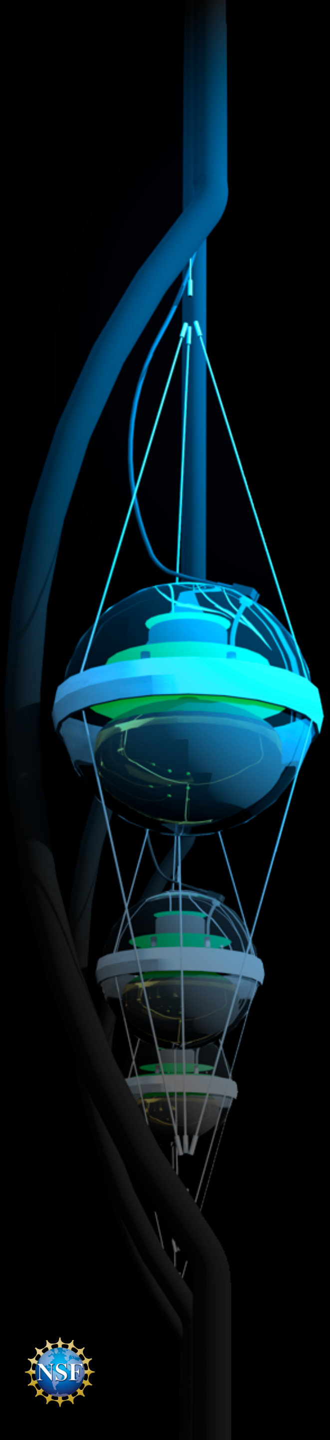

IceCube: Evolving Work Flows

Steve Barnet
University of Wisconsin – Madison
WIPAC

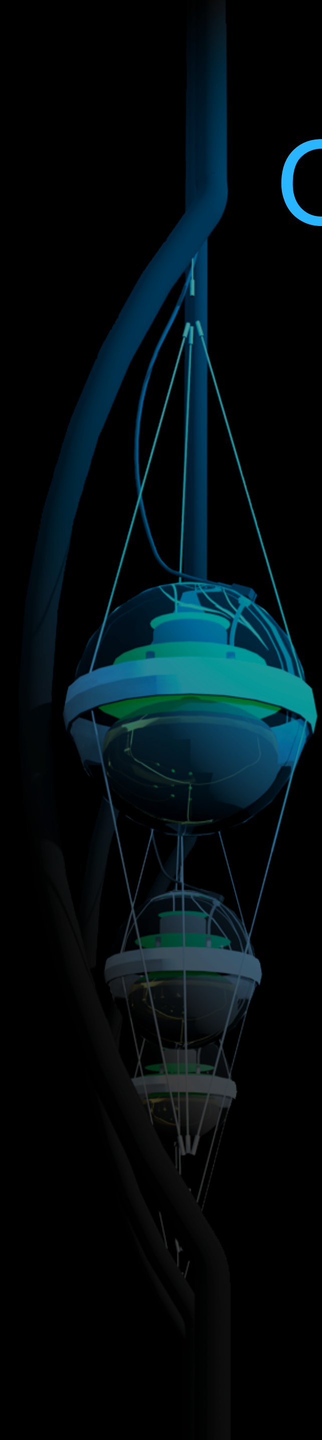

Overview

- The Basics of IceCube
- Data Handling
- Computing Model
- Summary

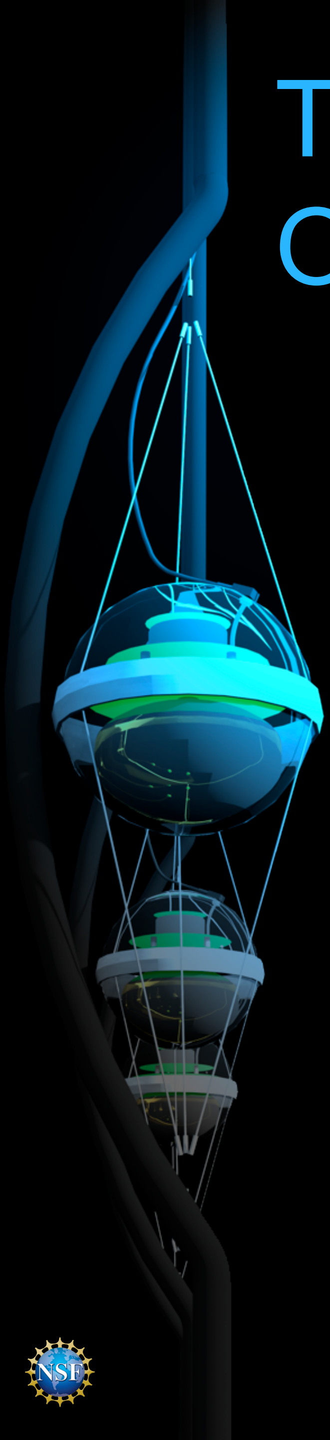

The IceCube Neutrino Observatory

- A kilometer scale neutrino detector
- Located at geographic South Pole
- Detects Cherenkov light from neutrino interactions

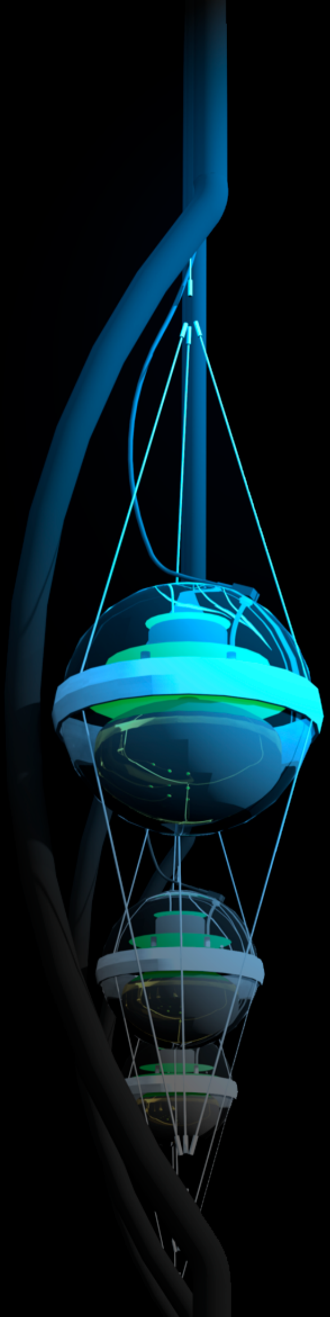

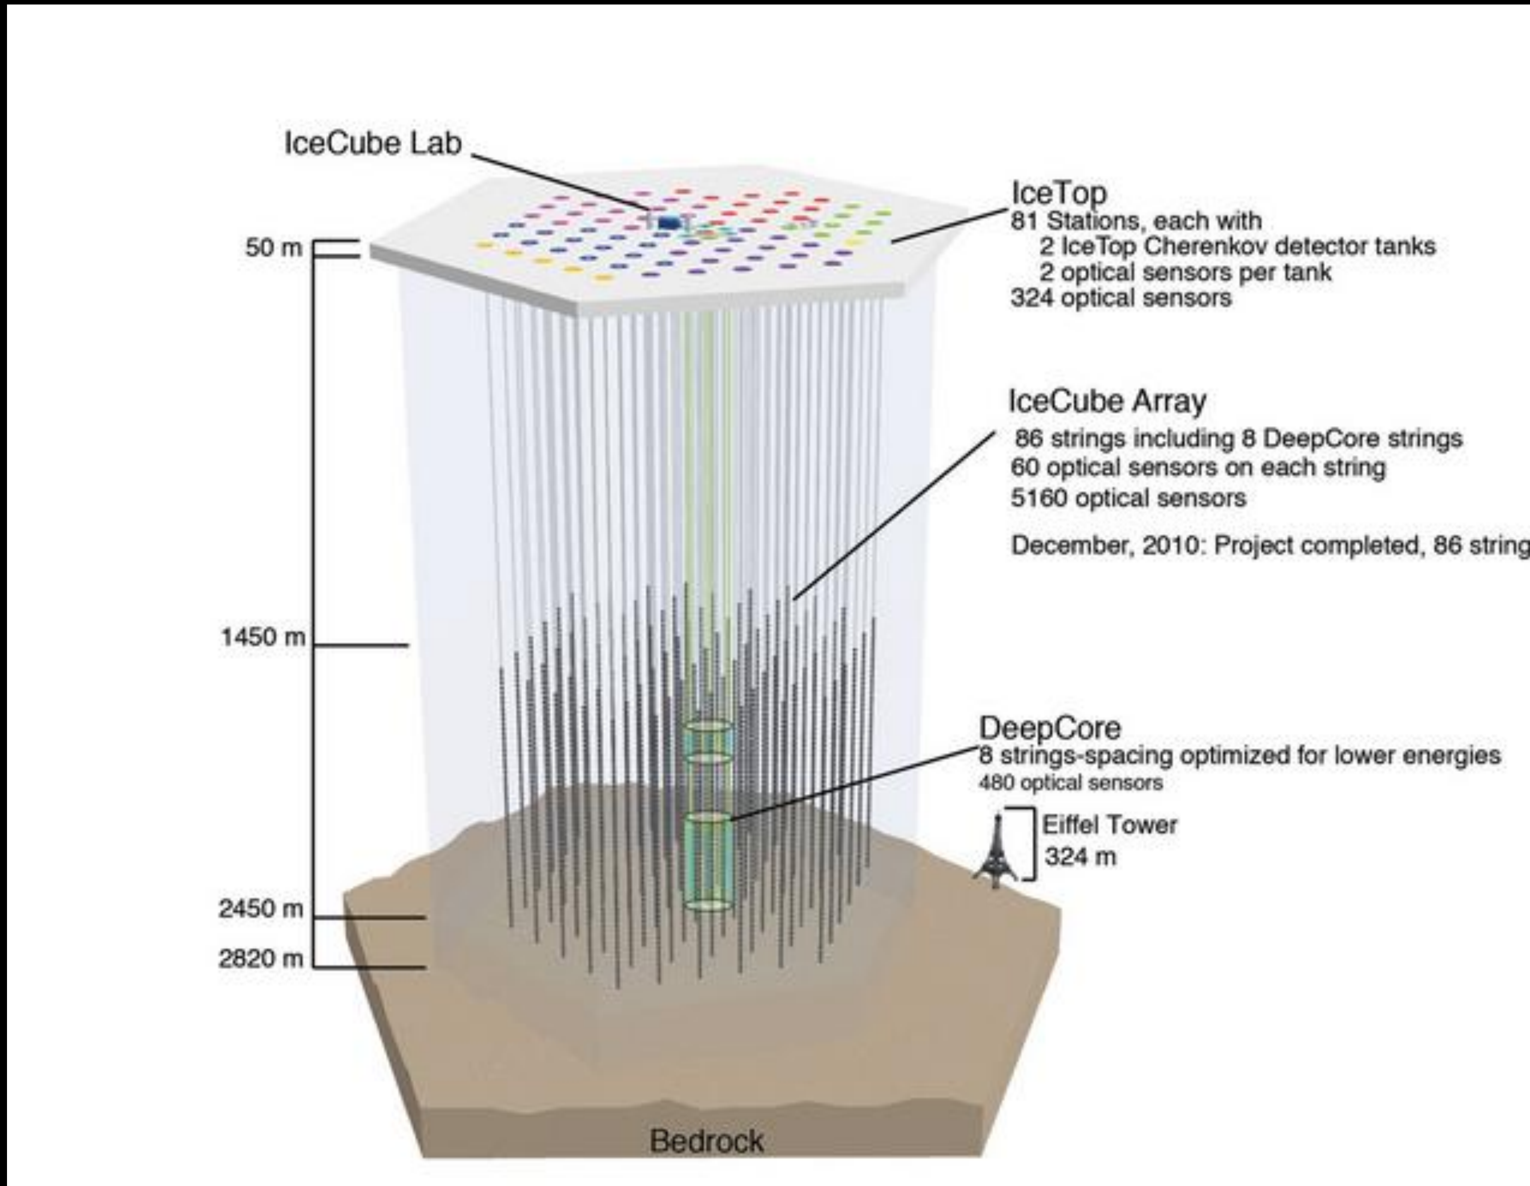
IceCube Lab
IceTop
81 Stations, each with
2 IceTop Cherenkov detector tanks
2 optical sensors per tank
324 optical sensors
50 m
IceCube Array
86 strings including 8 DeepCore strings
60 optical sensors on each string
5160 optical sensors
December, 2010: Project completed, 86 strings
1450 m
DeepCore
8 strings-spacing optimized for lower energies
480 optical sensors
Eiffel Tower
324 m
2450 m
2820 m
Bedrock

Location, Location, Location

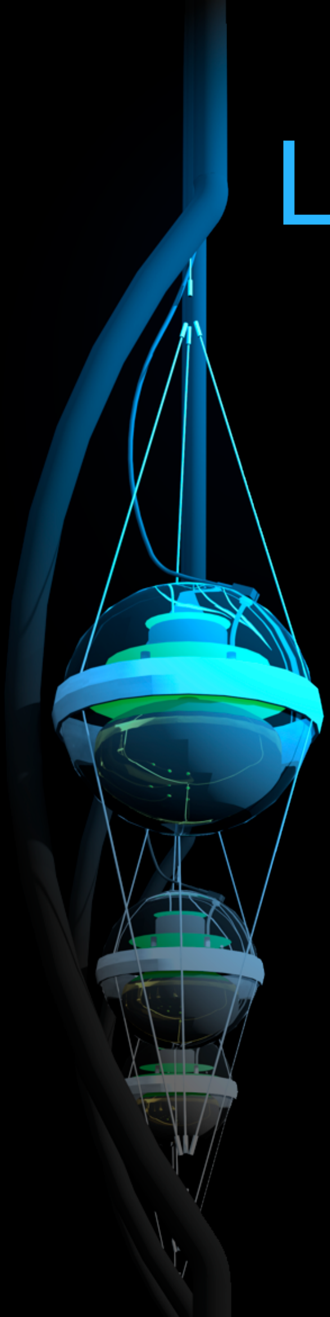

- Why the South Pole?
- Lots of ice – a great detection medium
- The ice is very clear
- Thick ice sheet – sensors deep enough to provide significant background reduction

IceCube Overview

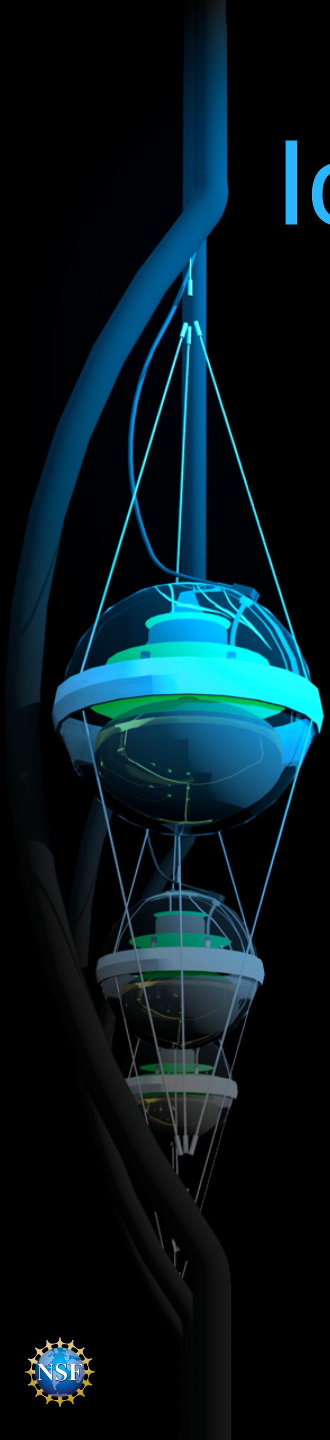

- 400 people
- 39 Institutions
- 11 Countries
- Exotic Locales

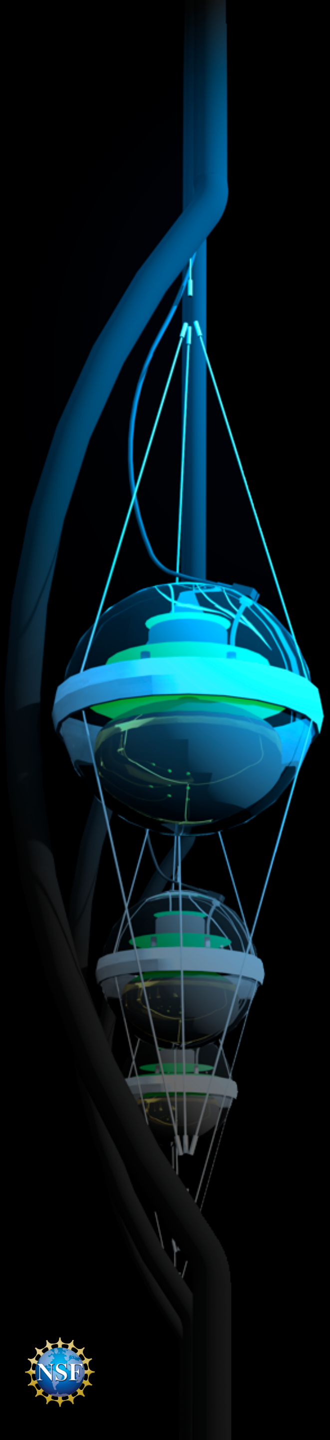


16

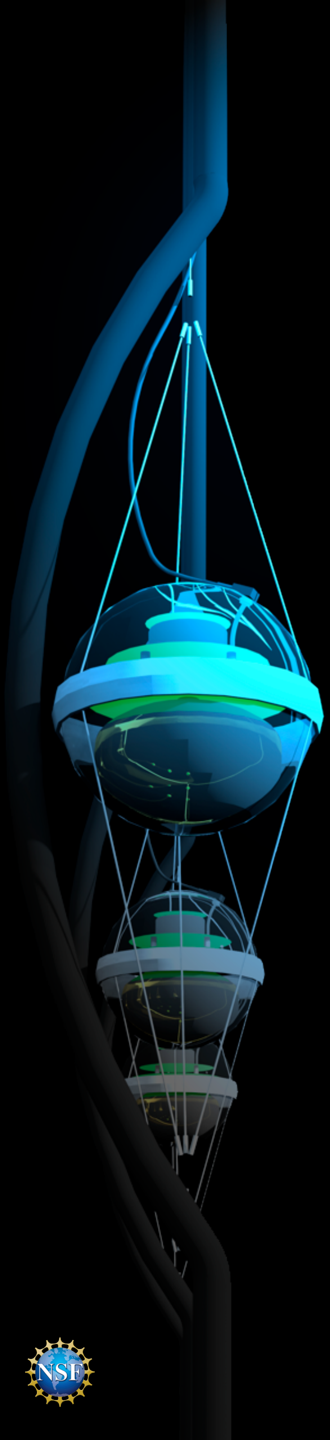

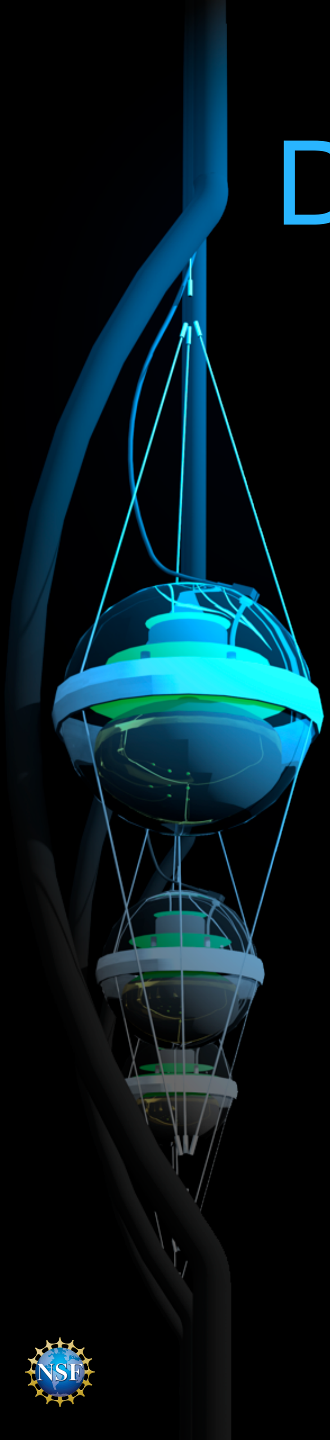

Data Volume

- Event rate of 3000 Hz
- 1 TB/day raw data
- Reduced to 100 GB/day
  - Based on available satellite bandwidth
- Tapes shipped North every year

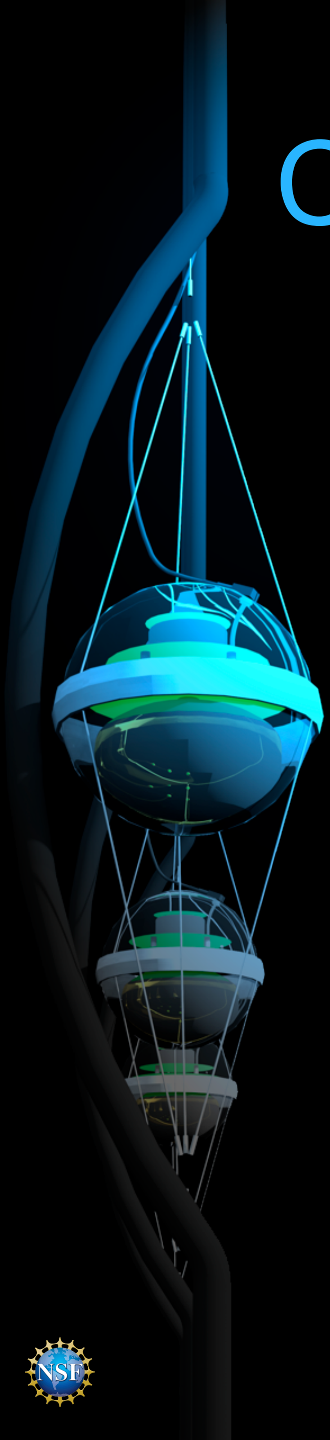

Computing Model

- University of Wisconsin - Madison – Tier 0
  - – Raw data collection and archive
  - – Data production to Level 3
  - – Coordinate simulation production
- DESY-Zeuthen – Tier 1
  - – Second copy of Level 2 data
  - – Hold simulation data sets in Europe

Tier 0 Capacity

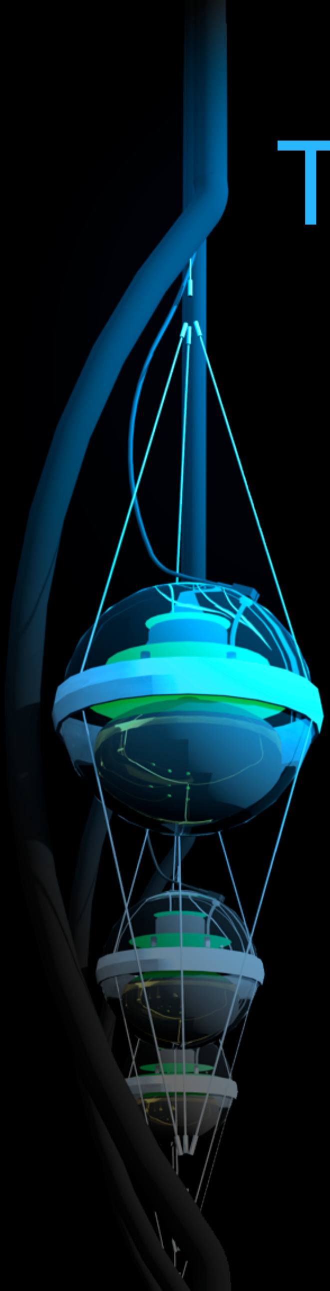

- Compute cluster
  - 1400 cores, memory at least 4 GB/core
  - HTCondor managed
- Experimental data
  - 1 PB Lustre filesystem
- Simulation data
  - 1.2 PB Lustre filesystem
- Analysis data sets
  - 250 TB Lustre filesystem

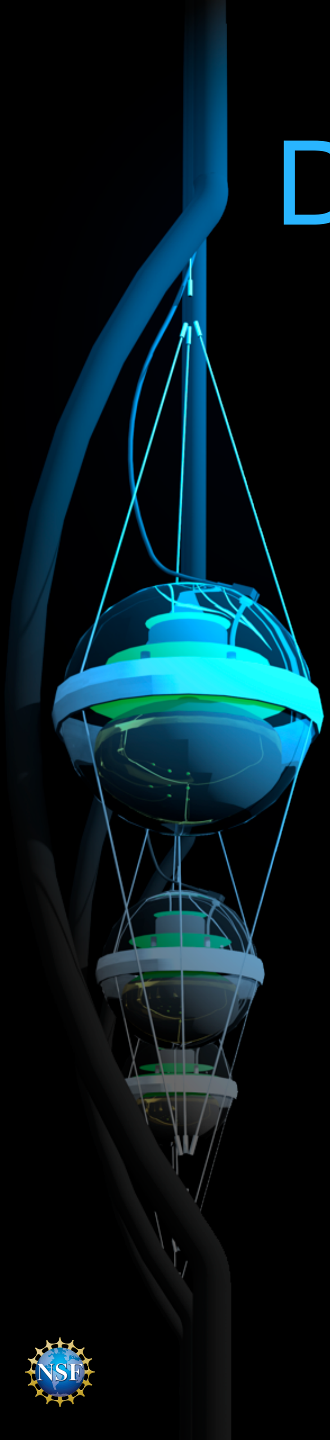

Data Production

- Level 1
  - Filtered stream from Pole
- Level 2 – Science Ready
  - Calibrations, basic reconstructions
  - Common processing needed by all analyses
- Level 3 – set by each working group
  - Event selection
  - Stream separation

Simulation Production

- Driver of our distributed computing
- Collaboration institutions contribute to simulation production
- Production is coordinated with a central DB at UW Madison
- Output is collected at either DESY-Zeuthen or UW Madison

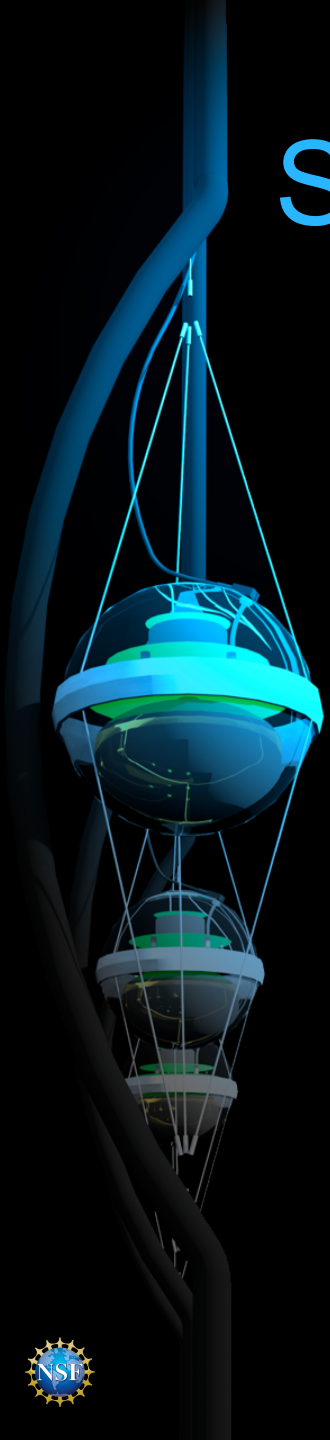

Simulation Framework

- Software framework to coordinate components
- Central DB at UW-Madison to coordinate and track production
- First versions were monolithic – photonics tables made grids difficult
- Newer versions broken into finer steps

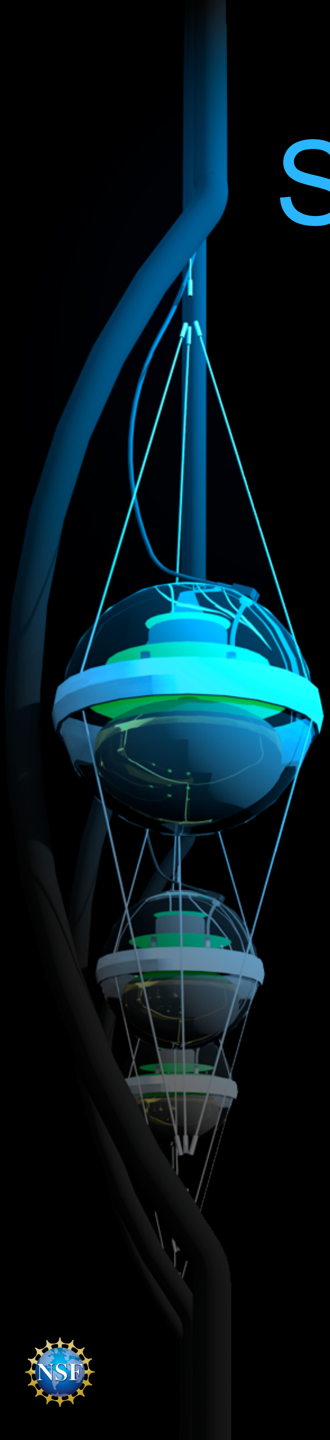

Simulation

- Event generation
  - CORSIKA for cosmic ray background
  - Several neutrino generators
- Photon propagation
  - Lookup tables
  - Direct simulation on GPUs
- Detector simulation

First Monolithic WorkFlow

- This was a single job
- A data set is O(100,000) jobs
- Long Photonics chain
- Lots of waiting
- Multiple detectors simulated
- Worked Locally, grids difficult

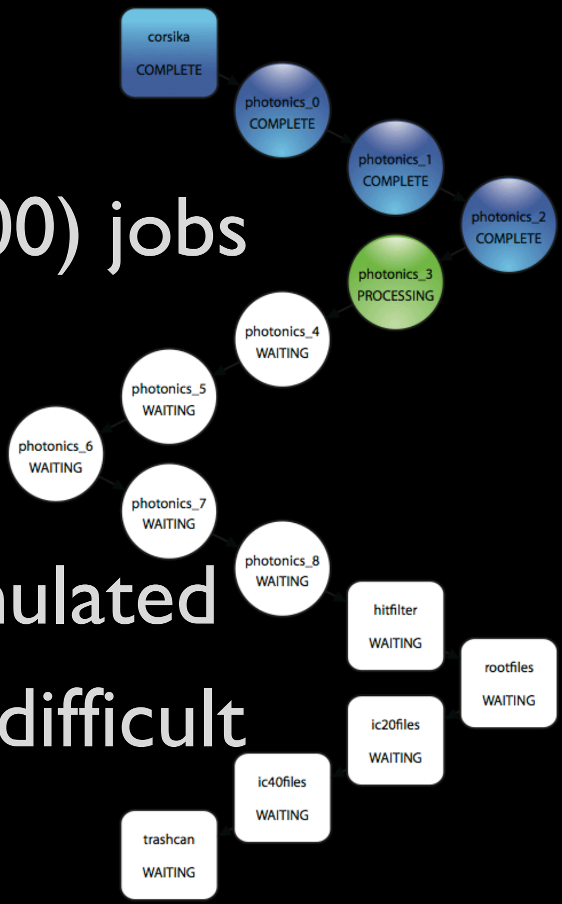

OSG – Attempt 1

- Refactor work flow using DAGMan
- glideinWMS is now on the scene
- + lots of local caching and clever job routing

Workable with caching

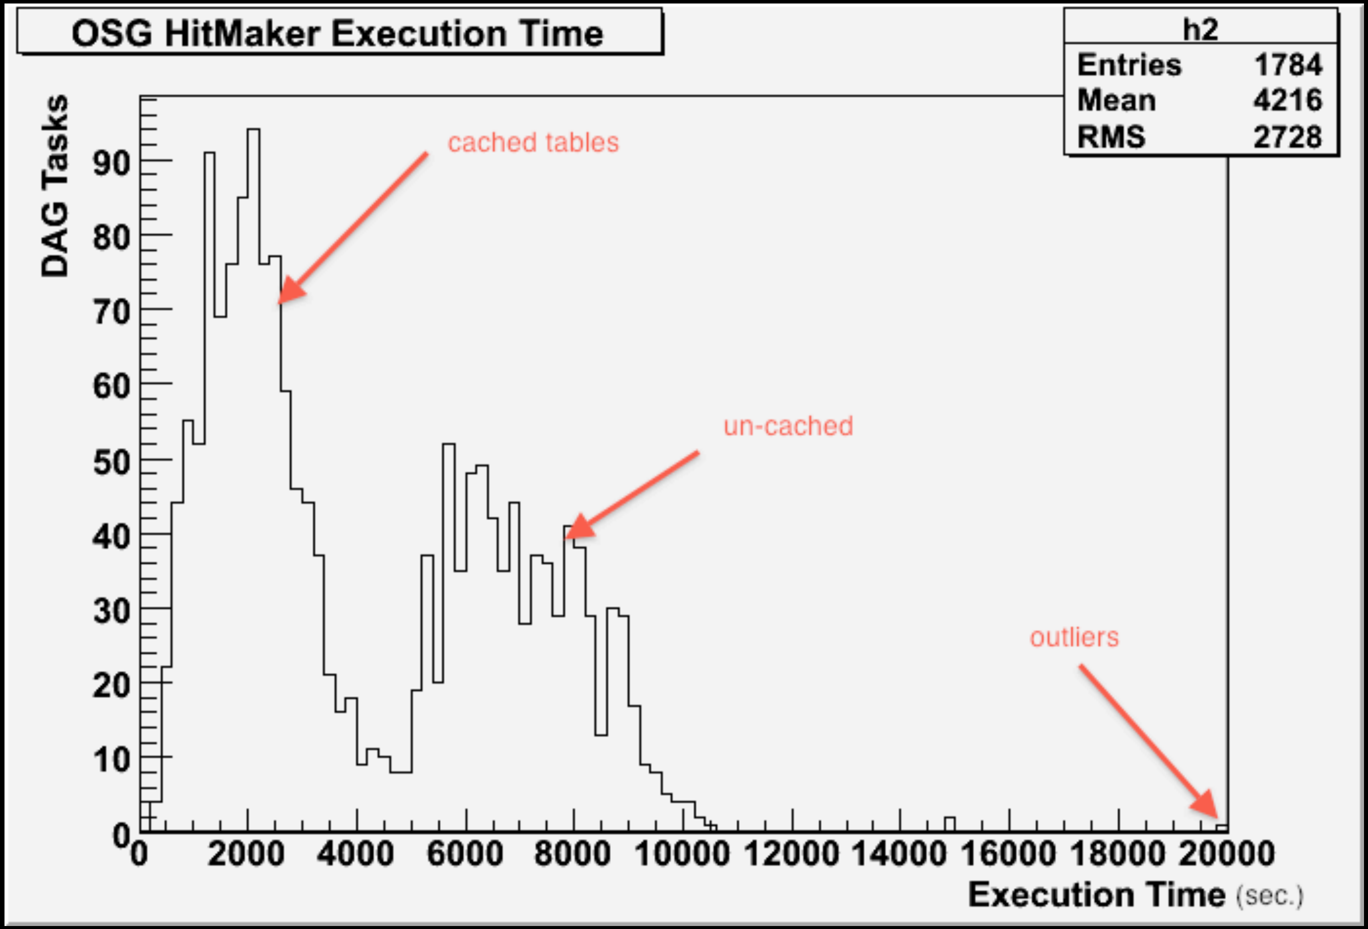

Enter GZK-9000

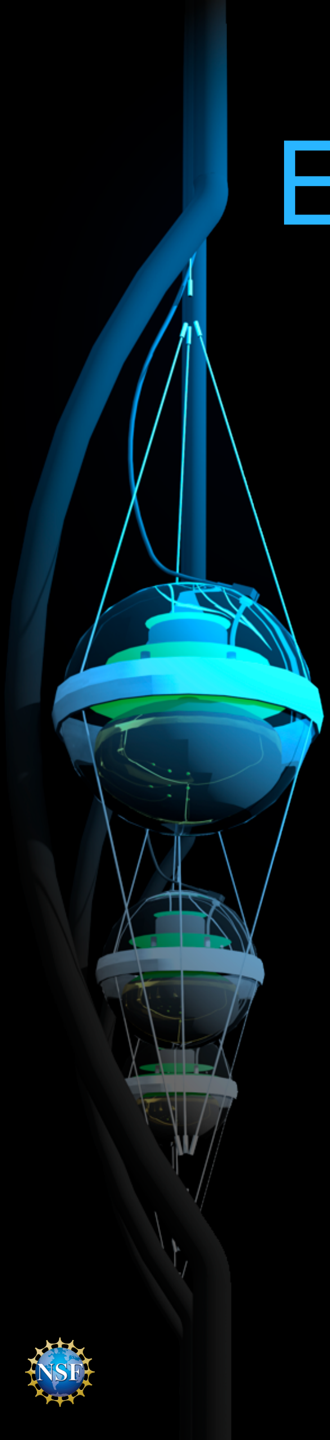

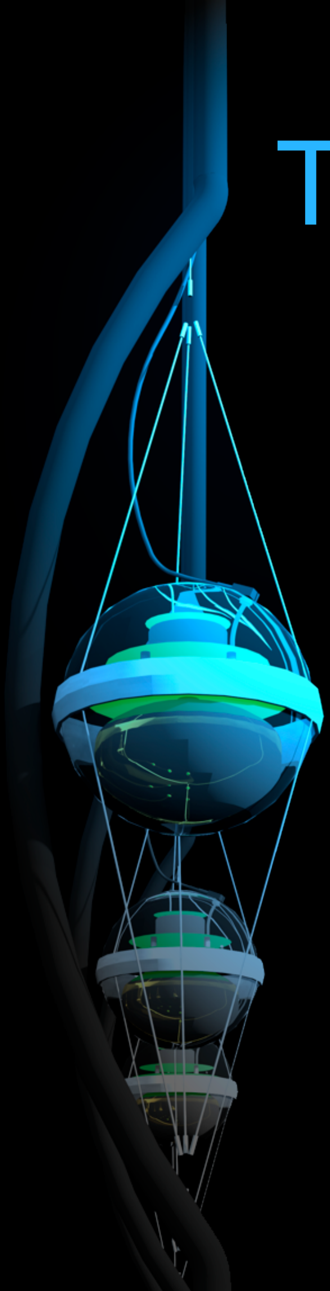

The Name

- Or: what happens when you give clever grad students a mission and some time
- Derived from HAL-9000 (2001: A Space Odyssey)
  - IBM << 1 == HAL
  - HAL << 1 == GZK
- GZK is a predicted limit on the energy of cosmic rays

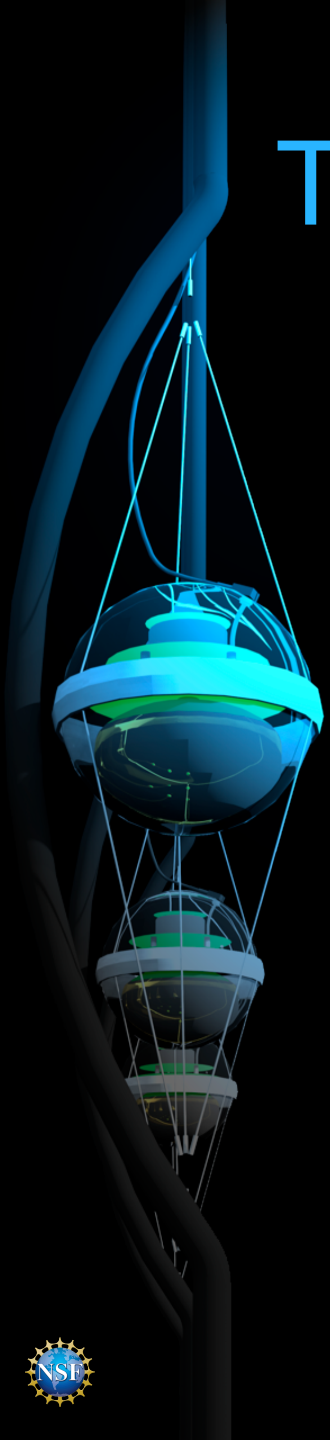

The Hardware

- 6x Dell C6145 series chassis for CPUs
- 3x Dell C410 Chassis for GPUs
  - 16x GPUs each
  - nVidia Tesla M2070

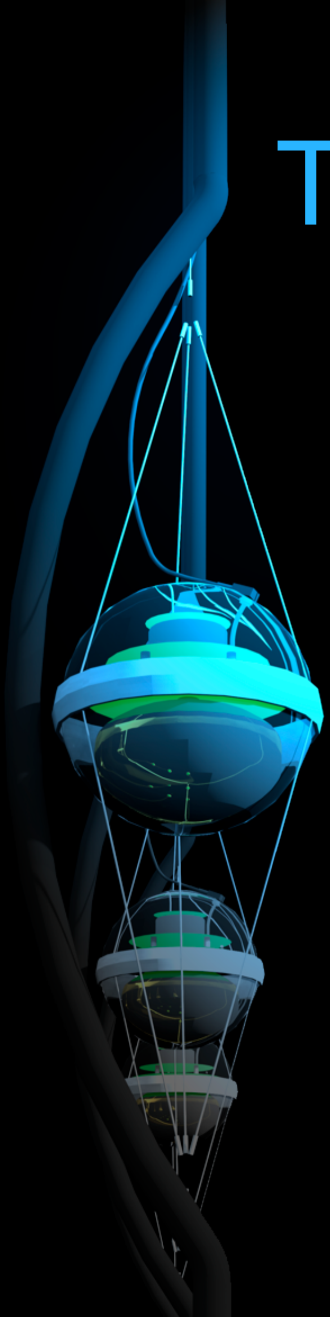

The Place

- CHTC pool
- Shared resource – Available to other CHTC users
- WID/MIR data center
- 29 U of rack space
- 15 KW power dissipation

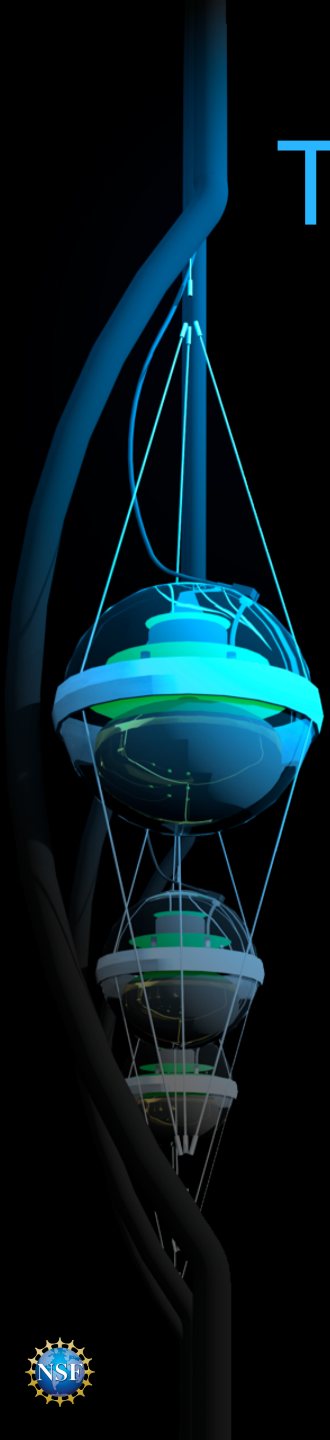

The Reason

- Improves the science
  - Eliminates many binning artifacts associated with tables
  - Required many improvements to our understanding of the ice
- Greatly reduces the need for Photonics tables

Simplifies the Work Flow

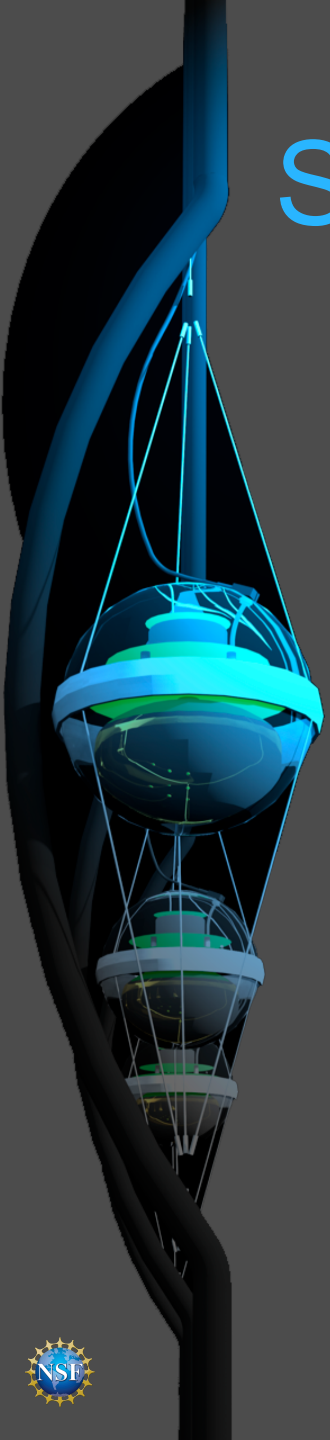

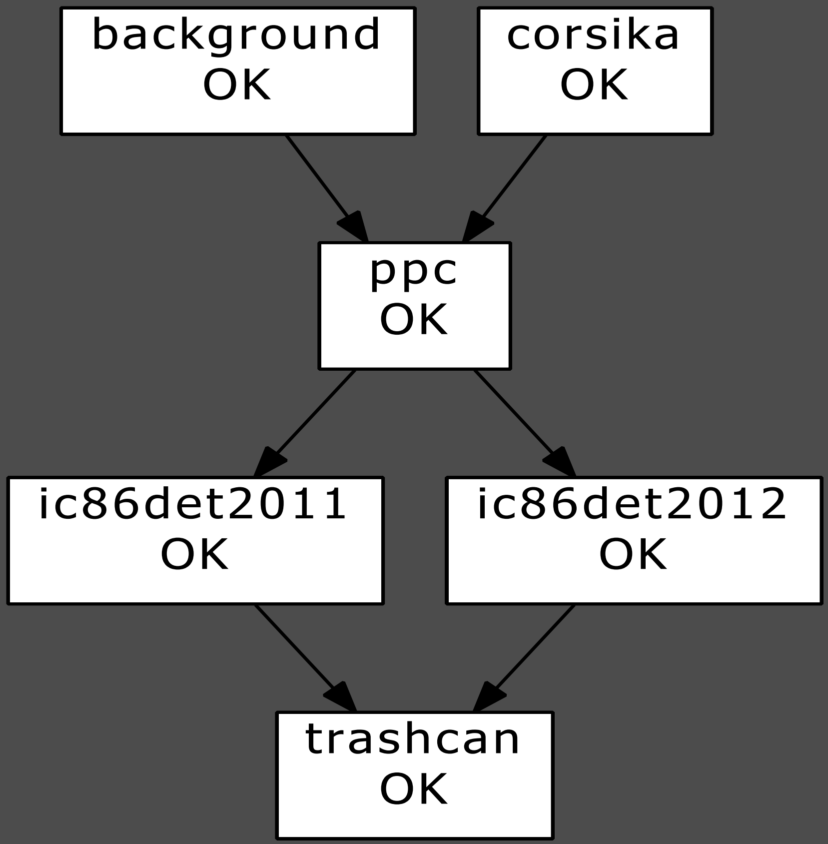

OSG Attempt 2

- Move top tier of DAG to OSG
- Results:
  - GPU processing at UW
  - Use local GPUs if available
- Turns our “tables problem” into a data movement problem – may be more tractable

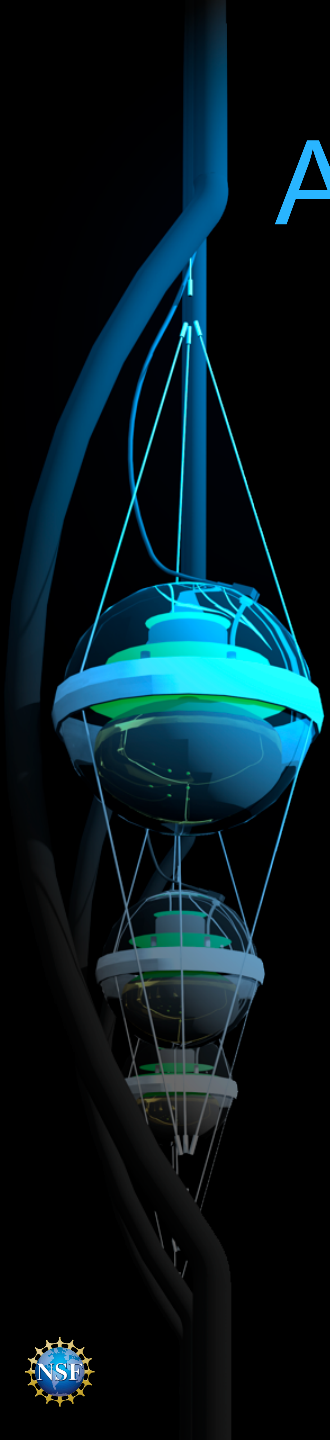

Approach

- Start small scale with CHTC resources
  - Submit machine and glideinWMS
  - Shake out the bugs and surprises
- Add IceCube resources as needed to feed the beast
  - CVMFS – software distribution, remaining tables
  - Squid caching
  - Expand gridFTP service

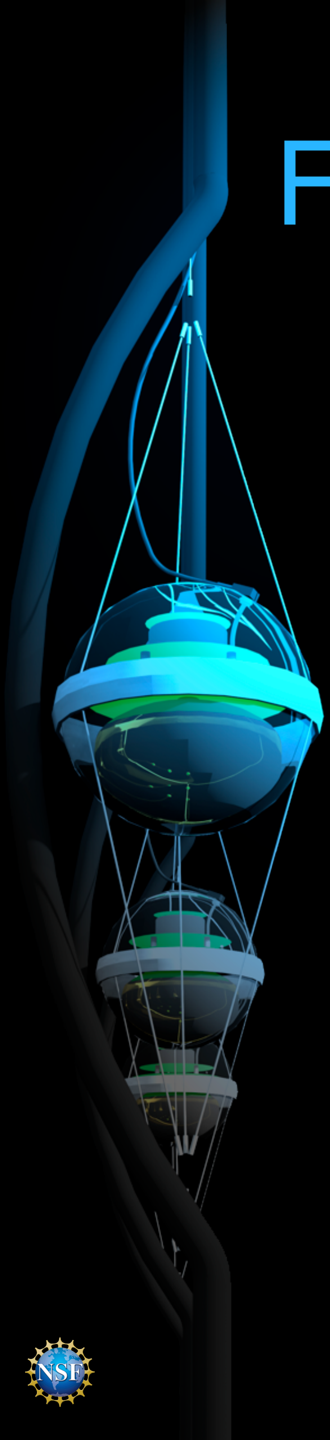

Future Workloads

- Analysis – get Jane the Physicist doing regular analysis using OSG
- Data processing (Level 2 production)
- Reformatting old data into archival format

GPUs

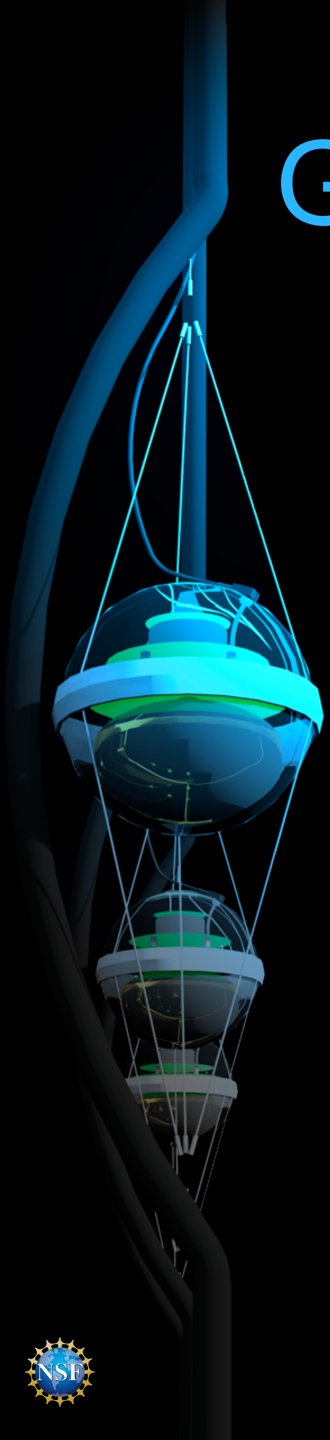

- Better integration with HTCondor
- CPU/GPU ratios – 2:1, 1:4 - ?
- Limit jobs to assigned GPUs

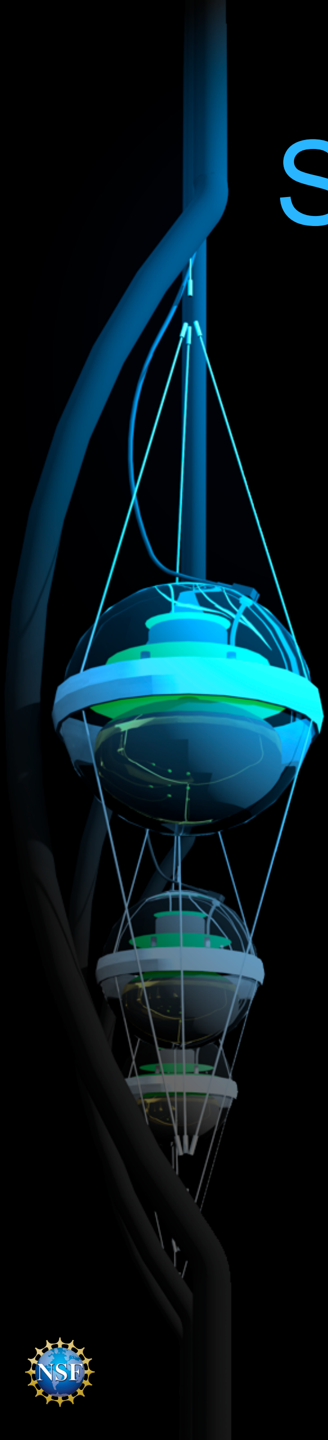

Summary

- IceCube is a km scale neutrino detector
- Work well locally with HTCondor
- Expanding to OSG – involves work
- Extensive use of GPUs
- We may not be the coolest experiment on the planet, but we are one of the coldest.

Credits

- Juan Carlos Diaz-Velez, David Schultz – IceProd framework
- Dima Chirkin, Claudio Kopper – GPU simulation development
- Heath Skarlupka – GPU wrangling
- CHTC – Hosting the GZK cluster, OSG support, HTCondor GPU efforts
- UW Madison – Fast networks
- NSF – IceCube M&O funding

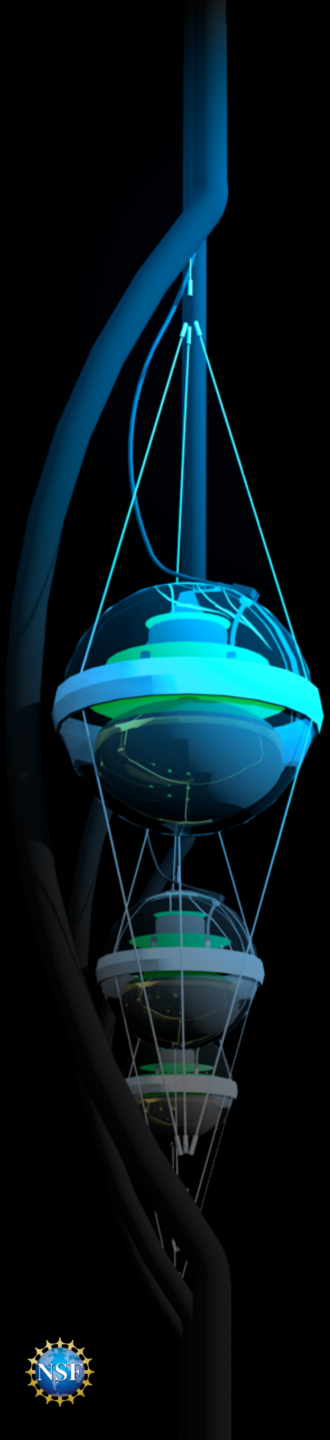

Questions?