Production Software since SCAP 2021

David Schultz

What is "Production Software"?

- Primarily, computing middleware in support of science
  - Not directly science software, but adjacent
  - Sometimes designed specifically for the needs of our science
  - Examples:
    - Job submission and file transfer services
    - Metadata handling
    - Data archival
- Also includes collaboration services
  - User management
  - Author lists
  - MoU tools

Software Issues from Production POV

2025-2

Changes in both physics software and grid computing have affected our performance and hardware utilization metrics

- GPU photon propagation
  - Snowstorm's rapid switching between ice parameters causes reduced GPU utilization, especially noticeable on modern, faster GPUs (as low as 30%)
  - Dependence on OpenCL is becoming an issue at more and more sites
    - Some sites will no longer run our code, and we blacklist them
    - Software team working on a rewrite using C++ std::par
- ML inference
  - Poor utilization because each event is too "small" so the workload is very bursty
  - Batching helps, but requires tuning. Not used by default

Software Issues from Production POV

Changes in both physics software and grid computing have affected our performance and hardware utilization metrics

- ARM architecture
  - Several external systems now deployed with ARM CPUs, more in planning
  - We purchased a local ARM machine for testing
  - Fixed minor issues to get physics software to run
  - Open question: how stable are our numerical libraries and algorithms?
    - We've noticed issues with fp math in the past
  - For our production software services, they all work on ARM now

Other Issues from Production POV

2025-2

Changes in who submits jobs are both good and bad

- Good: lots of people gaining access to wider computing resources
  - Something we've been pushing for years
  - Allows larger analyses to proceed faster
- Bad: poorer resource usage
  - Submitters less familiar with how OSG works, potential restrictions
  - Often times submitting jobs that are really short (seconds) or really long (days)
    - Short jobs stress the system, cause throughput to go down
    - Long jobs are restricted on where they run, and often get evicted and retry (high badput)

IceProd

Production job submission framework and historical provenance.
Also used by users for working group and analysis level work.

Lots of maintenance updates, and a new major release

- Many improvements:
  - Use modern HTCondor python bindings
  - Remove custom pilot, allow using HTCondor file transfer
  - Allow running in containers
  - Migrate to Kubernetes for backend services
  - New database hardware
- Completed 2024Q4

Big negative diff!
Less code to maintain

Long Term Archive

2025-3

Copying data to tape - raw, several filter levels, completed analyses

Currently have ~7 PB at NERSC

- Reduced time to archive retro disks to NERSC
  - Significant rewrites in 2022
  - Optimizing the process in 2023
  - 30 days to transfer in 2024, until discovering missing disks
- Added TACC archive in late 2024 - early 2025

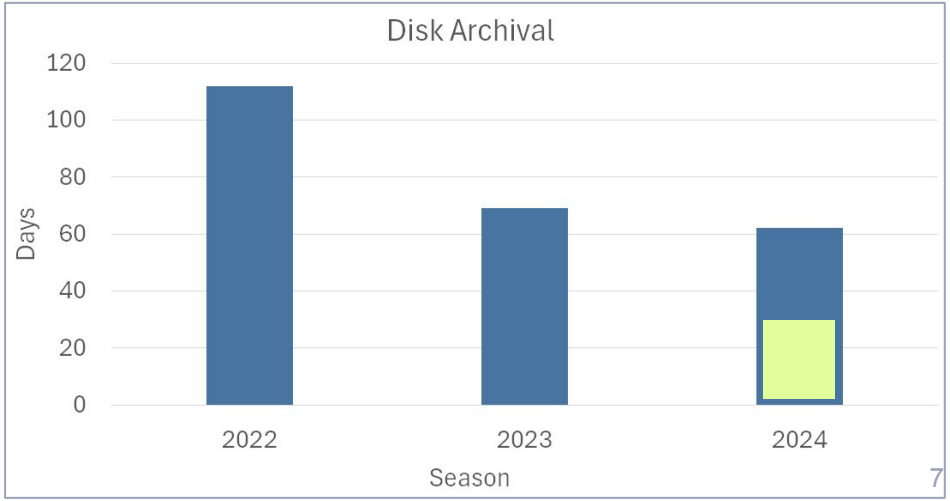

File Transfer

Have used GridFTP for file transfers for many years, based on X509

- Lots of investigation and dabbling in HTTP + tokens
  - Prototype embedding uid/gid in a Keycloak token in 2023
  - Used nginx + a lua script to change user
  - Integrated with HTCondor OAuth CredMon
  - It worked, but was complicated to use and never put in production
- Current effort based off of Pelican
  - Keycloak with a custom claims script to generate WLCG-compliant tokens
  - Set up a Pelican origin
  - Use HTCondor native support for Pelican
  - In progress now, scheduled to enter production in 2 months

Event Workflow Management Service

2021-4

NSF grant in partnership with CHTC / Morgridge Institute to build an event-based processing framework, running on top of HTCondor.

- Skymap scanner for realtime events
  - Offline reconstruction of community alert events, time sensitive
  - Beta used in production - 2023Q4
  - Complete - 2025Q1
- Adapting IceCube simulation framework to use EWMS
  - Goal for 2025-2026

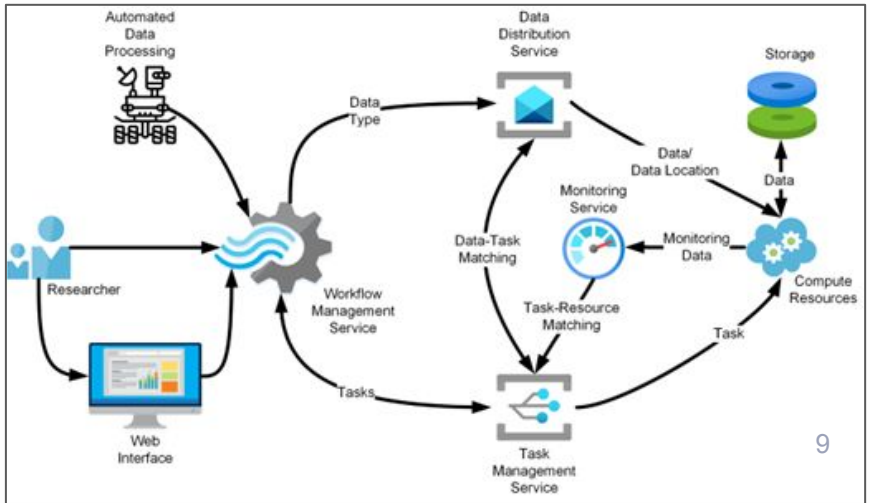

ML Inference-as-a-Service

2021-4
2025-2

Using SuperSONIC implementation in partnership with A3D3

"Services for Optimized Network Inference on Coprocessors" from CMS

- Triton inference servers running in kubernetes (NRP)
  - Automatically scales up with load
- Client in IceTray software, requests inferences as needed
  - Optimized for asynchronous operation, continuing other processing while waiting
- Currently in wide-scale testing, hoping to use for offline processing

User Management

New user management portal in 2022

- Gives institution leads power to manage their people
  - Create new user accounts
  - Add/remove people from their author lists
- Some group memberships now editable by managers
- Users can edit their attributes
  - author list name
  - ORCID
  - mailing list email
- Overall, decreased load on IT with more self-service

Metrics provided to NSF

Milestone/Metric	Target	Status
Data Transfer Delay	< 2 Days	PY1-3: ✅, PY3: 🚧/✅, PY4: ✅
Replication of Processed Data	<= 7 Days	PY1-2: ✅, PY3: 🚧, PY4: ❓
Replication of Raw Data	<= 90 Days	PY1-4: ✅
IceProd Uptime	>= 90%	PY1-4: ✅
Non-production IceProd users	20 Users	PY1-3: 🚧, PY4: ✅

- Processed Data: offline processing stopped in PY3, so nothing to replicate in PY4
- 23 non-production IceProd users in PY4
  - This is a little blurry since WGs now submit simulation datasets themselves

Bonus: HPC Wire Award

Best Use of HPC in the Cloud (Use Case)

Readers' Choice: The IceCube Neutrino Observatory is a gigaton-scale neutrino detector located at the South Pole. In a collaboration between the San Diego Supercomputer Center and the Wisconsin IceCube Particle Astrophysics Center at the University of Wisconsin–Madison, the team used Google Cloud, Google Kubernetes Engine, and GPU sharing with Nvidia GPUs in Google Kubernetes Engine to expand the Open Science Grid to help detect neutrinos at the South Pole. Sharing increased job throughput by about 40%.