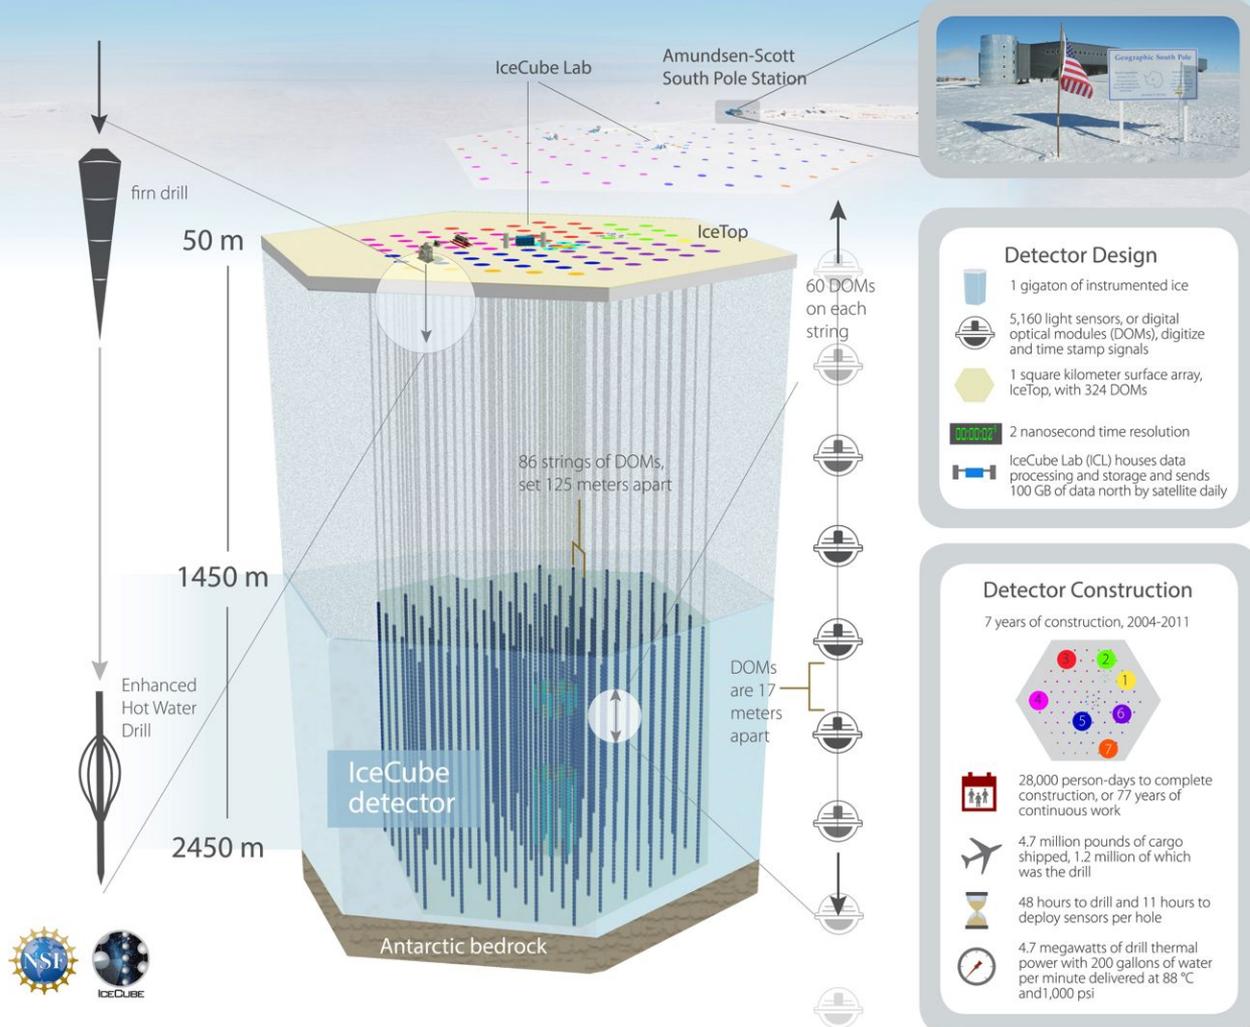


# IceCube Data Management

Rucio Community Workshop  
March 2018, CERN

G. Merino, P. Meade



### Detector Design

-  1 gigaton of instrumented ice
-  5,160 light sensors, or digital optical modules (DOMs), digitize and time stamp signals
-  1 square kilometer surface array, IceTop, with 324 DOMs
-  2 nanosecond time resolution
-  IceCube Lab (ICL) houses data processing and storage and sends 100 GB of data north by satellite daily

### Detector Construction

7 years of construction, 2004-2011



-  28,000 person-days to complete construction, or 77 years of continuous work
-  4.7 million pounds of cargo shipped, 1.2 million of which was the drill
-  48 hours to drill and 11 hours to deploy sensors per hole
-  4.7 megawatts of drill thermal power with 200 gallons of water per minute delivered at 88 °C and 1,000 psi

Neutrino observatory at the geographic South Pole

5483 Digital Optical Modules (DOMs) in 1km<sup>3</sup> of ice

7 years after completion the detector continues to operate nominally.

- More reliable than expected
- Only 32 of 5483 DOMs have failed

In Apr 2016: NSF renewed UW-Madison M&O contract until 2021



**49 institutions**  
**12 countries**  
**>300 researchers**

# THE ICECUBE COLLABORATION

 **AUSTRALIA**  
University of Adelaide

 **BELGIUM**  
Université libre de Bruxelles  
Universiteit Gent  
Vrije Universiteit Brussel

 **CANADA**  
SNOLAB  
University of Alberta–Edmonton

 **DENMARK**  
University of Copenhagen

 **GERMANY**  
Deutsches Elektronen-Synchrotron  
ECAP, Universität Erlangen–Nürnberg  
Humboldt–Universität zu Berlin  
Ruhr–Universität Bochum  
RWTH Aachen University  
Technische Universität Dortmund  
Technische Universität München  
Universität Mainz  
Universität Wuppertal  
Westfälische Wilhelms–Universität  
Münster

 **JAPAN**  
Chiba University

 **NEW ZEALAND**  
University of Canterbury

 **REPUBLIC OF KOREA**  
Sungkyunkwan University

 **SWEDEN**  
Stockholms Universitet  
Uppsala Universitet

 **SWITZERLAND**  
Université de Genève

 **UNITED KINGDOM**  
University of Oxford

 **UNITED STATES**  
Clark Atlanta University  
Drexel University  
Georgia Institute of Technology  
Lawrence Berkeley National Lab  
Marquette University  
Massachusetts Institute of Technology  
Michigan State University  
Ohio State University  
Pennsylvania State University  
South Dakota School of Mines and  
Technology

Southern University  
and A&M College  
Stony Brook University  
University of Alabama  
University of Alaska Anchorage  
University of California, Berkeley  
University of California, Irvine  
University of California, Los Angeles  
University of Delaware  
University of Kansas  
University of Maryland  
University of Rochester

University of Texas at Arlington  
University of Wisconsin–Madison  
University of Wisconsin–River Falls  
Yale University

## FUNDING AGENCIES

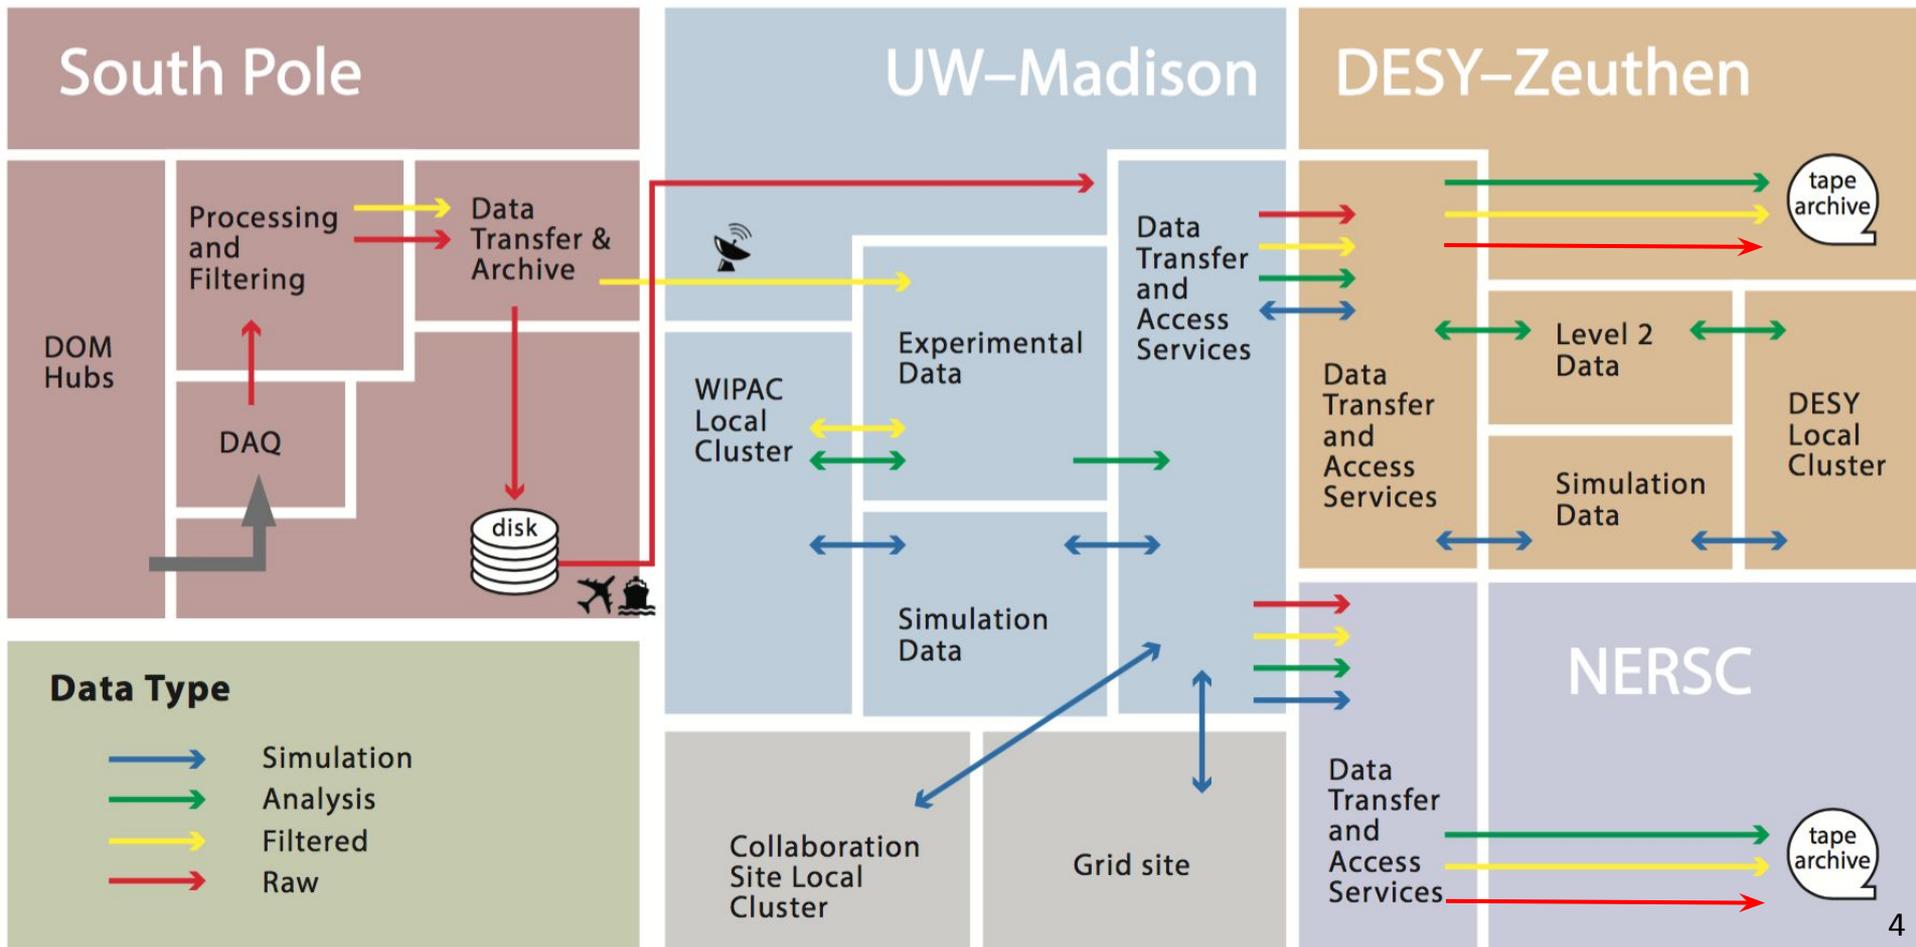
Fonds de la Recherche Scientifique (FRS-FNRS)  
Fonds Wetenschappelijk Onderzoek-Vlaanderen  
(FWO-Vlaanderen)

Federal Ministry of Education and Research (BMBF)  
German Research Foundation (DFG)  
Deutsches Elektronen-Synchrotron (DESY)

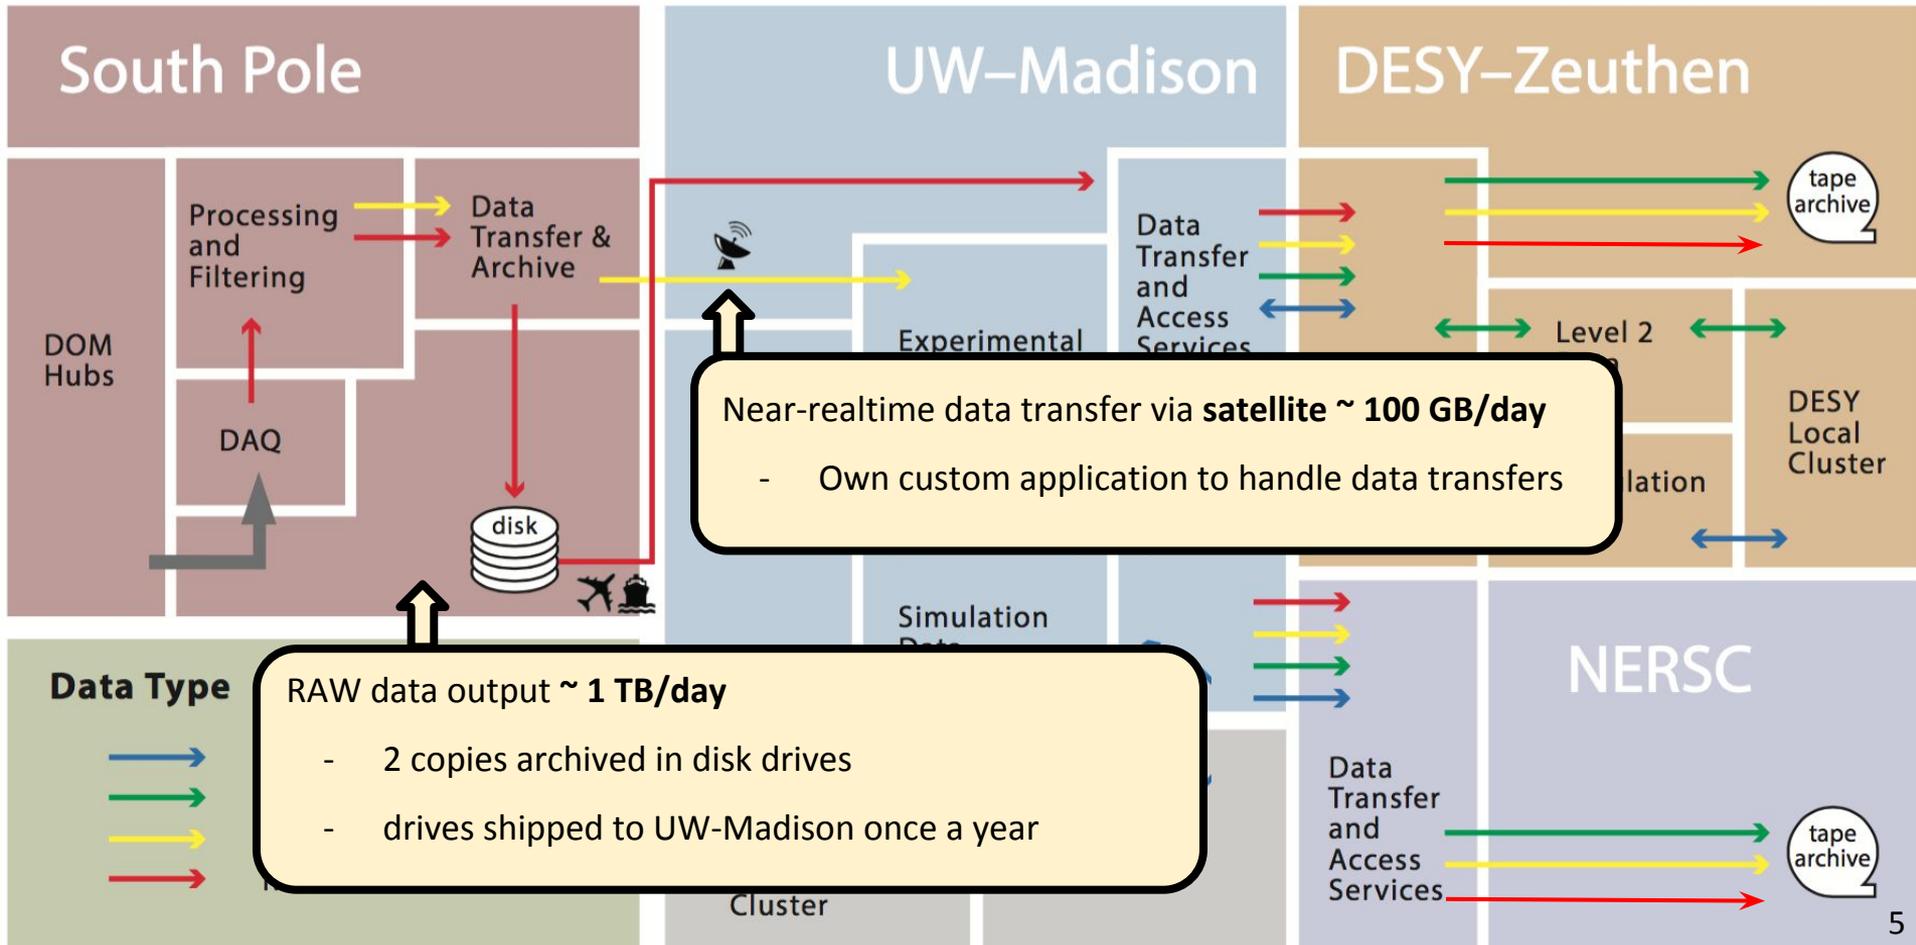
Japan Society for the Promotion of Science (JSPS)  
Knut and Alice Wallenberg Foundation  
Swedish Polar Research Secretariat

The Swedish Research Council (VR)  
University of Wisconsin Alumni Research Foundation (WARF)  
US National Science Foundation (NSF)

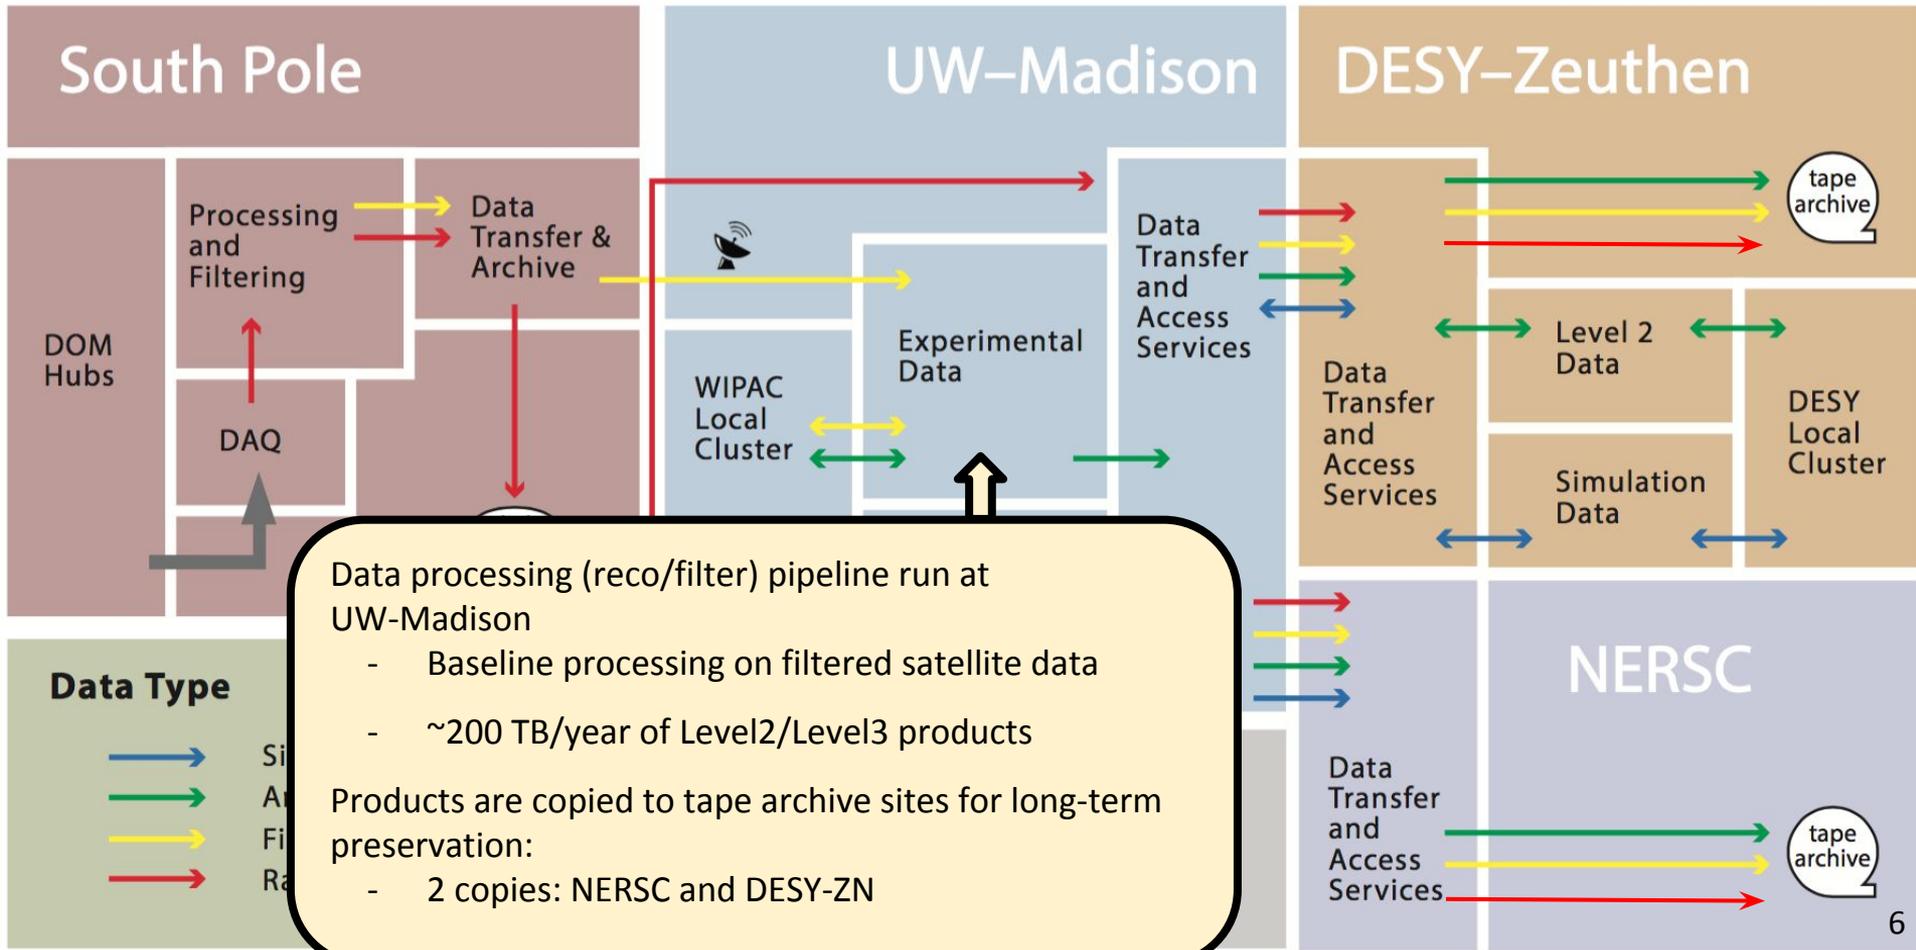
# IceCube Data Flow



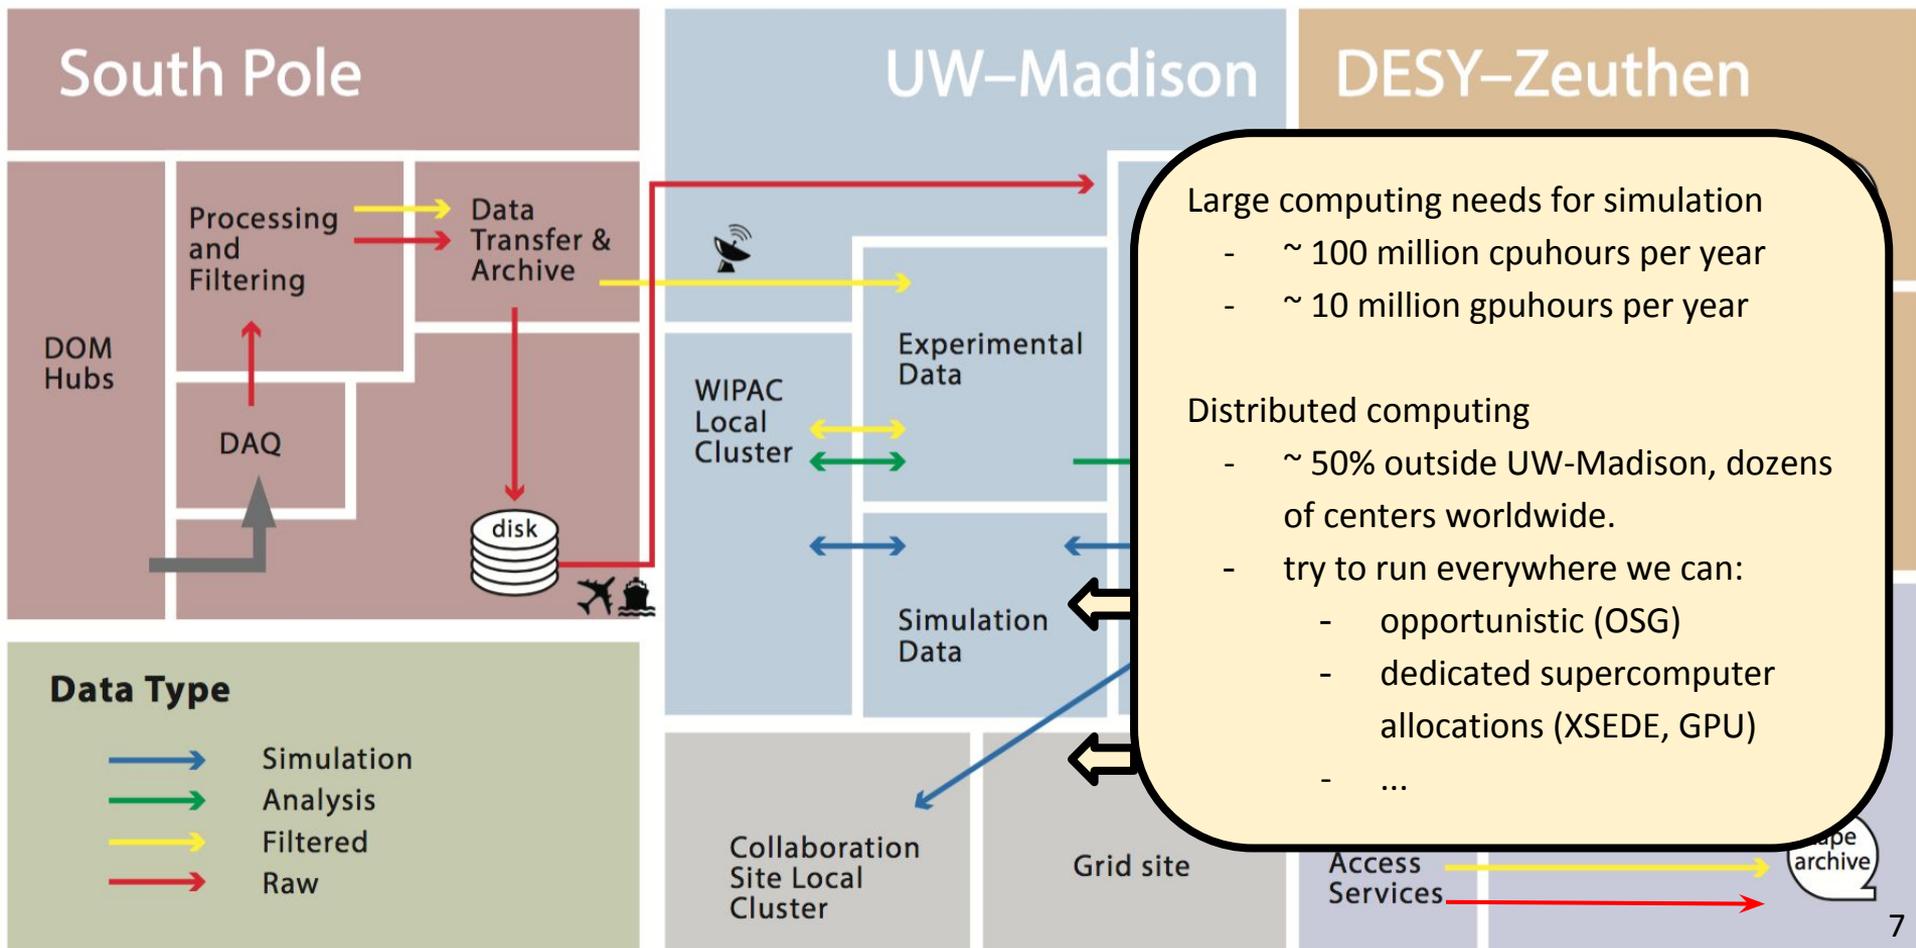
# IceCube Data Flow



# IceCube Data Flow



# IceCube Data Flow



# Data Management: JADE

In-house developed application to handle data files: transfer, archive, bookkeeping, etc.

**South Pole JADE:** consumption from DAQ, packaging, archive and satellite transfer

**JADE North:** Indexing and warehousing satellite data to the disk at UW-Madison

**JADE LTA** (Long Term Archive):

- Indexing and bundling (>500GB zip bundles) archival data
- Transfer to tape archives: UW-Madison → DESY-ZN, NERSC
  - Currently using Globus for handling GridFTP transfers to tape archives

# Data retention/archival policies

Data type	Subtype	Growth (TB/yr)	DESY-ZN tape	NERSC tape	Years on disk at UW-Madison
Experimental	RAW	320		yes	2
	SuperDST	70	yes	yes	2
	Filtered	40	yes	yes	2
	Level2	100	yes	yes	3
	Level3	100		yes	10
Simulation	Level2	400			3
	Level3	100		yes	10
	Photon tables	8			5

The total file count since detector started is ~500M (not counting private user files)

Current data model is very centralized: only 4 Storage Elements

- Primary Disk copy at UW-Madison (main processing and analysis facility)
- Disk storage at DESY-ZN with a replica of L2,L3 data and some simulation data
- Two archival (Tape) sites: DESY-ZN and NERSC

# IceCube files

IceCube data is organized in files. Each file contains a number of “frames” (~events ~triggers).

One RAW file for instance contains ~2.5 minutes of detector data (~1.2 GB)

Data taking is organized in “runs”: run changes every 8 hours, or if detector configuration changes.

Some experimental files examples:

```
/data/exp/IceCube/2016/filtered/level2/0101/Run00127347/Level2_IC86.2015_data_Run00127347_Subrun00000147.i3.bz2
```

```
/data/exp/IceCube/2016/unbiased/PFRaw/0601/PFRaw_PhysicsFiltering_Run00127988_Subrun00000000_00000012.tar.gz
```

# File Metadata

The path/filename encodes a number of file metadata fields:

```
/data/exp/IceCube/2016/filtered/level2/0101/Run00127347/Level2_IC86.2015_data_Run00127347_Subrun00000147.i3.bz2
```

```
data_type = experimental
create_date = 2016-01-01T20:34:54.38Z
processing_level = level2
run_number = 127347
season = 2015
```

Users often discover data by browsing the filesystem.

# File Metadata

There are several additional metadata fields that we want to keep record of.

These are currently in bookkeeping DBs internal to the various s/w systems that handle files.

We are in the process of building a “file metadata catalog” that will be the core repository for this information.

<code>logical_name</code>	<code>start_datetime</code>	<code>software</code>
<code>file_size</code>	<code>end_datetime</code>	<code>validation.valid</code>
<code>checksum</code>	<code>first_event</code>	<code>dated</code>
<code>content_status</code>	<code>last_event</code>	<code>validation.date</code>
		<code>...</code>

# Workflows

Two different methods used by two different groups:

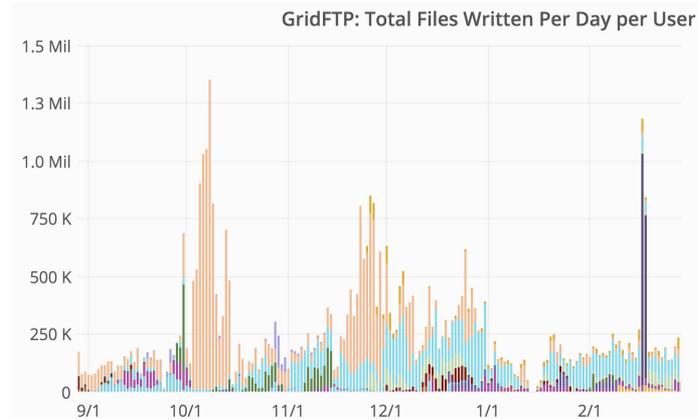
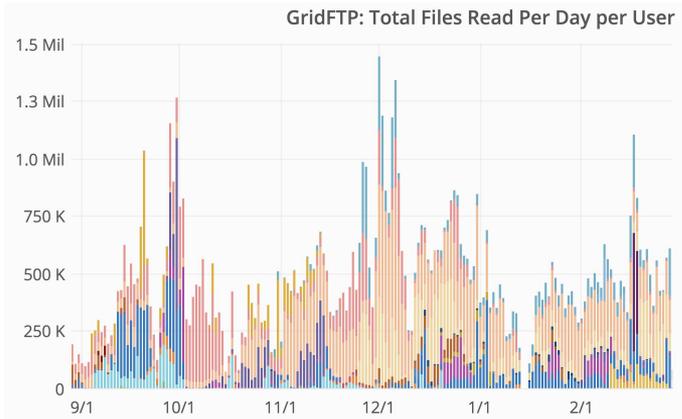
- Organized simulation production and processing
  - Using in-house developed framework called IceProd
  - Mainly “sits on top of HTCondor” for job submission, but handles dataset prioritization, bookkeeping, etc.
  - Few users (~2-5) driving most of the data “writes”
- User analysis
  - Mostly use HTCondor DagMan for managing workflows
  - Mostly reading the “official” data produced by the 1st group, and writing private analysis user output (no bookkeeping)

# Data access

Jobs running in the central cluster at UW-Madison access the data directly (Lustre mounted filesystem)

Jobs running in the Grid (~50% of total) access data at UW-Madison via GridFTP

- x509 certs and VOMS are the technologies currently in use
- interested in contributing to/prototyping alternatives (e.g. SciTokens, etc)



# Data sharing

Public data releases are still limited to targeted publications

- Published data varies from publication to publication: histogram bins, contour values, event lists ...

A generic data release containing standard samples of reconstructed events is work in progress.

<http://icecube.wisc.edu/science/data>

## Download Datasets

The pages below contain information about the data that were collected and links to the data files. We ask that you provide your contact information so that we may notify you if we revise the datasets in the future.

Dataset	Release Date
<a href="#">Measurement of atmospheric neutrino oscillations with three years of data from the full sky</a>	2018 Feb 13
<a href="#">A combined maximum-likelihood analysis of the astrophysical neutrino flux</a>	2016 Nov 15
<a href="#">Search for point sources with first year of IC86 data</a>	2016 Nov 01
<a href="#">Search for sterile neutrinos with one year of IceCube data</a>	2016 Jun 24
<a href="#">The 79-string IceCube search for dark matter</a>	2016 Jan 05
<a href="#">Observation of Astrophysical Neutrinos in Four Years of IceCube Data</a>	2015 Oct 21
<a href="#">Astrophysical muon neutrino flux in the northern sky with 2 years of IceCube data</a>	2015 Aug 20
<a href="#">IceCube-59: Search for point sources using muon events</a>	2015 Jun 19
<a href="#">Search for contained neutrino events at energies greater than 1 TeV in 2 years of data</a>	2015 Feb 19
<a href="#">IceCube Oscillations: 3 years muon neutrino disappearance data</a>	2015 Jan 27
<a href="#">Search for contained neutrino events at energies above 30 TeV in 2 years of data</a>	2013 Nov 21
<a href="#">IceCube-40 String Data</a>	2011 Sep 05
<a href="#">IceCube-22 Solar WIMP Data</a>	2009 May 21
<a href="#">AMANDA 7 Year Data</a>	2008 Sep 11

# Data traceability

Q: How do you link your scientific publications to the data that you used?

- (Great question!) unfortunately, for now, we don't
- This is a HOT topic inside the collaboration
  - We plan to develop tools for users to be able to register analysis datasets linked to publications and record provenance information
  - Needless to say that we are happy to share ideas and discuss with anyone out there doing this

# Rucio Evaluation

We are interested in understanding the details of Rucio capabilities for data management

- The goal is to explore integration possibilities (avoid reinventing the wheel)

Some parts of our data management system (South Pole, satellite) are custom to the specific environment, but some others are doing very standard things

Our first goal is to make a limited scope evaluation by using Rucio to perform a task that we currently do with JADE LTA

- Replication of L2/L3 processed data from UW-Madison disk to DESY-ZN disk

With this exercise we hope to learn about the Rucio internals for then being able to propose potential further integrations with our system.

thanks

questions?