NOAA's
NATIONAL CLIMATIC DATA CENTER's

Plan for Reprocessing Large Datasets

27th IEEE MSST

May 2011

NCDC Headquarters in Asheville, NC

- NCDC fulfills much of the Nation's climate data requirements
- NCDC is the steward of the Nation's *in-situ* and satellite data and information
- Co-located with Air Force 14th Weather Squadron.
- NCDC's Federal presence extends to:
  - Asheville, NC
  - Boulder, CO
  - Honolulu, HI
  - Silver Spring, MD
  - WI, NY, TX, UT, MO, and AK

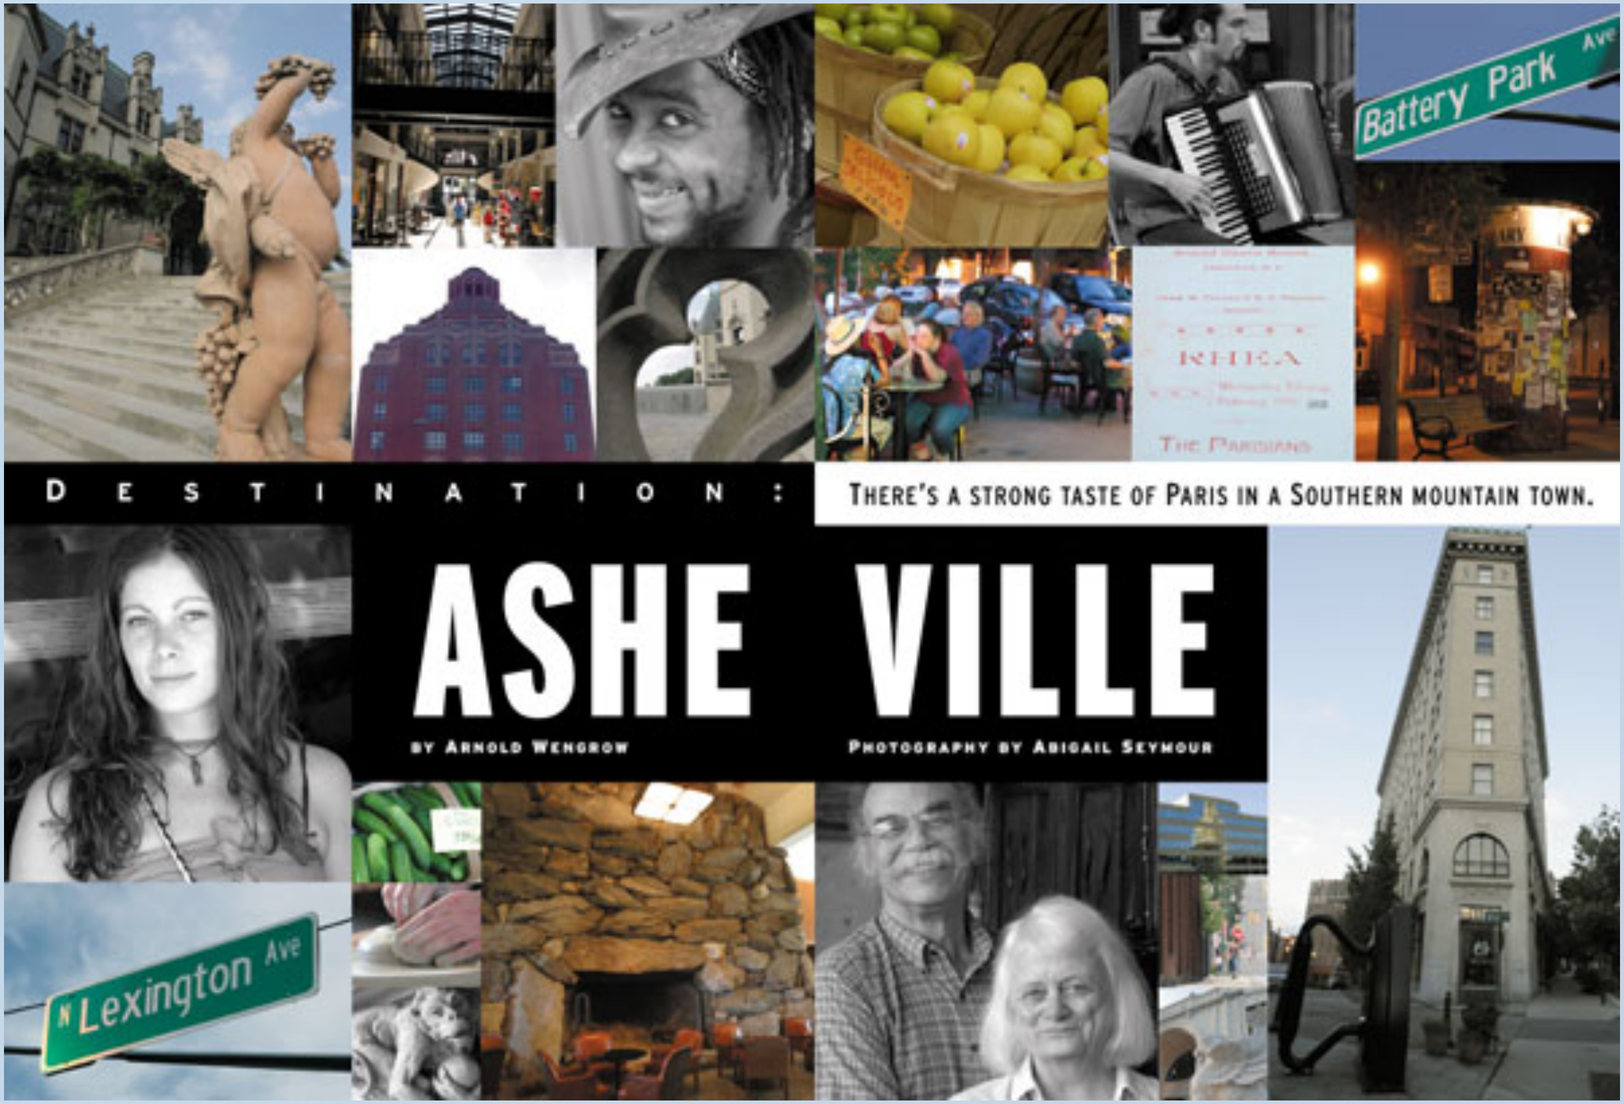
Battery Park Ave
RHEA
THE PARISIANS
DESTINATION:
THERE'S A STRONG TASTE OF PARIS IN A SOUTHERN MOUNTAIN TOWN.
ASHE VILLE
BY ARNOLD WENGROW
PHOTOGRAPHY BY ABIGAIL SEYMOUR
N Lexington Ave

Why Asheville?

- The U.S. Government took over the **Grove Arcade** during World War II because it was large and located in a safe, remote place. In 1950, the National Weather Records Center (NWRC)was established. By 1964, the U.S. **Weather Bureau** consolidated regional record centers with the NWRC. In 1982, the National Climatic Center becomes the **National Climatic Data Center** (NCDC). The Grove Arcade housed NCDC from 1951 until 1995, when it moved to its current location.

Data Received from Many Sources

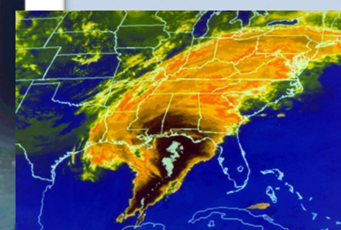

Forecast Warning Analysis

Voluntary U.S. Observers

Global Weather Reports

NCEP Weather Charts & Models

Ship, Buoy Reports

Rocketsonde

Weather Balloons

Storm Data

Doppler Radar

(GOES, POES, NPOESS, many other) Satellites

Aircraft Observations

Wind Profiler

Airport Weather Reports (ASOS)

U.S. Climate Reference Network

Reanalysis & Climate Models

Archive, Access and Assessment

Recent Highlights

Safe Storage of over 5.35 Petabytes of climate data

- Two copies: Primary and Offsite

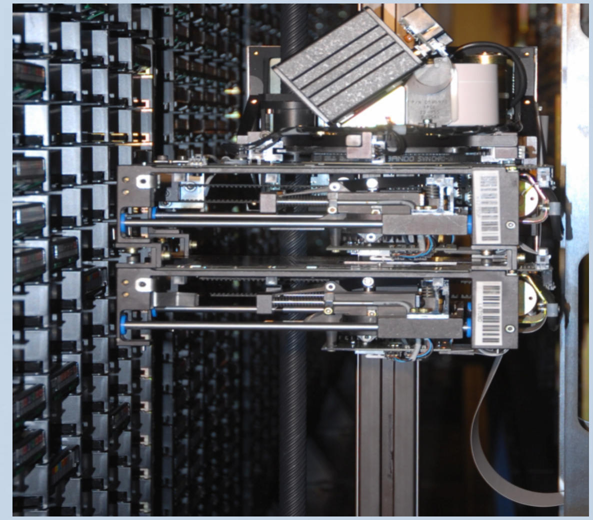

2007 Nobel Peace Prize Scientists

- NCDC scientists contributed to the IPCC report (Intergovernmental Panel on Climate Change), which shared the Nobel Peace Prize with Al Gore

Other Key Scientific Assessments

- Climate Change Science Program (CCSP) Assessments
- Weekly / Monthly / Annual "State of the Climate" Reports
- National Assessments

Mandated Functions of NCDC's Mission

Archive and Provide Scientific Stewardship of the Nation's Meteorological Data – National and International

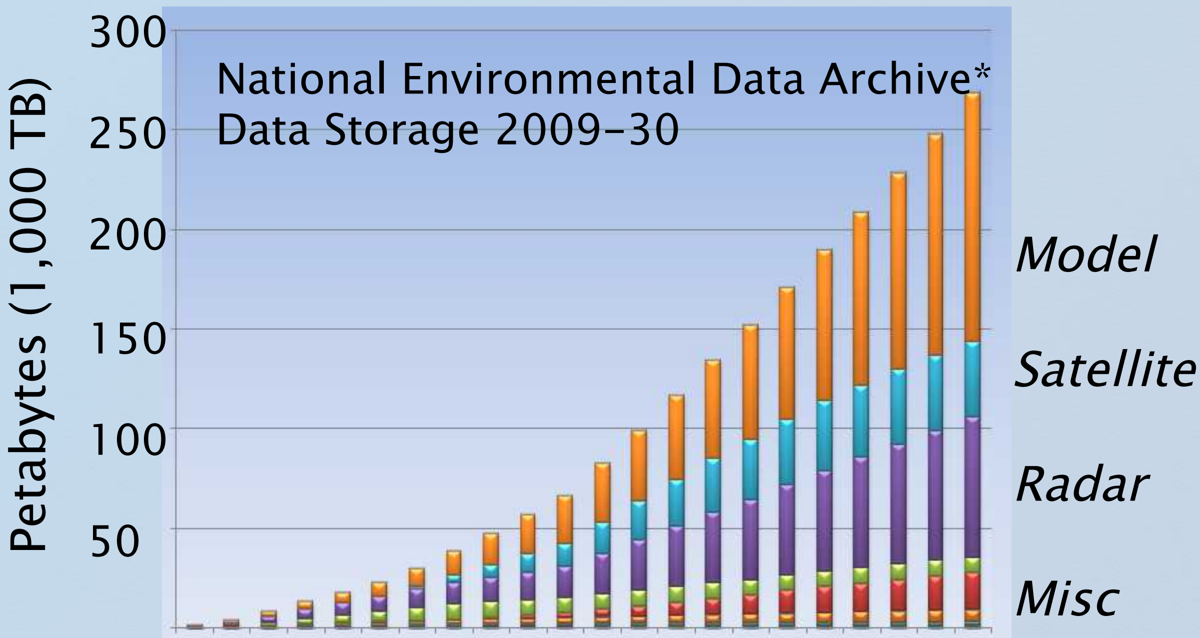

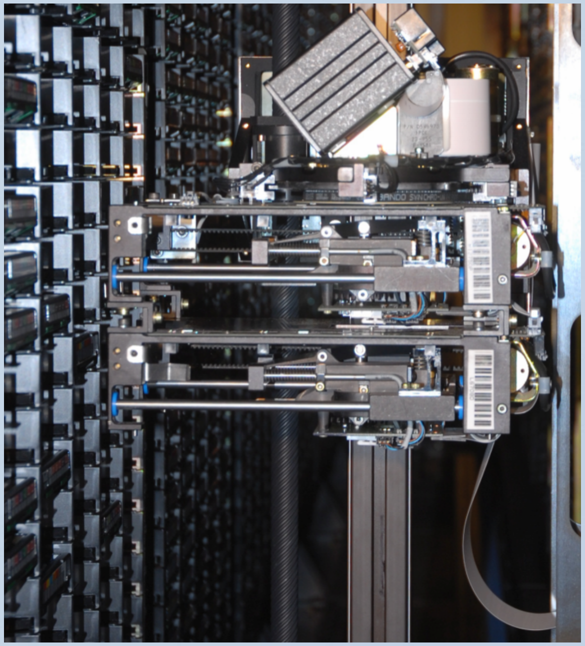

*Total cumulative date volume projected for storage by the NOAA Data Center (single copy, backup not included)

Stewardship vs Computing

Preservation and Storage are not the problems

February 16, 2011
"Tape is Big Data: 80% of the world's data is stored on tape and tape is the only media that can scale to exabyte(s) and still be cost effective."

Matt Starr
Spectra Logic CTO

Data Intensive Scientific Computing is

Storage bandwidth hasn't kept pace with capacity
Network bandwidth hasn't kept pace with data volumes

Bring computations to the data
- "Eliminate gratuitous data movement"

Go from "working to working"
- Build & test early & often

Photo by flickr.com/dullhunk via Creative Commons Attribution 2.0 Generic (CC BY 2.0)

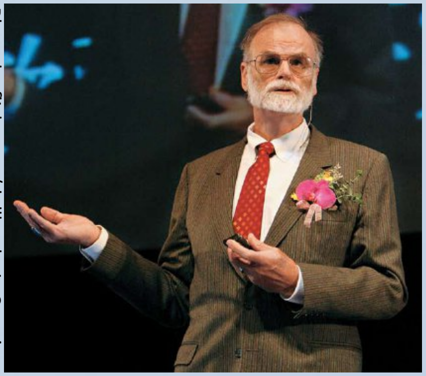

Jim Gray
cf."The Fourth Paradigm: Data-Intensive Scientific Discovery"

Remote Sensing and Applications Division

- Provides scientific leadership in the use of NCDC's satellite and radar data sets and their applications, particularly uses in numerical weather and climate prediction." (2005)
- A Climate Data Record (CDR) is a time series of measurements of sufficient length, consistency, and continuity to determine climate variability and change

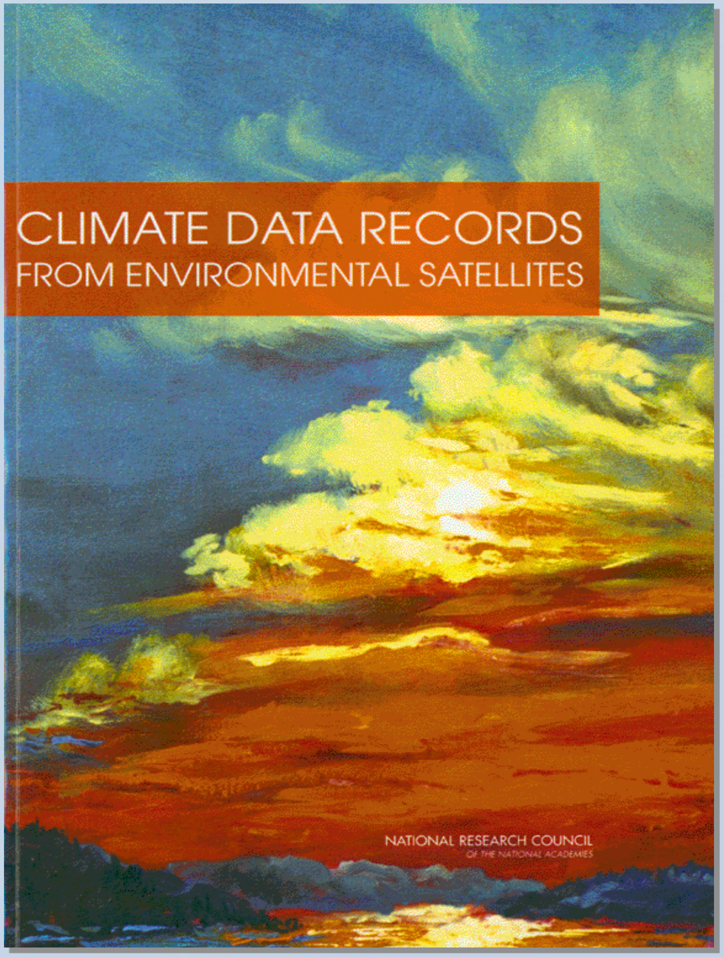

CDRs Are A Core Business at NCDC

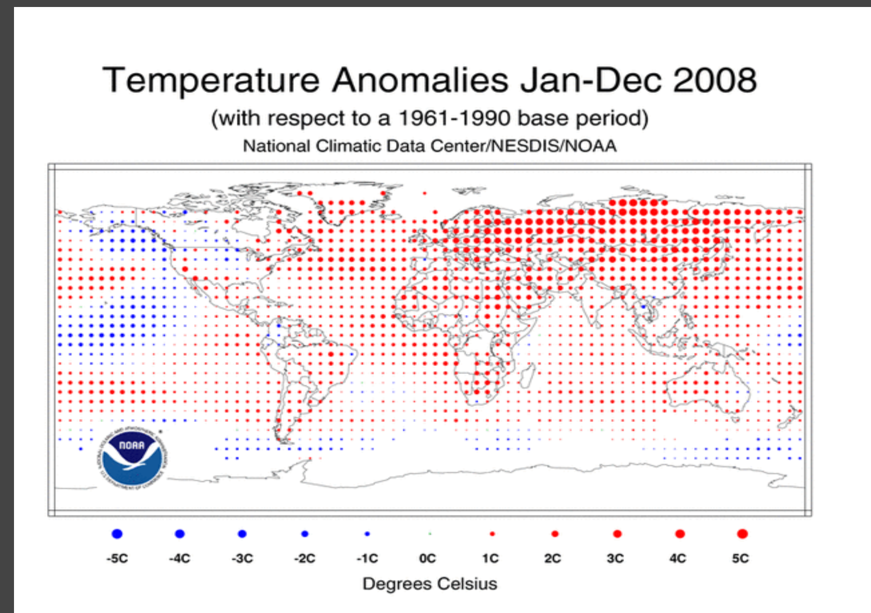

Mission:

To provide access and <u>stewardship to </u>the Nation's resource of <u>global climate </u>and weather related <u>data </u>and information, and <u>assess and monitor climate variation and change</u>

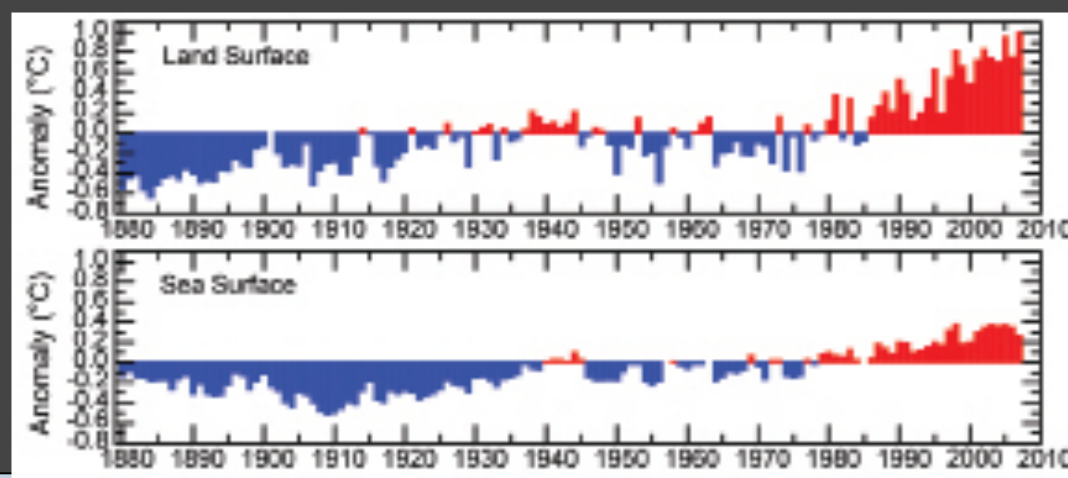

NOAA (and others) has Collected Decades of Satellite Data

Using Same or Similar Observing Systems

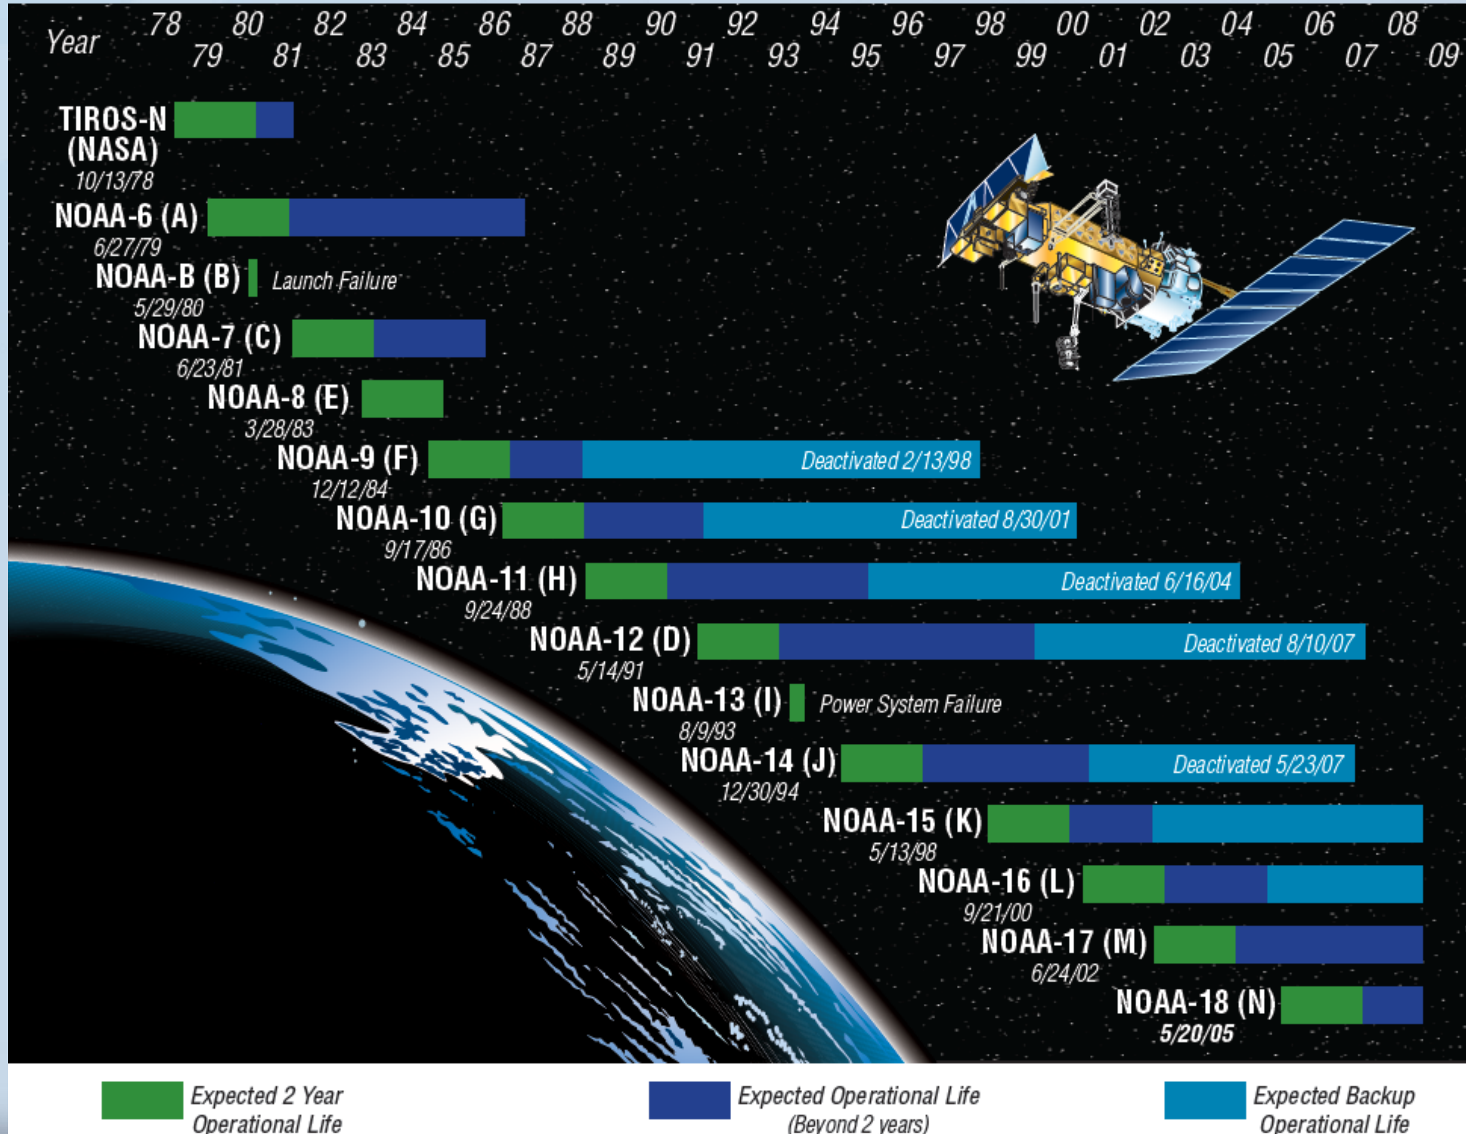

CDR Process

Input

1970	1980	1990	2000	2010	2020	2030

POES/GOES/DMSP/EOS (others as appropriate) | NPP | JPSS/GOES-R/JASON-x/DWSS (others as appropriate)

Competitively-selected Community Experts

Output

Performance Metric	2010	2011	2012	2013	2014	2015
CDRs in Operations	3	10	14	18	18	18

Reprocessing Data From Different Sensors
Expert Knowledge and Retrospective Insights

Uncorrected Data Time Series Contain Both Environmental Information and Satellite-induced Artifacts

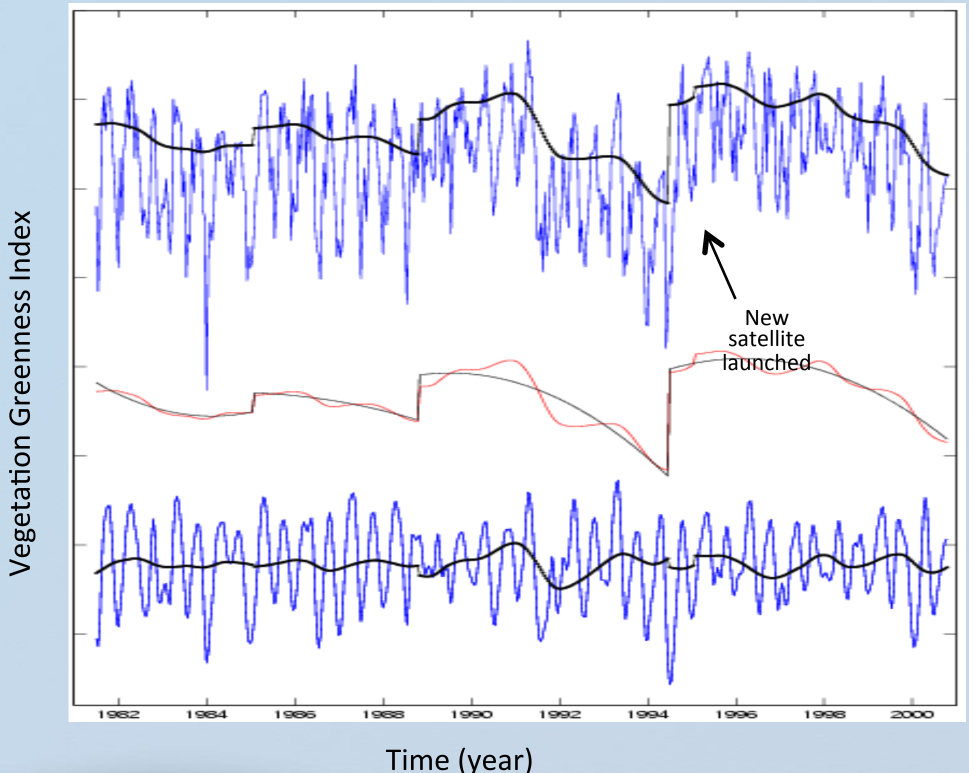

<u>Operational weather and hazard</u> products are produced rapidly to potentially save life and property

<u>Climate Data Records (CDRs)</u>
provide long term product consistency through rigorous reprocessing with advanced algorithms, ancillary data and evolved instrument understanding.

<u>Climate Information Records (CIRs)</u>
provide specific information about environmental phenomena of particular importance to science and society (*e.g.*, hurricane trends, drought patterns)

Reprocessing Results

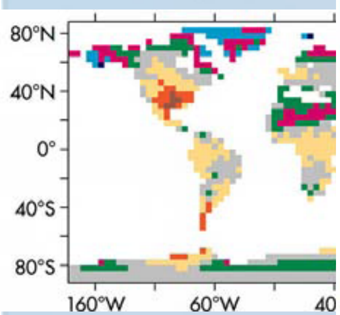

~1995

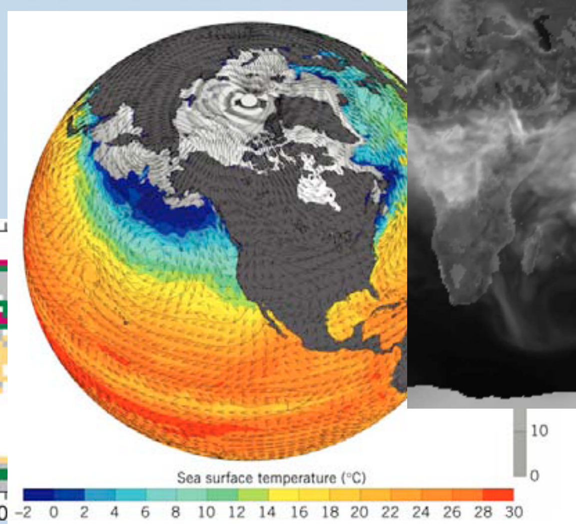

~2000

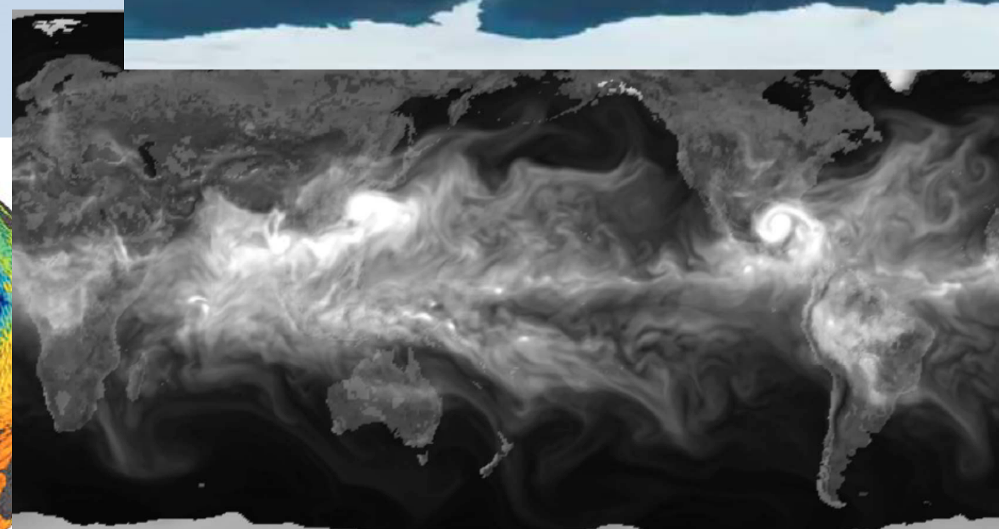

~2005

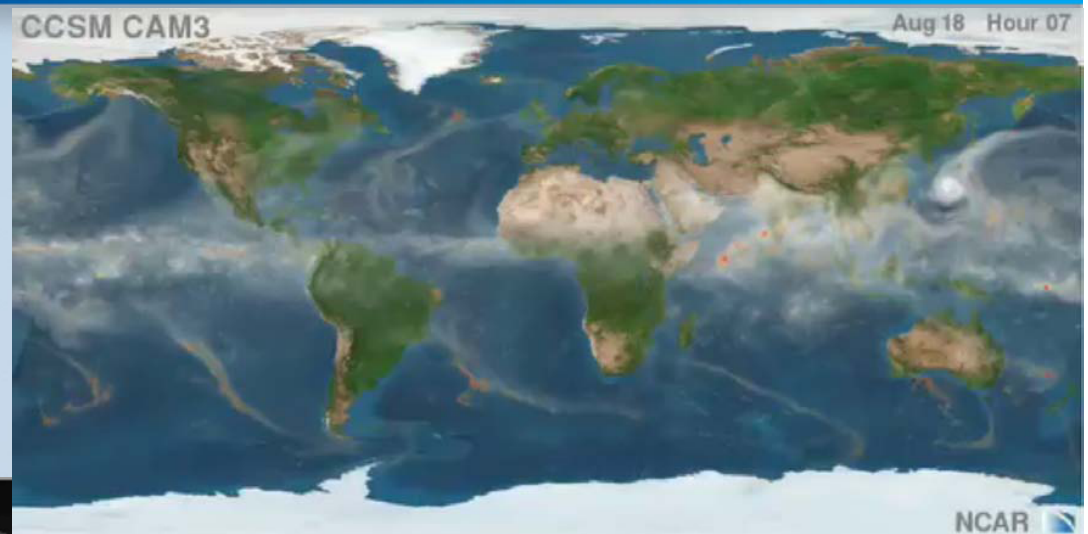

~2009

Can be dramatic

Reprocessing Reveals Key Information From Past Data

29 Operational Geostationary Satellites + 16 Operational Polar Satellites

Raw Satellite Data

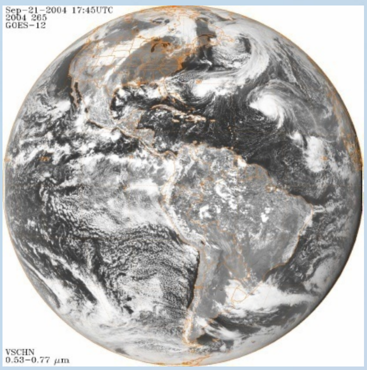

Geostationary Imagers
- Infrared window
- Visible
- Infrared water vapor

Format? Navigation?
Remapping?
Data volume?

30 Year FCDRs

HURSAT

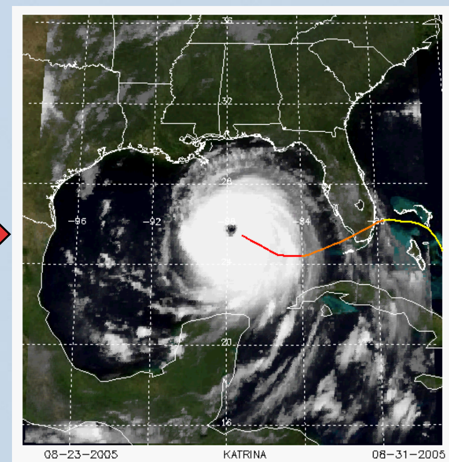

HURSAT-B1
- Brightness Temperatures
- Geo-located
- Calibrated
- netCDF
- 8km, 3-hourly

Global and Basin Objective Hurricane Intensity Trends

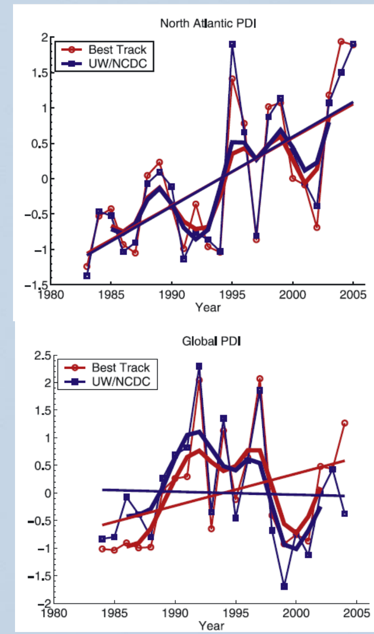

- New intensity estimates (e.g., Kossin et al. 2007)

Data Movement Challenges

Discover endpoints, determine available protocols, negotiate firewalls, configure software, manage space, determine required credentials, configure protocols, detect and respond to failures, determine expected performance, determine actual performance, identify diagnose and correct network misconfigurations, integrate with file systems, ...

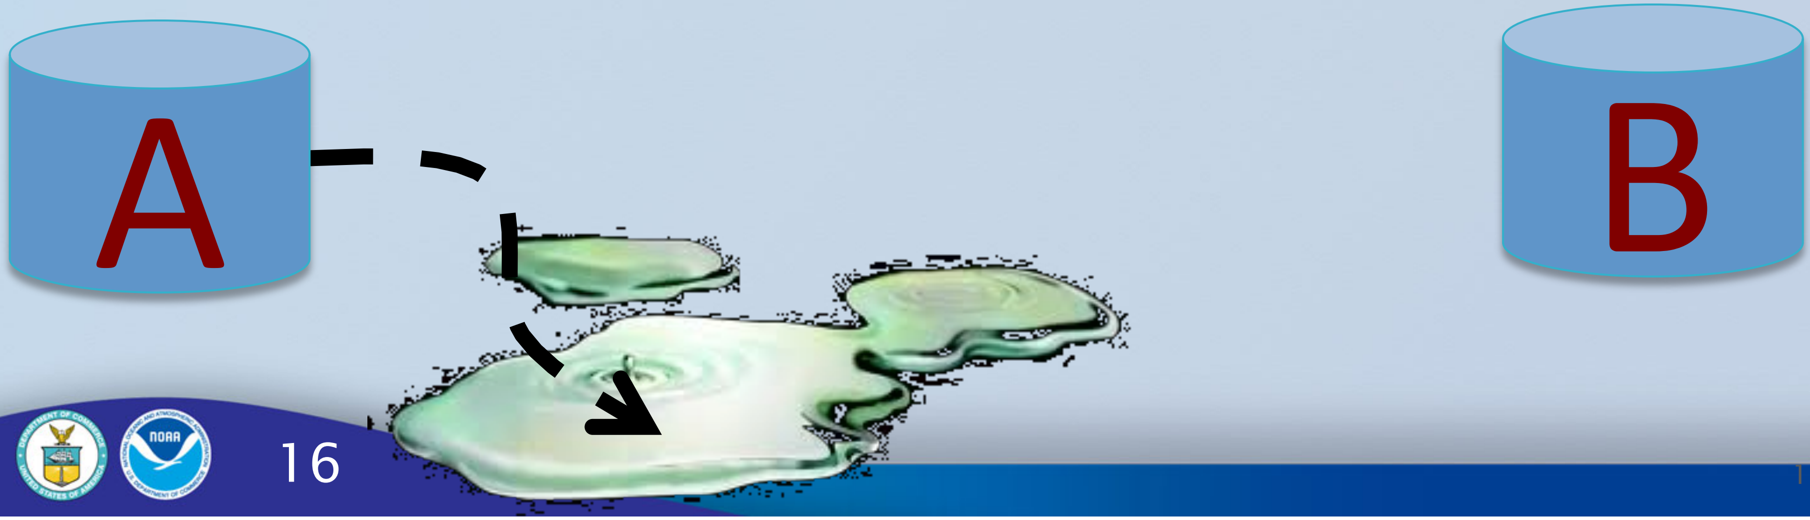

Transfer Rates

2002: Slow

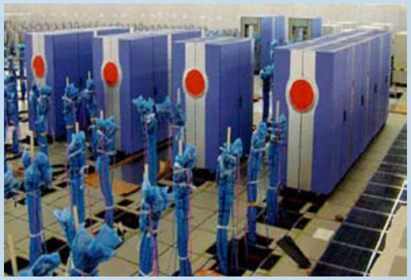

NEC Earth Simulator
no optics

2005: 2Gb/s

55 miles Optical Cables

2008: 5Gb/s

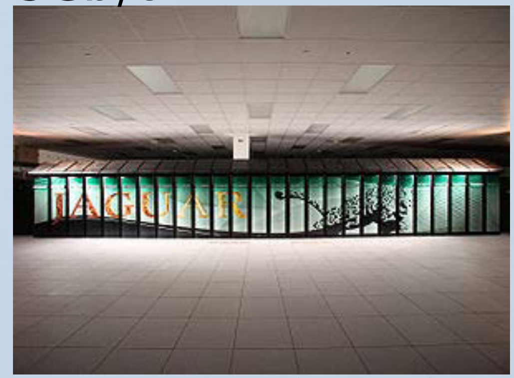

Cray Jaguar(ORNL)
3 miles Optical Cables

IBM Federation Switch for ASCI Purple (LLNL)
-Copper for short-distance (≤10 m)
-Optical for longer links (20-40m)

Reprocessing Problem

Modern Satellite Raw Data Records (RDRs) - 600GB/day

- Climate-Raw Data Record (C-RDRs) - 600GB/day
  - CDR - Best guess of 300GB/day

1.5TB/day total daily volume

545TB/year

5.45Petabytes/decade

GOAL: Reprocess one year in a day

Reprocessing Problem

GOAL: Reprocess one year in a day

545TB per year:

- 1 Gbit Ethernet (~125MB/sec) - ~1,200hrs (~50.5 days)
- 8G Fiber Channel (~850MB/sec) - ~178hrs (~7.5 days)
- 10 Gbit Ethernet (~1.2GB/sec) - ~126hrs (~5.3 days)
- 12X QDR Infiniband (~11GB/sec) - ~13hrs

Moving the data is the problem

- storage/network bandwidth

Reprocessing Problem: Data Solution

Strategically Staging the data

- Disk storage is getting cheaper
- Data is protected in robust Archive

Federating Data (iRODS)

- High-performance network transfer
- Support for a wide range of physical storage
- Easy replication/synchronization of large collections

NCDC RDR/C-RDR Archive and Distribution

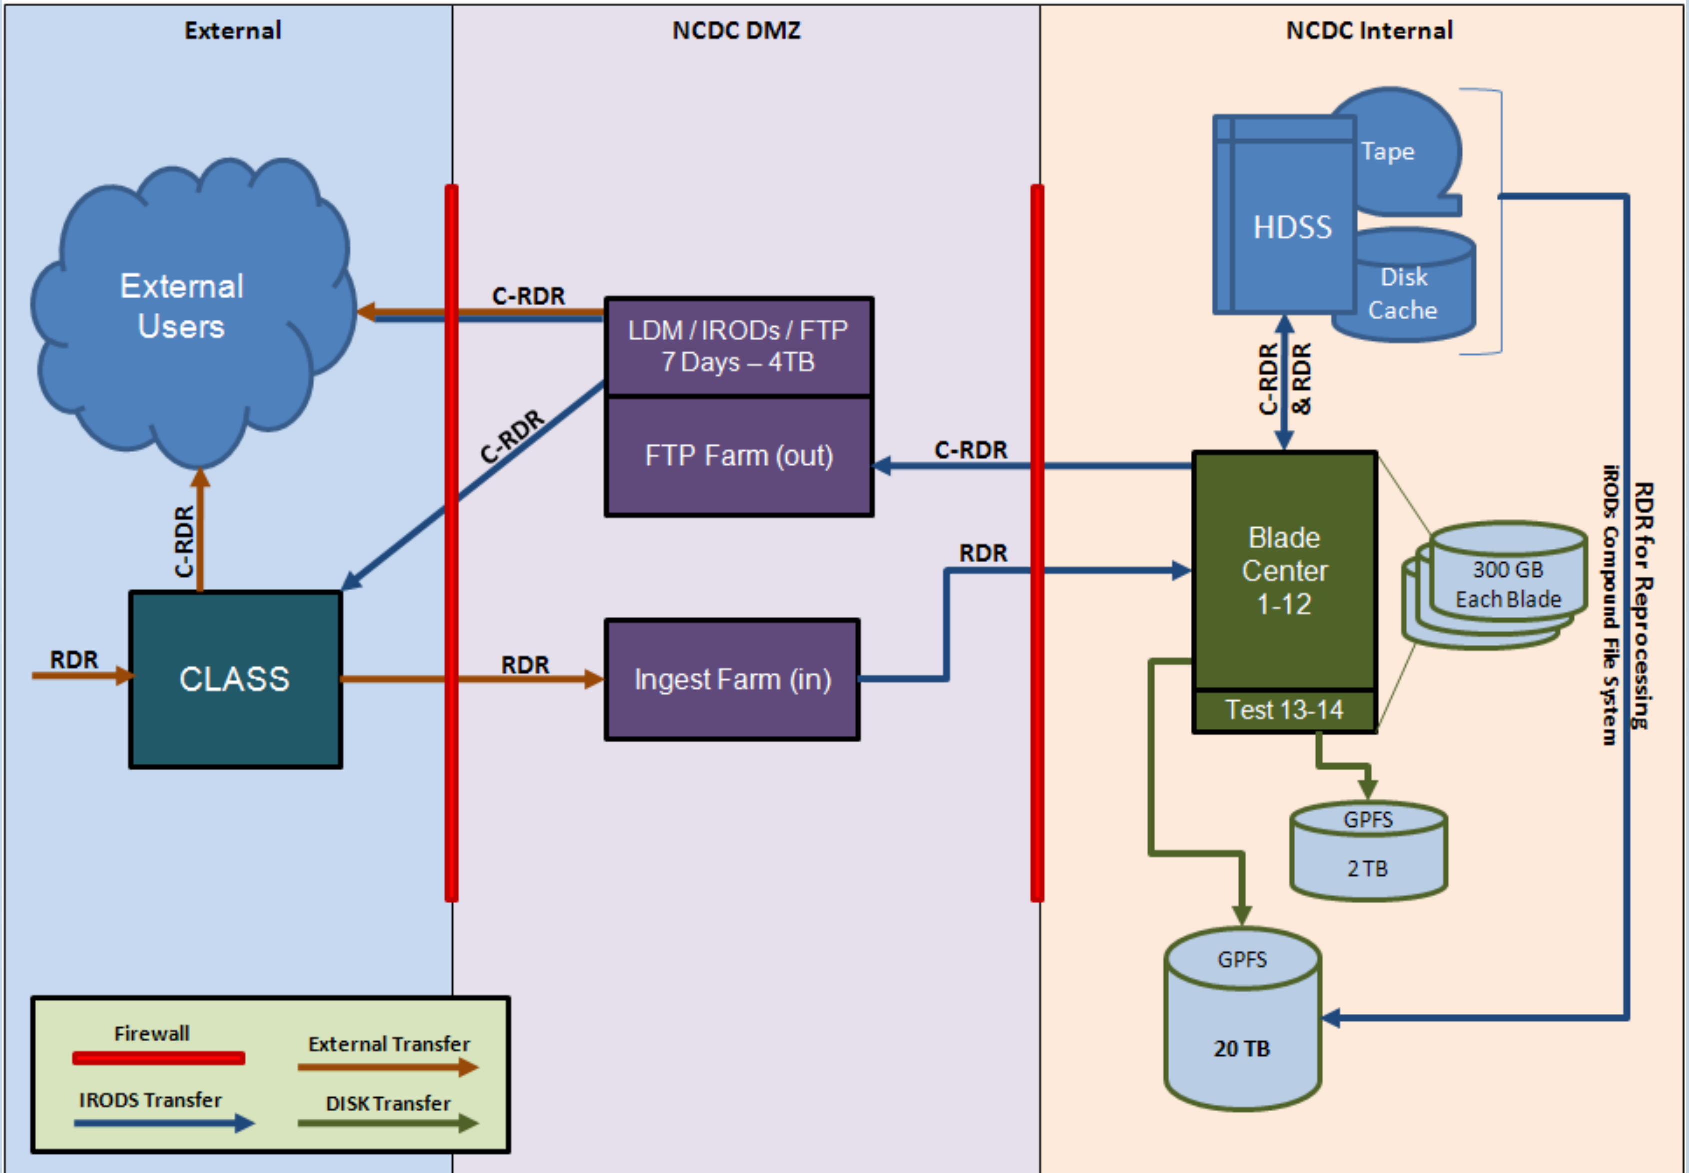

Reprocessing Problem: Computation Solution 1

Once the data is staged:

- Computations are moved to the data

Packaging of processes and dependencies

- Code; Ancillary data; Libraries; etc

Package is moved to data platform

- Compiled and installed

Scheduled to run

CDR Research to Production Roadmap

CDR Development	Operational Readiness	Production Package	Reprocessing Package
Maturity Matrix	Deployed to Integration System	Source Code Archived with Data	Strategically Staged Data
Code in Repository	Transformed to run in cluster	Deployed to Production cluster	Package Deployed for Reprocessing
Original Code Versioned	New production version tagged	Self Contained Deployment Package Created	New version of product is returned

Advantages/Disadvantages: Computation Solution 1

Advantages:

- Dedicated known platform
- Can duplicate platform locally for development/testing
- Complete control over entire compute environment
- Fixed cost

Disadvantages:

- High fixed cost
- Not scalable/flexible

Reprocessing Problem: Computation Solution 2

The Cloud?

Photo by flickr.com/null0 via Creative Commons Attribution – Share Alike 2.0 Generic (CC BY-SA 2.0)

John McCarthy, 1961:

If computers of the kind I have advocated become the computers of the future, then **computing may someday be organized as a public utility** just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry.

Reprocessing Problem: Computation Solution 2

Image instead of a package?

Virtual Machine **Image**

- Operating system + any installed server or application software
- Complete machine state, bundled & stored (~ "hibernate file")

Virtual Machine **Instance Type**

- "Resource bundles" of a particular size
- May be abstracted from hardware – or not
  - Bandwidth to storage; bandwidth between cores
  - Exclusive use of CPU, RAM, *etc.*

Virtual Machine **Instance**

- A Machine Image running on a particular Instance Type
- A running OS w/ RAM, filesystem, IP address, services, etc.

Elastic IT provisioning

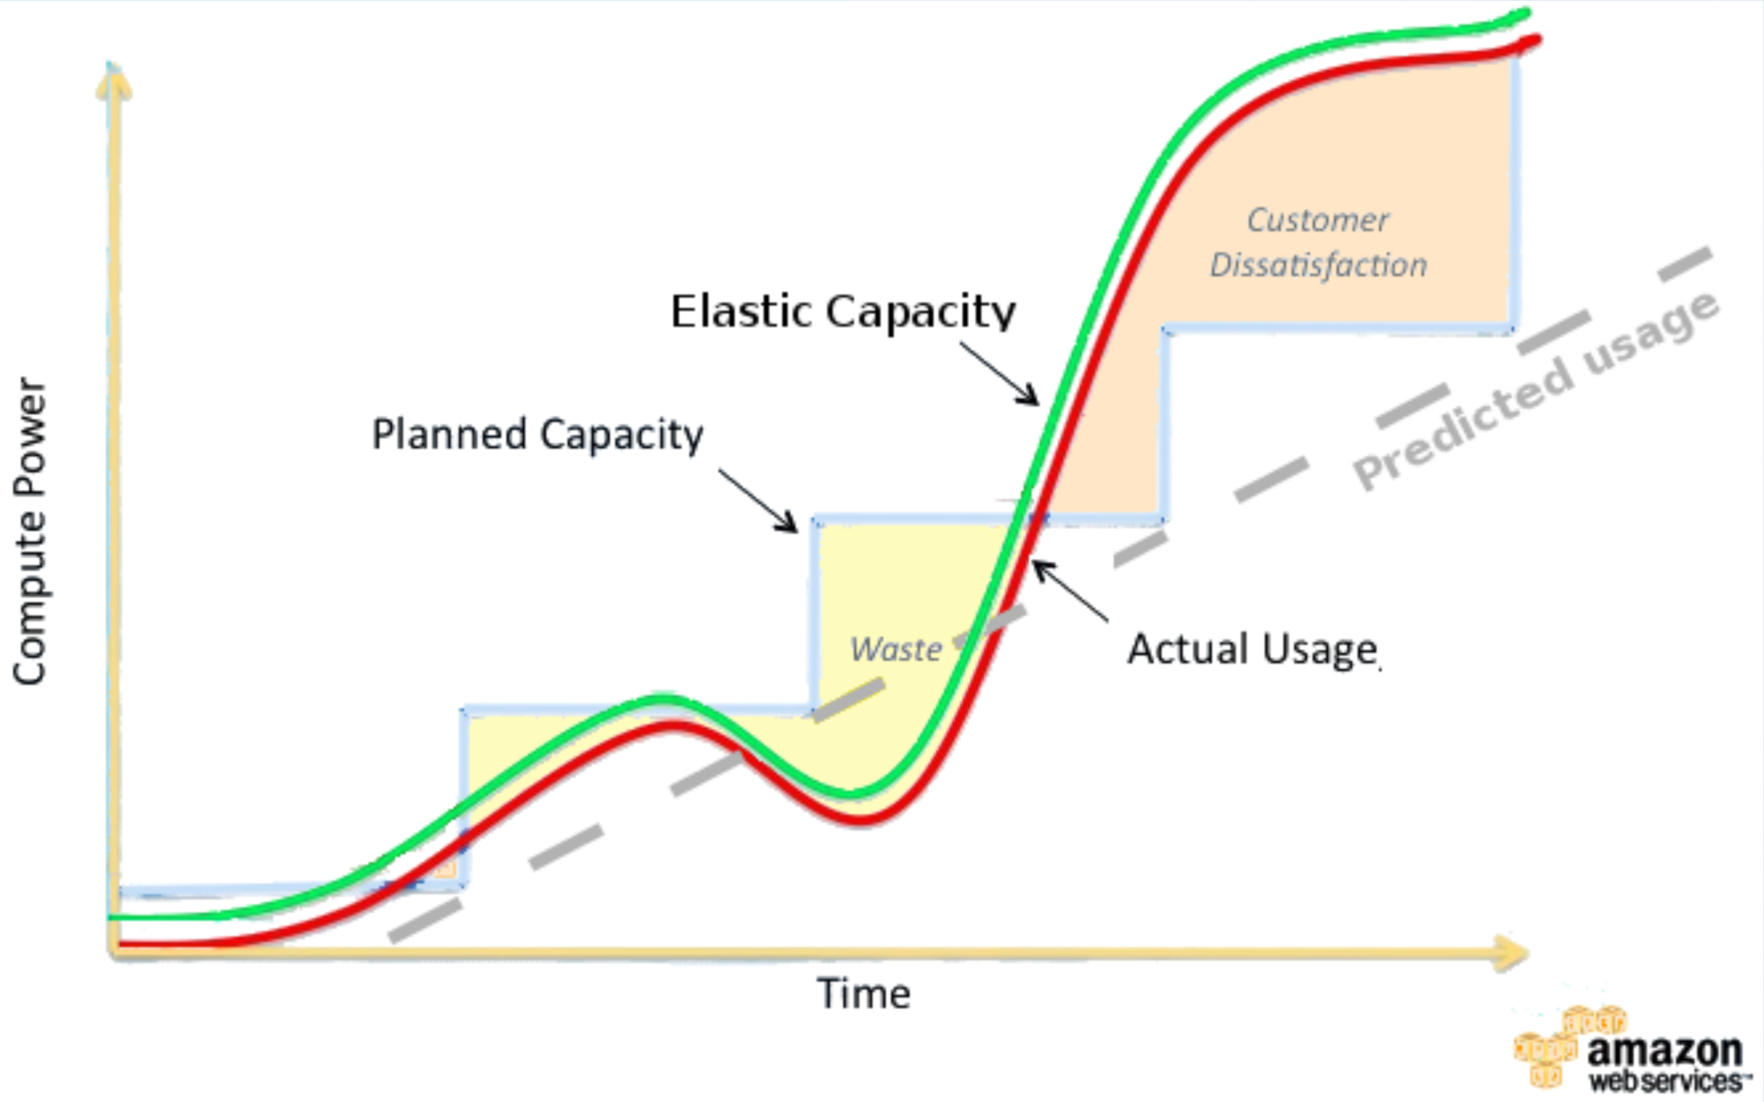

Workloads optimal for cloud

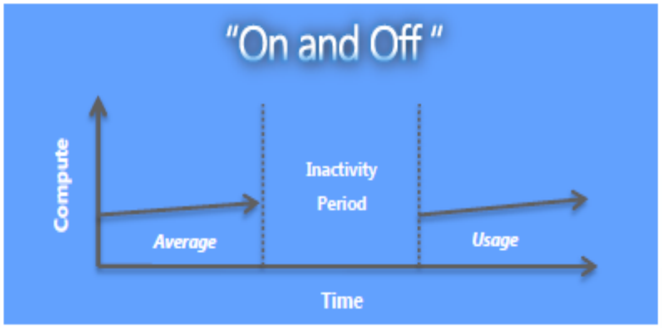

- On & off workloads (e.g. batch job)
- Over provisioned capacity is wasted
- Time to market can be cumbersome

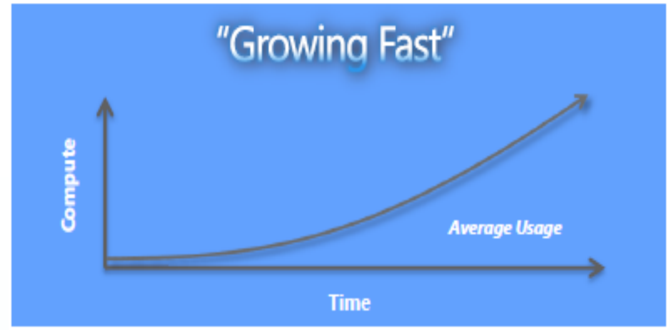

- Successful services needs to grow/scale
- Keeping up w/ growth is big IT challenge
- Complex lead time for deployment

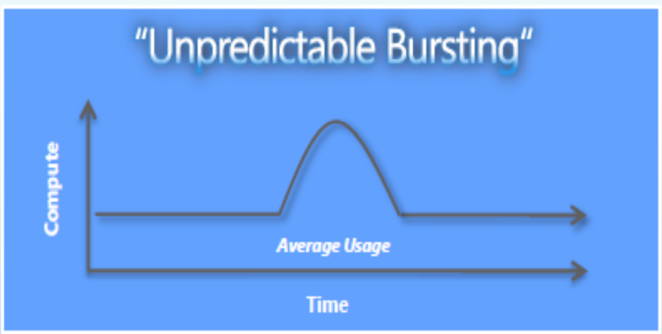

- Unexpected/unplanned peak in demand
- Sudden spike impacts performance
- Can't over provision for extreme cases

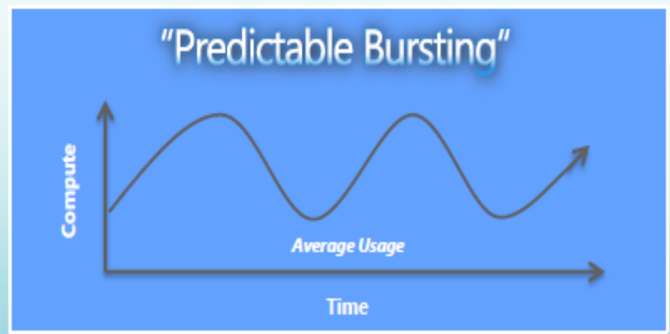

- Services with micro seasonality trends
- Peaks due to periodic increased demand
- IT complexity and wasted capacity

Windows Azure

Cloud computing costs *(Amazon EC2)*

- CPU usage *(see table)*
  - Pennies per hour add up!
  - Note 100:1 cost ratio
- Data Storage
  - EBS: $100/TB/month
  - S3: $37-140/TB/month
- Data Transfer
  - In: $100/TB
  - Out: $80-150/TB
- Other
  - SQL queries; I/O requests
  - Snapshot GETs/PUTs

Instance type	cost/hr	/day	/mo	/yr
t1.micro	$0.02	$0.48	$15	$175
m1.small *(default)*	$0.085	$2	$62	$745
m1.large	$0.34	$8	$248	$2,978
m1.xlarge	$0.68	$16	$496	$5,957
c1.xlarge	$0.68	$16	$496	$5,957
m2.xlarge	$0.50	$12	$365	$4,380
m2.2xlarge	$1.00	$24	$730	$8,760
m2.4xlarge	$2.00	$48	$1,460	$17,520
cc1.4xlarge *(cluster compute)*	$1.60	$38	$1,169	$14,026
cg1.4xlarge *(GPU cluster)*	$2.10	$50	$1,534	$18,409

Advantages/Disadvantages: Computation Solution 2

Advantages:

- Scalable on demand – only pay for what you use
- Control over image to be deployed (hibernate)
- Flexible/Upgradable hardware arrangements

Disadvantages:

- Nickel and dime you for everything – for profit!
- Unknown hardware dependences (floating point?)

Questions?