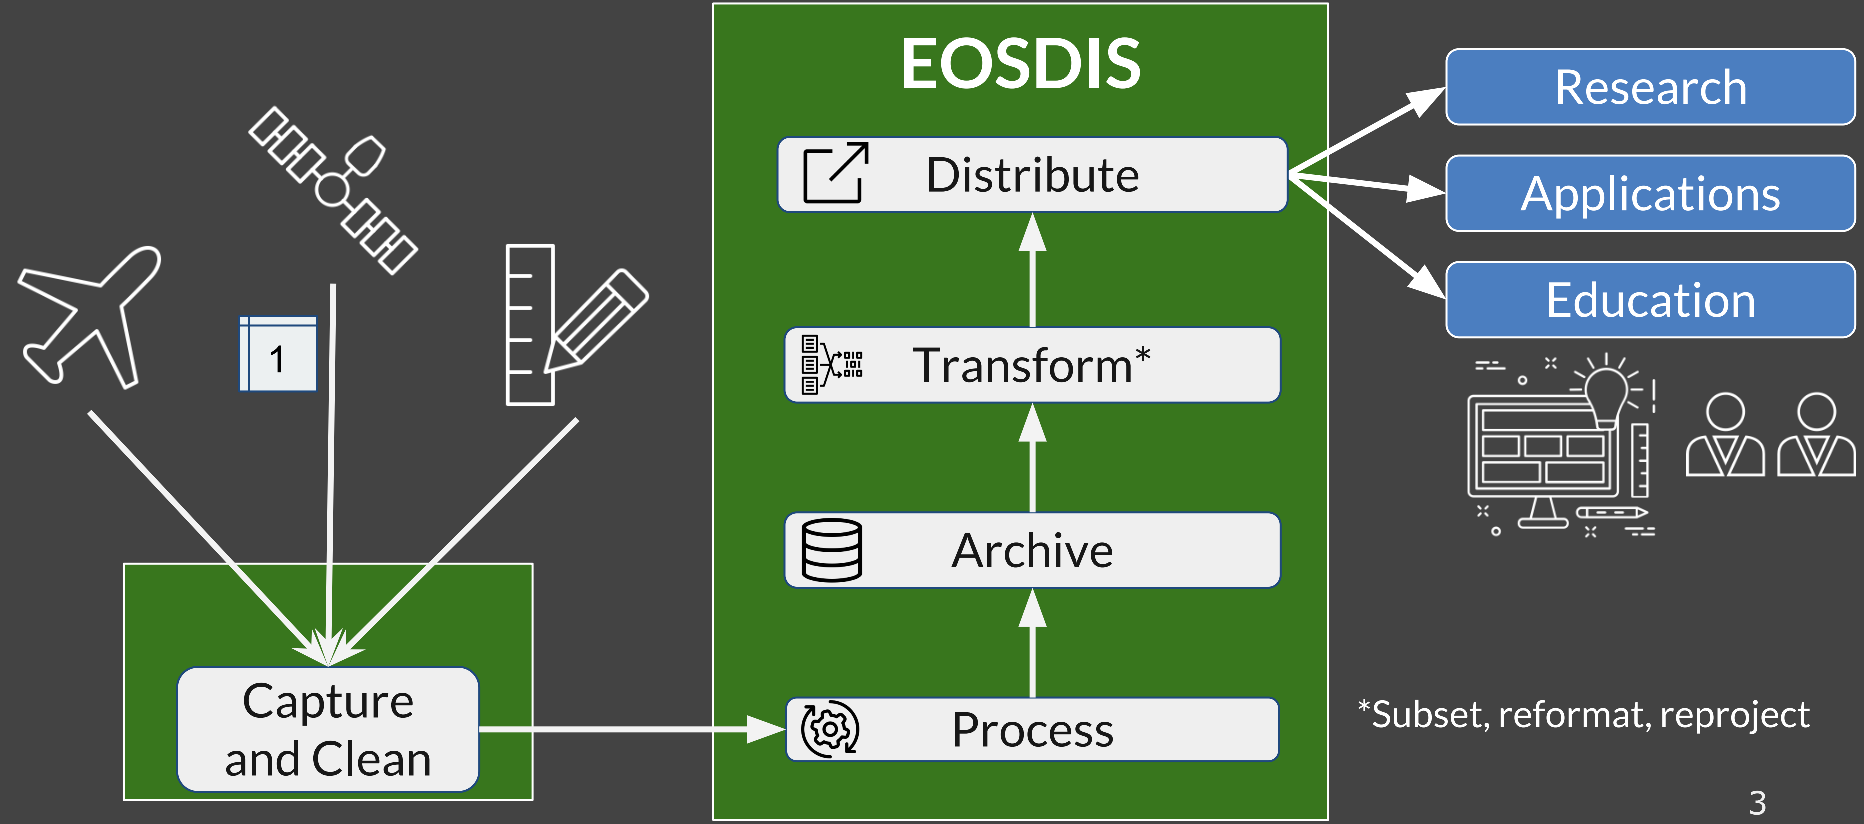

Earth Observing System Data and Information System (EOSDIS)

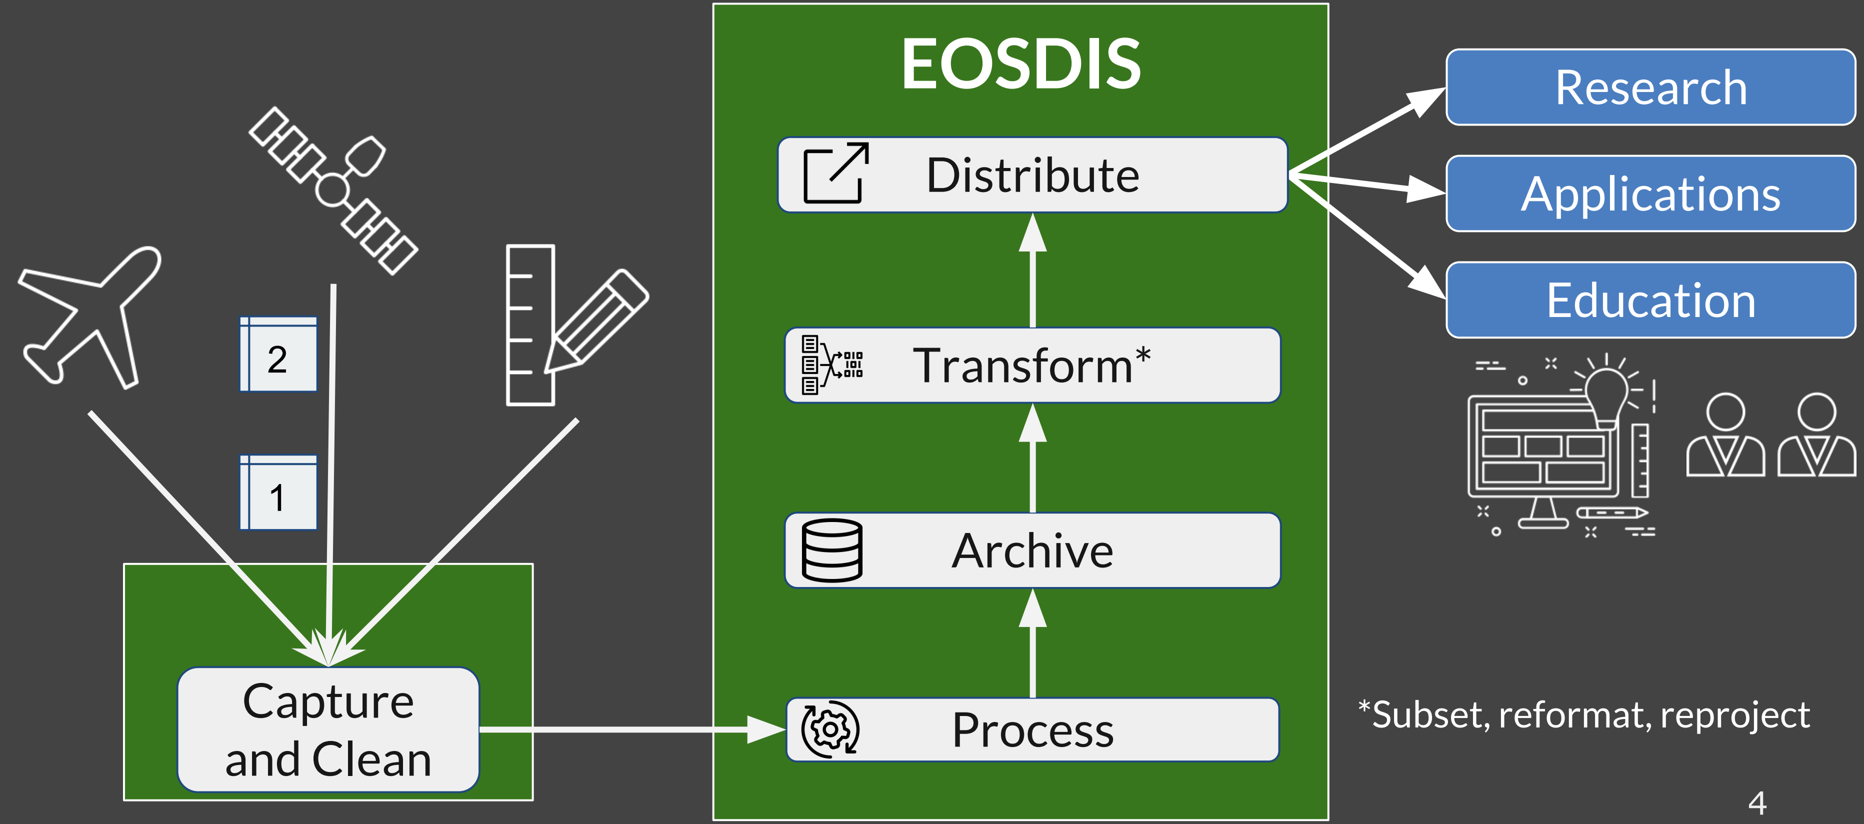

Earth Observing System Data and Information System (EOSDIS)

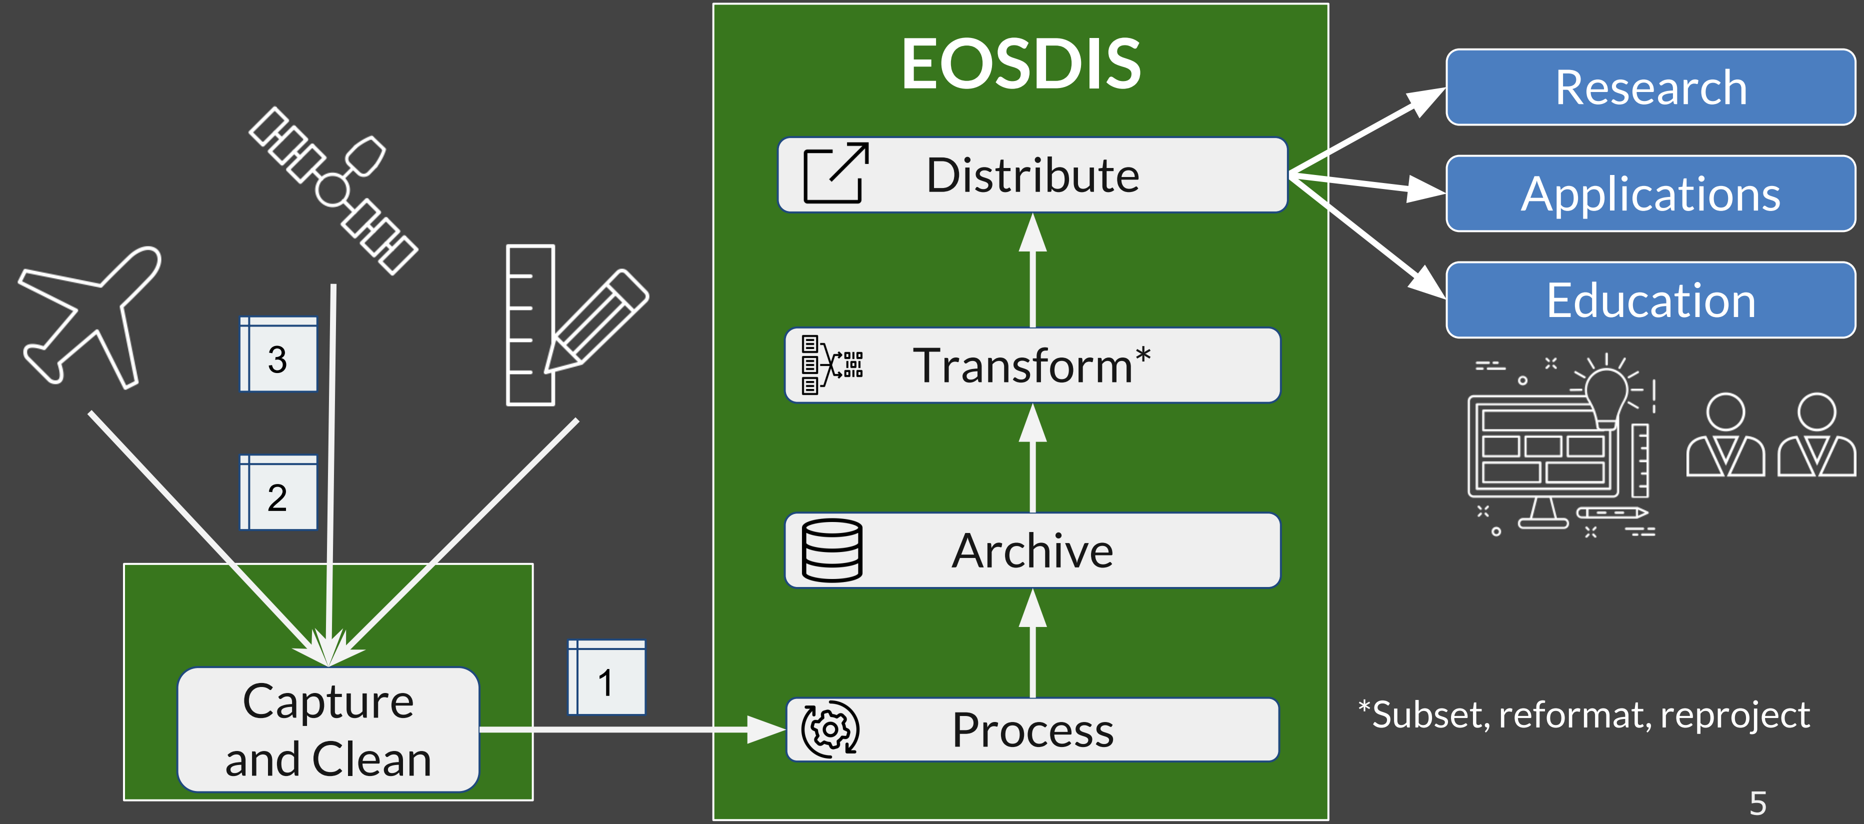

Earth Observing System Data and Information System (EOSDIS)

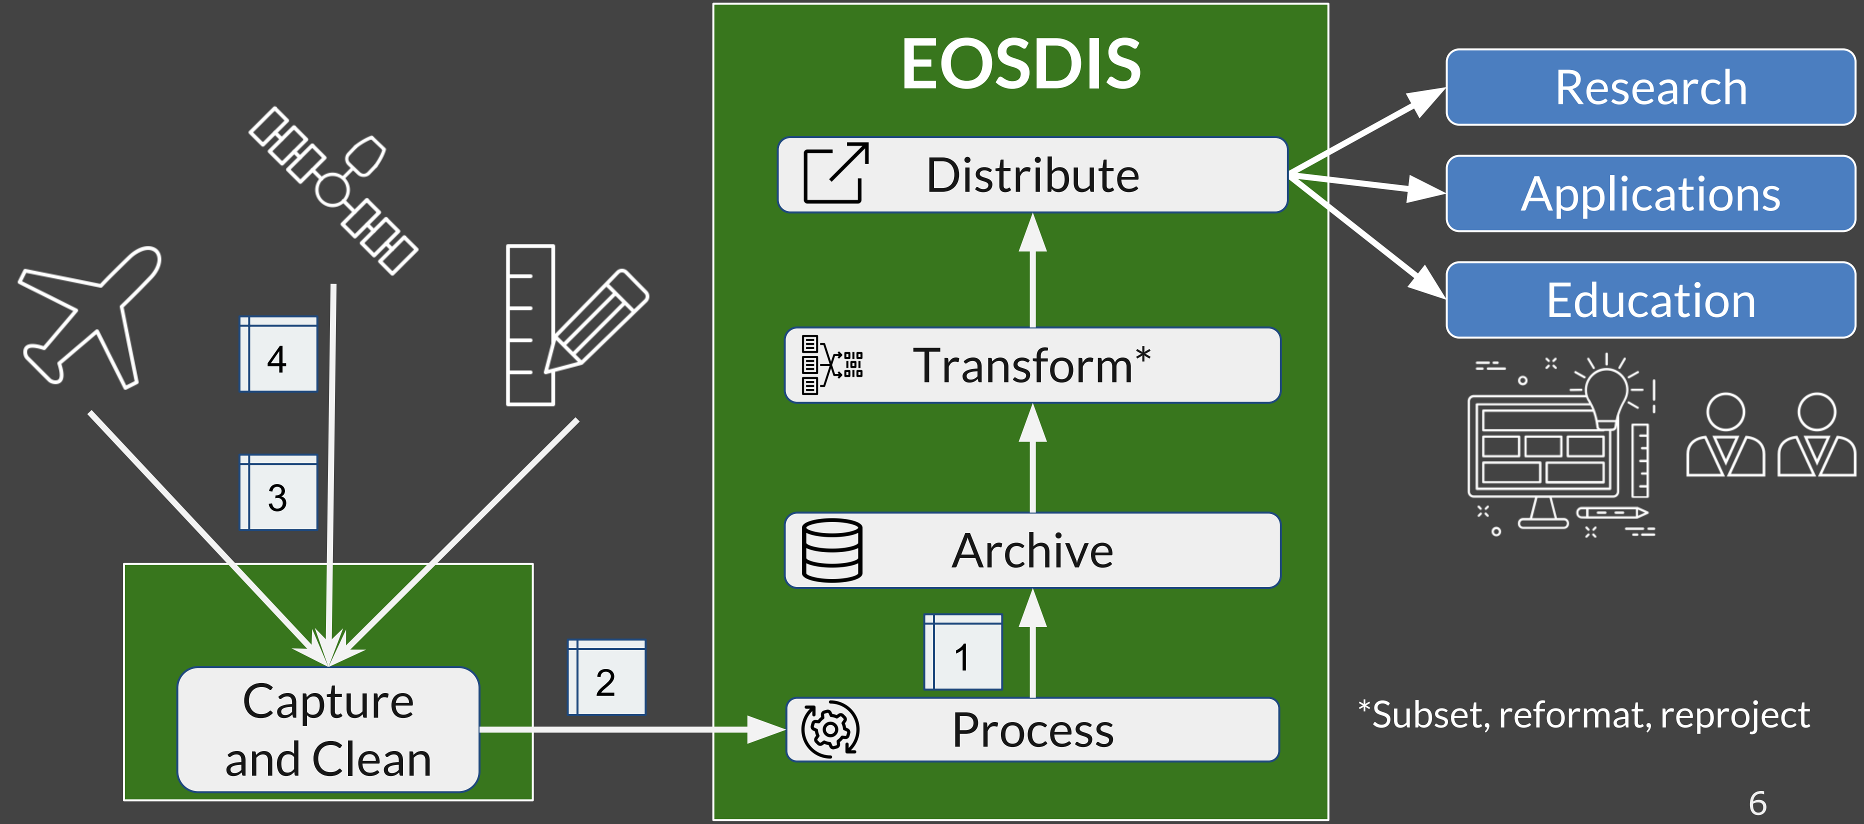

Earth Observing System Data and Information System (EOSDIS)

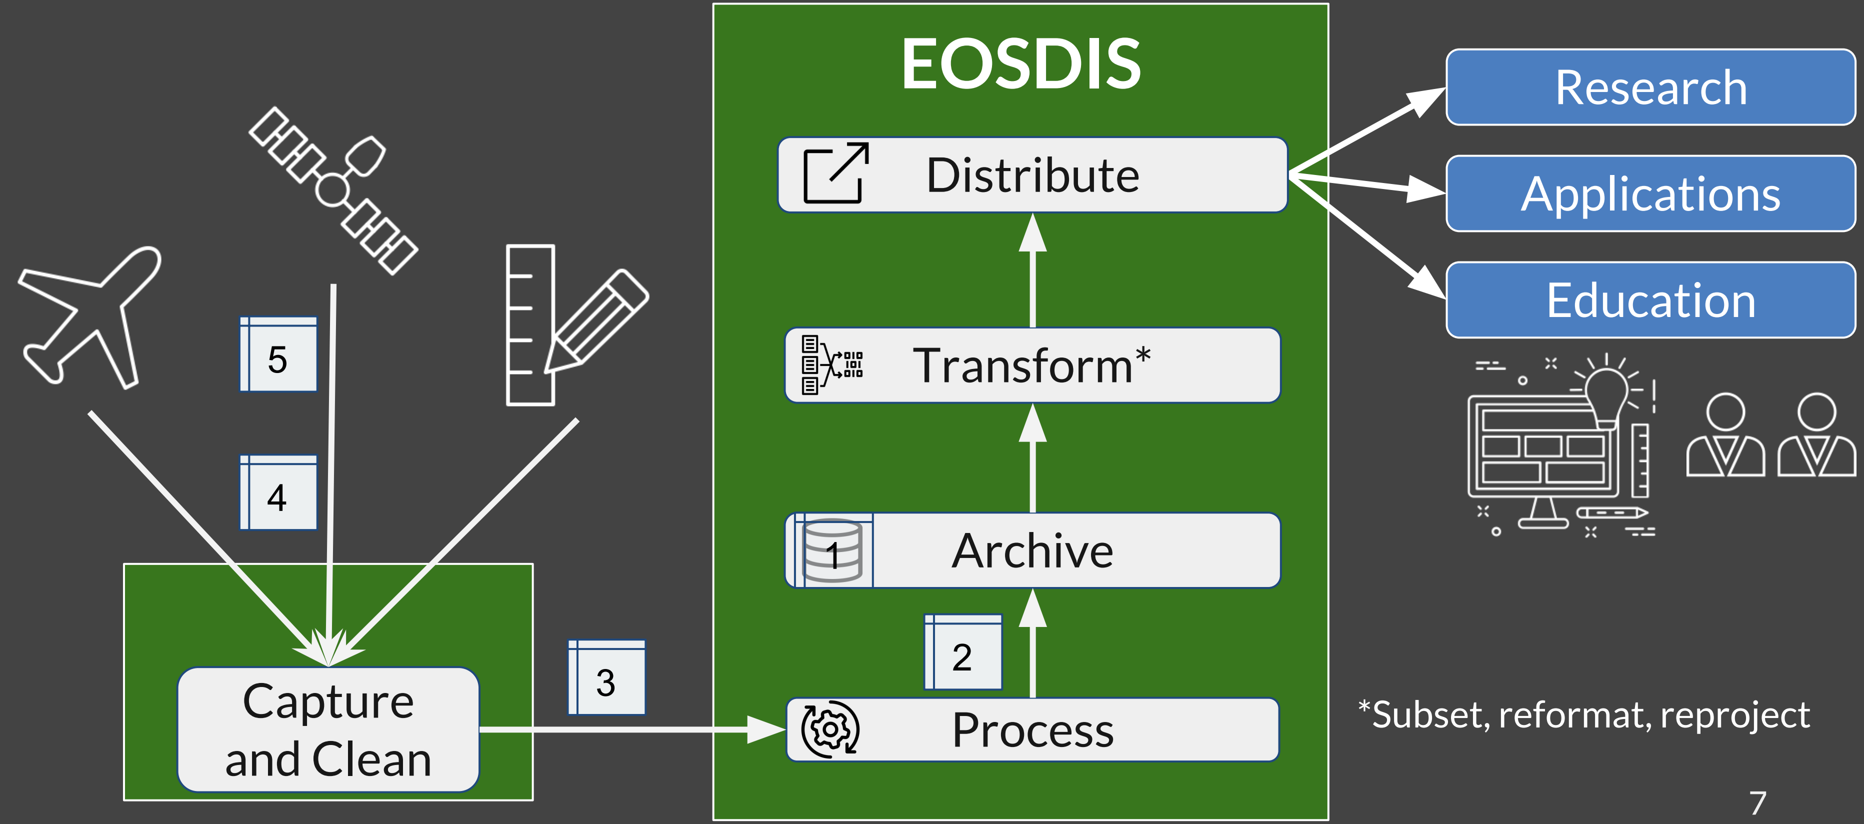

*Subset, reformat, reproject

Current Data Organization

AIRS.2017.10.29.L3.RetStd_IR001.v6.0.31.0.G17303161840.hdf	2017-10-30T20:25:44
AIRS.2017.10.30.L3.RetStd_IR001.v6.0.31.0.G17304144754.hdf	2017-10-31T18:56:04
AIRS.2017.10.31.L3.RetStd_IR001.v6.0.31.0.G17305141729.hdf	2017-11-01T18:26:14
AIRS.2017.11.01.L3.RetStd_IR001.v6.0.31.0.G17306150758.hdf	2017-11-02T19:11:25
AIRS.2017.11.02.L3.RetStd_IR001.v6.0.31.0.G17307140216.hdf	2017-11-03T18:11:46
AIRS.2017.11.03.L3.RetStd_IR001.v6.0.31.0.G17310121421.hdf	2017-11-06T17:17:47
AIRS.2017.11.04.L3.RetStd_IR001.v6.0.31.0.G17310142829.hdf	2017-11-06T19:32:48
AIRS.2017.11.05.L3.RetStd_IR001.v6.0.31.0.G17311141745.hdf	2017-11-07T19:32:58
AIRS.2017.11.06.L3.RetStd_IR001.v6.0.31.0.G17313131129.hdf	2017-11-09T18:18:20
AIRS.2017.11.07.L3.RetStd_IR001.v6.0.31.0.G17313124354.hdf	2017-11-09T17:48:19
AIRS.2017.11.08.L3.RetStd_IR001.v6.0.31.0.G17313144044.hdf	2017-11-09T19:48:20
AIRS.2017.11.09.L3.RetStd_IR001.v6.0.31.0.G17317101251.hdf	2017-11-13T15:19:09
AIRS.2017.11.10.L3.RetStd_IR001.v6.0.31.0.G17315162221.hdf	2017-11-11T21:33:51

Current Data Organization

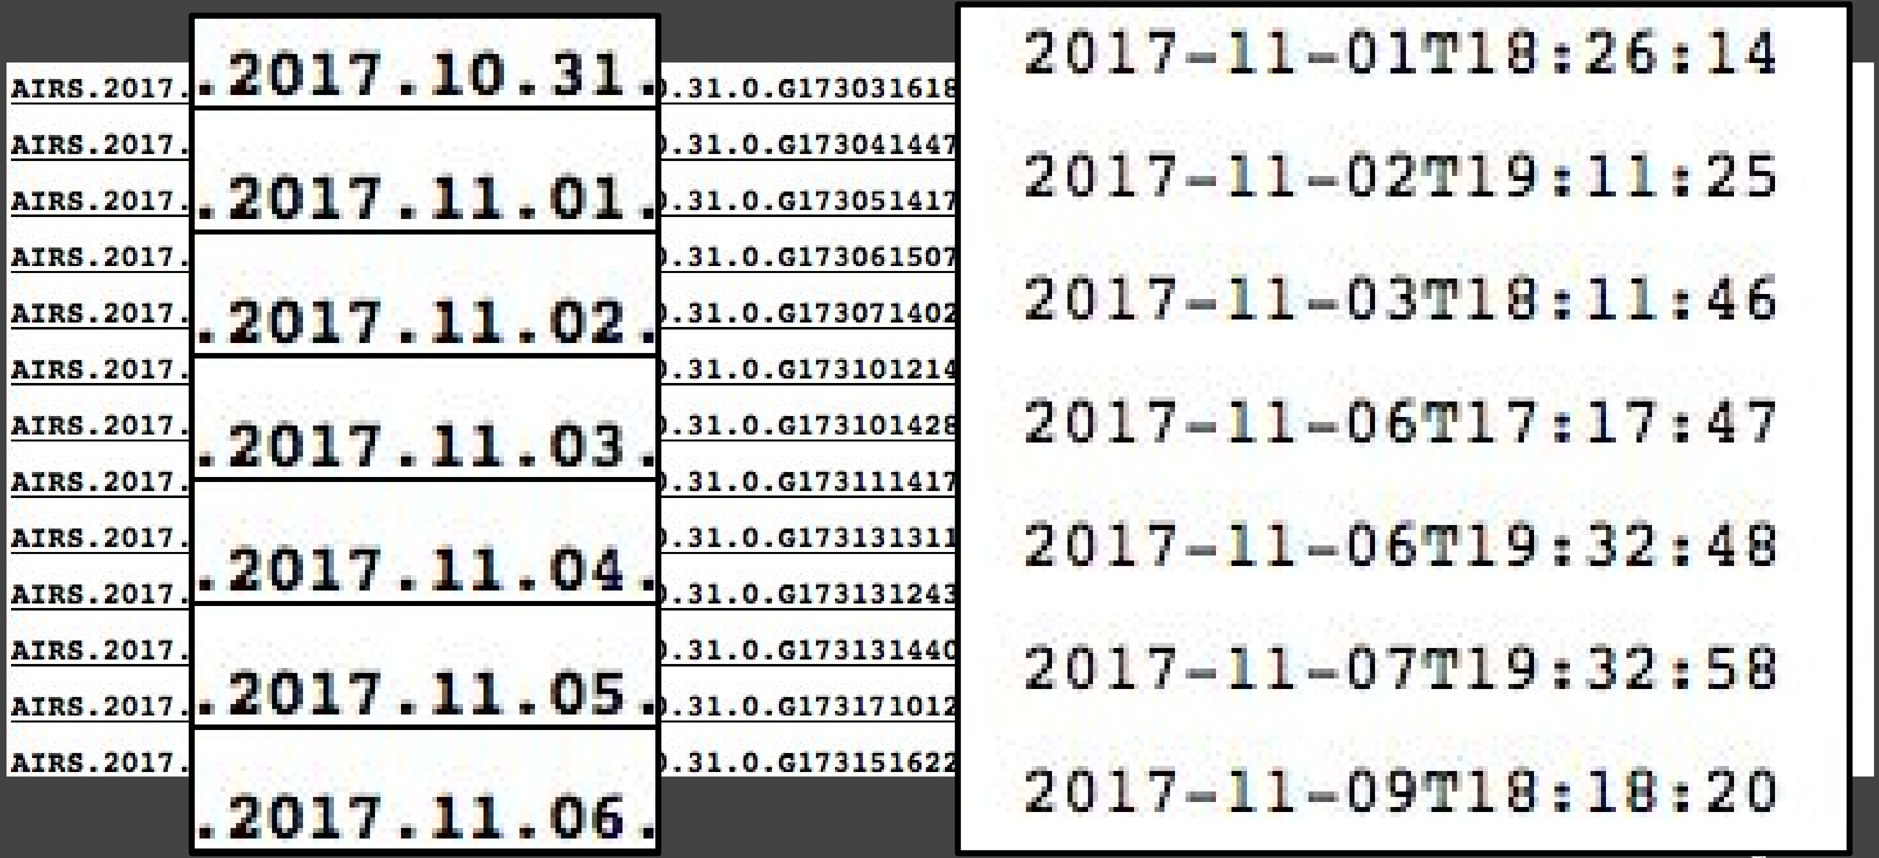

How does time slice organization affect analysis performance?

Data Set: North America Land Data Assimilation System

Temporal Resolution: Hourly

Spatial Resolution: 0.125 deg resolution (464 x 224)

Variable: Air Temperature @ 2m

Calculation: Average over time at each grid point

Hardware: MacBook Air

Software: *ncra* from netCDF Command Operators (nco)

How does time slice organization affect analysis performance?

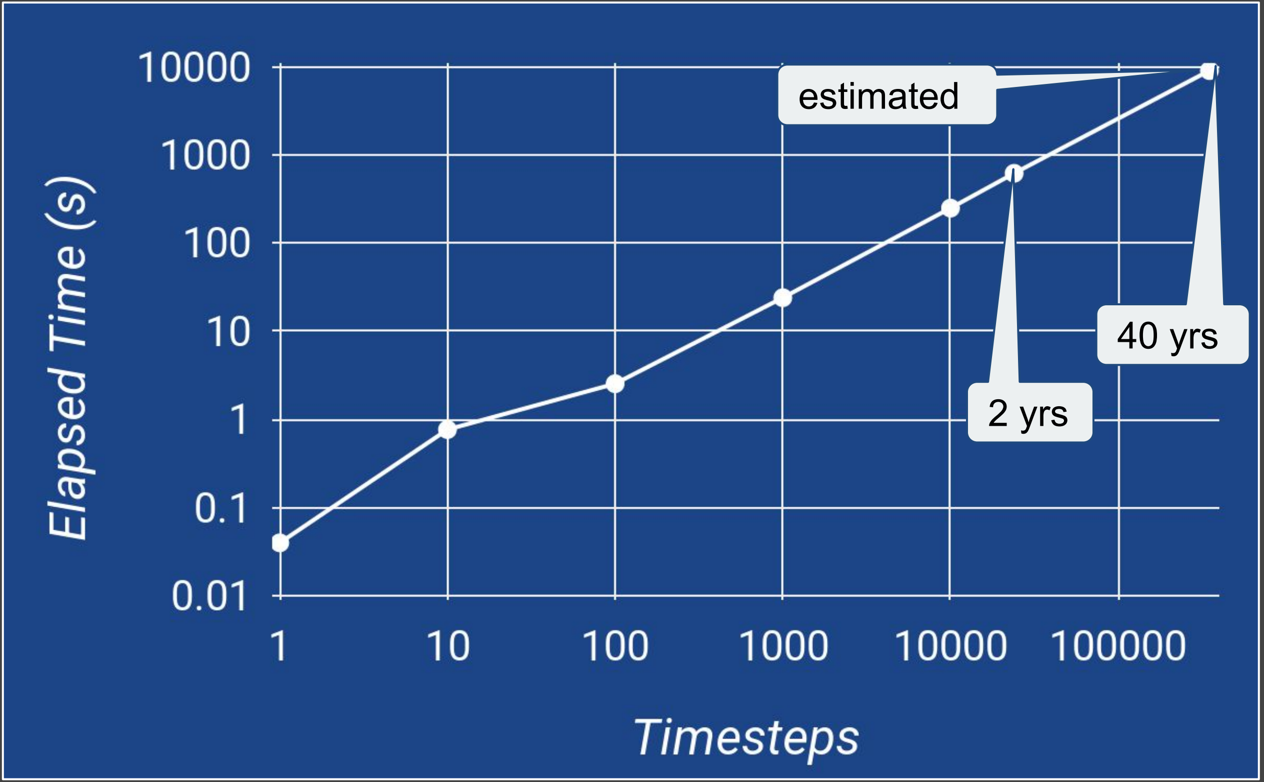

Hmmm...what if we pre-aggregate?

For 2 years of data...

Original Thin-sliced data: 17544 files

Aggregated into Yearly Files: 2 files

Data organization	Number of files	Elapsed time to process
1 Hour / File	17544	461 s
1 Year / File	2	66 s

Meanwhile, back at the ~~ranch~~ archive...

EOSDIS archive volumes are slated to grow quickly over the next several years

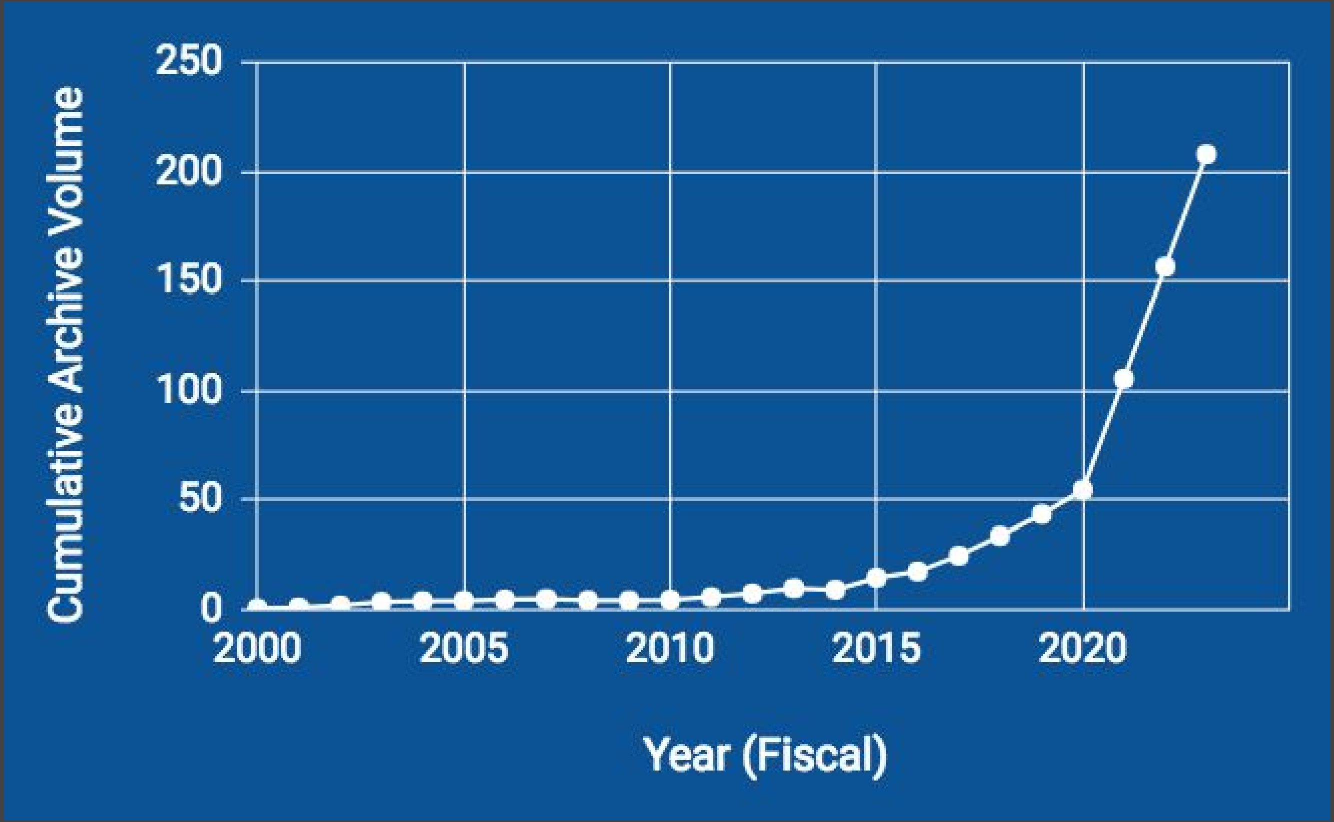

EOSDIS migration to the cloud brings several benefits

Large Volume Data Storage
All datasets stored in common Web Object Storage archive

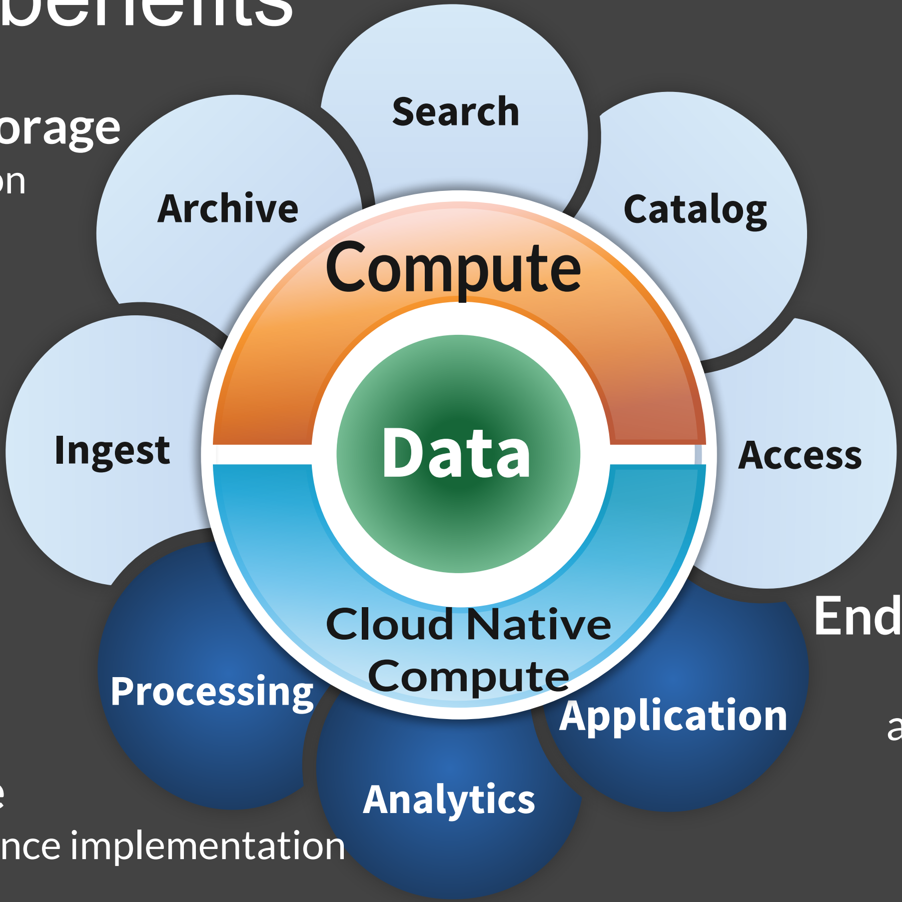

Scalable Compute
Provision based on need
Cost by use

End-User Processing
Scientists bring algorithms to the data.

Cloud Native Compute
Cloud compute services enhance implementation

“Scalable Compute” comes with a catch...

Cloud-based Data Parallelism

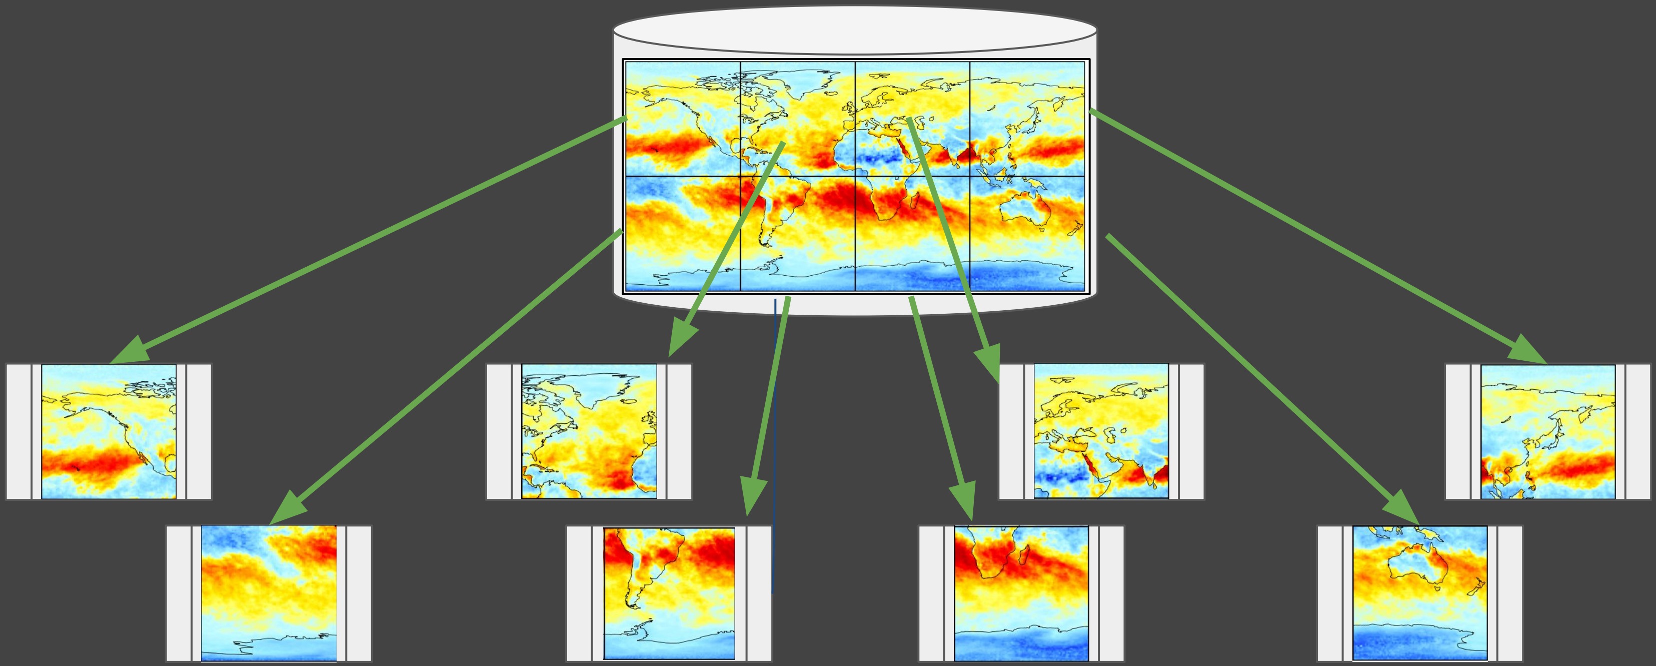

Cloud-based Data Parallelism

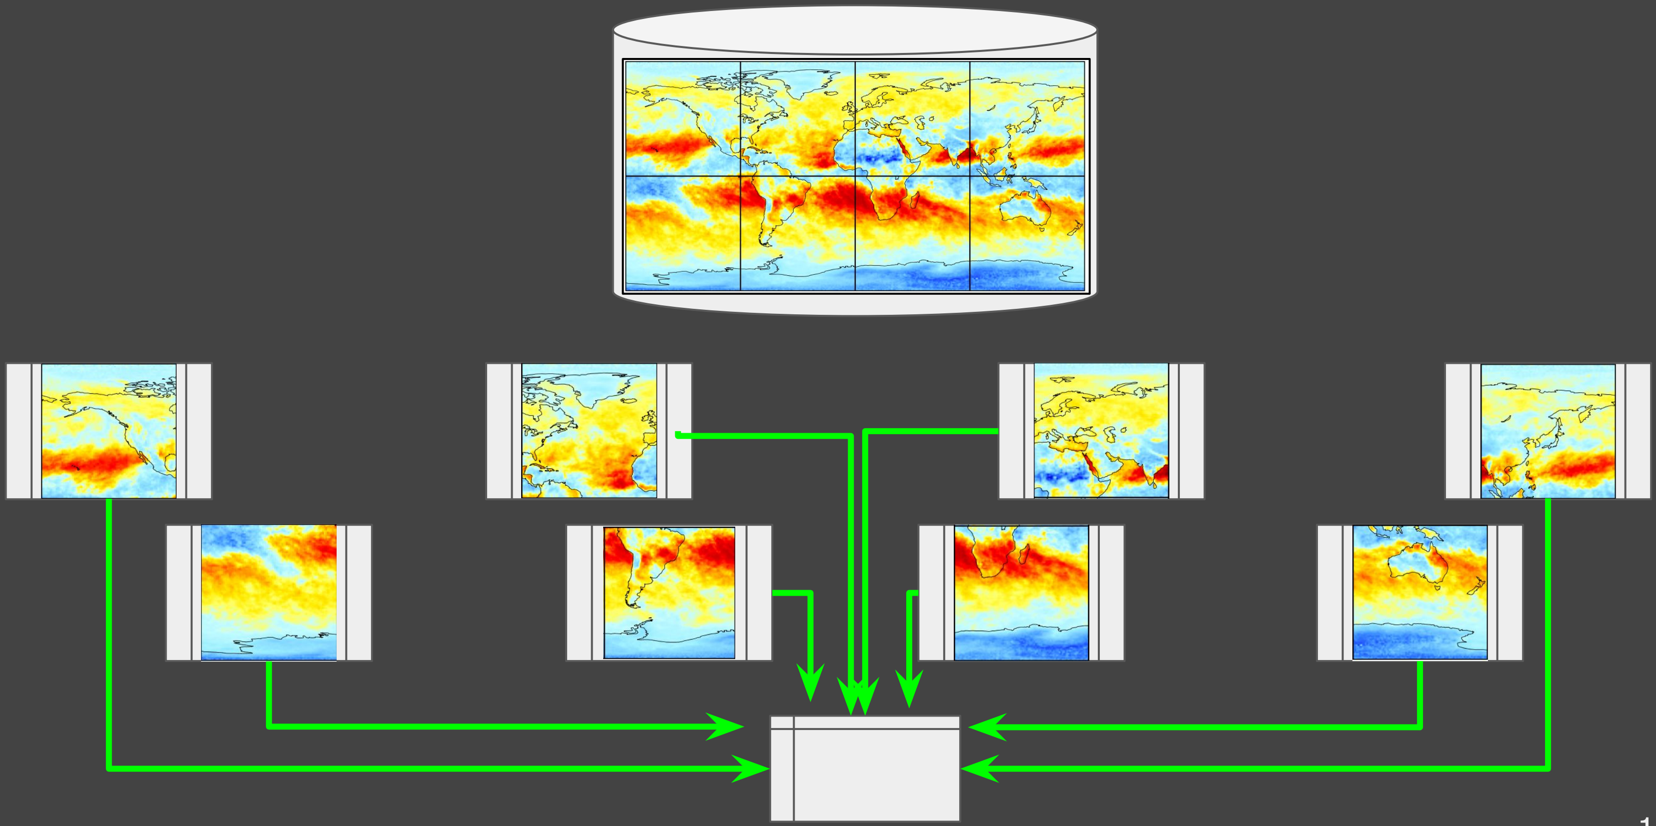

A user journey through data analysis on the cloud

Processor	Data org.	No. of files	Storage Type	Elapsed time
MacBook	1 Hr / File	17544	Local SSD	461 s
t2.xlarge	1 Hr / File	17544	Local SSD	97 s
MacBook	1 Hr / File	2	Local SSD	66 s
t2.xlarge	1 Yr / File	2	Network	56 s
t2.xlarge	1 Yr / File	2	Local SSD	39 s
t2.xlarge multi-proc.	1 Yr / File	2	Local SSD	20 s
2 * t2.xlarge multi-proc.	1 Yr / File	2	Local SSD	11 s

t2.xlarge = 4 vCPU, 8 GB memory, $0.1856/hr

SSD = Solid State Drive

Journey Cost in Time and Treasure

1.5 Days from a standing start*

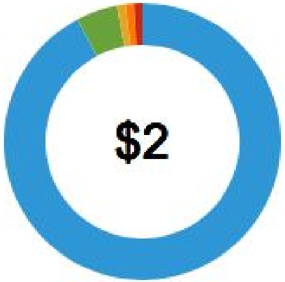

EC2	$2.10
EFS	$0.11
DataTransfer	$0.02
Budgets	$0.00
kms	$0.00
Tax	$0.00
Total	**$2.23**

*Thanks, Anaconda and nco!

Summary: How to run fast

1. Process on fast cloud CPUs
2. Reorganize the data (space-time tiles)
3. Get data onto fast storage
4. Use all the CPUs on the virtual machine
5. Use multiple virtual machines

Summary: How to run fast

1. Process on fast cloud CPUs
2. Reorganize the data (space-time tiles)
3. Get data onto fast storage
4. Use all the CPUs on the virtual machine
5. Use multiple virtual machines

Why not reorganize ALL the data in the cloud?

1. First Rule of Archive Club: Nobody modifies the original data in Archive Club.
2. But: a second copy of all the data costs a lot of money
3. Live data streams mean ever-changing tiles
4. Users may be confused by the quasi-duplication

Data Bursting

- ❖ Manual Curation
  - ➢ Burst Based on User Requests / Votes
  - ➢ Data Expeditions
- ❖ Automatic Curation
  - ➢ Event-triggered
  - ➢ “Data finds Data”*

*Jeff Jonas, http://jeffjonas.typepad.com/jeff_jonas/2009/07/data-finds-data.html

Data Bursting Opportunities

- Multi-dataset suites for studying Earth systems
- Bespoke gridding / projection schemes
- Rapid assembly of data suites in response to events

Data Bursting Challenges

- Reproducibility:
  - freeze-dry suites and store in low-temperature storage?
- Provenance:
  - bind to or place inside data?
- Choosing:
  - lightweight proposal process?
  - base on data impact?

Acronyms

CPU	Central Processing Unit
EOSDIS	Earth Observing System Data and Information System
nco	netCDF Command Operators
netCDF	Network Common Data Form
SSD	Solid State Drive