Title: Results and Responses for the 2022 User Forum Survey

Author(s): Josey, Colin James
Grieve, Tristan Sumner
Rising, Michael Evan

Intended for: 2022 MCNP User Symposium, 2022-10-17/2022-10-21 (Los Alamos, New Mexico, United States)

Issued: 2022-10-11

Los Alamos NATIONAL LABORATORY

Results and Responses for the 2022 User Forum Survey

Colin Josey, Avery Grieve, Michael Rising

XCP-3 Monte Carlo Codes
Los Alamos National Laboratory

2022 MCNP User Symposium
October 17, 2022

Managed by Triad National Security, LLC, for the U.S. Department of Energy's NNSA.

Introduction

In September, we sent out a survey asking how you all use the code. We had several motivations for this:

1. We know what LANL users want, but it is much harder to get a grasp what external users want
2. There are many features that we *genuinely* had no idea if anyone used
3. The code has gone into an “active period” where many big changes are being made, and without asking or releasing the code, we didn’t know how they would be received
4. There is no real archive of the forum to look back at previous concerns
5. There was no single place to look
6. Some people may have just never stated their concerns

Summary of Results

- We expected 50 results.
- Would have been happy with 100.
- Got 225!
- Results so far have been *extremely* informative!

We intend to send out surveys like these more regularly. Hopefully at least every users symposium or closely after new software releases.

In the future:

- Will probably request emails
  - Many people posted comments that needed following up.
  - People mentioned bugs we'd not heard of / were lost to the sands of time / had insufficient detail to act on.

Format

The general format will be:

1. Show the results of a question, and discuss the impact of the results
2. Go through some selected results from the text-comments
3. Open things up to the audience for a few minutes
4. Go to the next question

To keep things flowing somewhat fast, we'll take a few questions on each section, and then there will be a much longer block Q&A at the end.

To be clear, don't be shy about discussion, it's why we sent the survey in the first place!

A Note on the Team

One thing that needs to be made clear before we go through results:

- Our team is small, perhaps smaller than you expect.
- Work is often slow as we need to ensure features all work together.
- “License fees” through RSICC fund background checks, not us.

As such, we can’t get to every request, and big changes will require collaboration.

Which MCNP particle types do you typically use?

Why was this question asked?

Expertise on some particles is limited, and we need to know where to expand.

Expected answers:

Neutrons/Photons > Electrons > Light Ions > Other

Which MCNP particle types do you typically use?

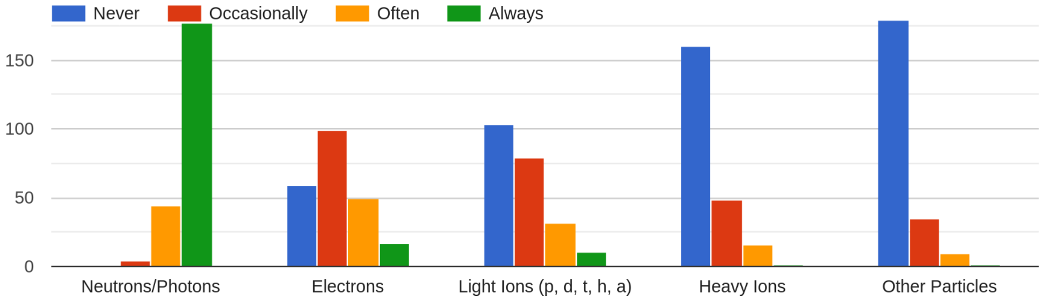

Selected Discussion

Thermal/Cold neutrons:

Improvements in $S(\alpha, \beta)$ in MCNP 6.3.0 and ENDF/B-VIII.1 will provide improved data in eV range. FIELD card improvements under investigation. No current team expertise on crystal lattice effects.

Improved physics models:

This would rely on collaboration between us and model developers, as well as an API – each new model makes quality assurance more difficult.

Energy deposition on mesh geometries:

Improvements planned for 6.4 – FMESH gaining most/all TMESH features.

Tabular data concerns:

Includes photonuclear concerns, heavy ion projectiles, collision-by-collision conservation of energy. These have been forwarded to XCP-5 (Nuclear Data Team).

Stopping powers:

Actively being researched.

Electron transport performance:

There are many places for improvements, but single-event transport and Landau straggling is expensive to model.

Discussion

What sort of simulations do you run most often?

Why was this question asked?

Outside the laboratory, we do not have a feel for who uses our code.

What sort of simulations do you run most often?

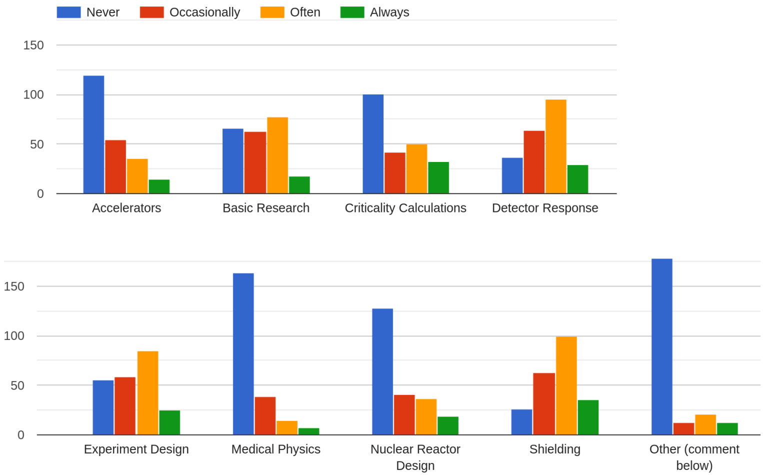

Simulation Type Results

- Shielding, detector response, and experiment design are unsurprisingly most common.
- Usage was surprisingly more even than we expected, with no abnormally uncommon work.

Other options listed:

- Activation and dosimetry analysis
- Well logging
- Fusion
- Microelectronics analysis

Discussion

Setup and Simulation Environment

Why was this question asked?

Sometimes, optimization on one operating system can harm optimization on another (particularly compiler choice). We need to know what will have the largest effect.

Expected answers:

Historically, macOS performance was given a much higher priority than other OSs.

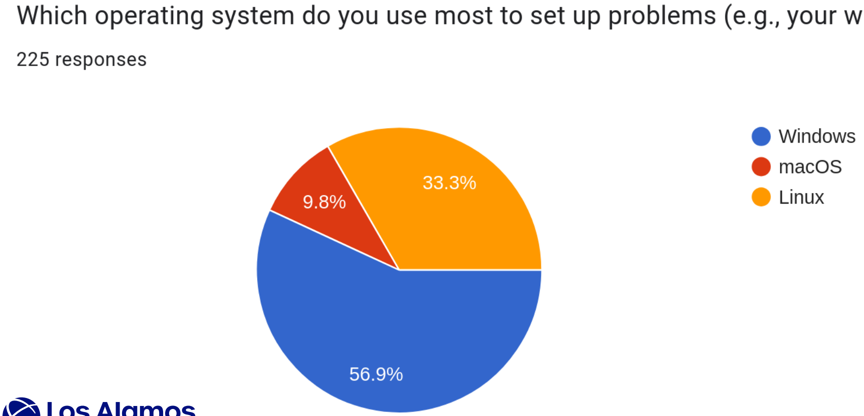

Which operating system do you use most to set up problems (e.g., your workstation)?

225 responses

Setup and Simulation Environment - 2

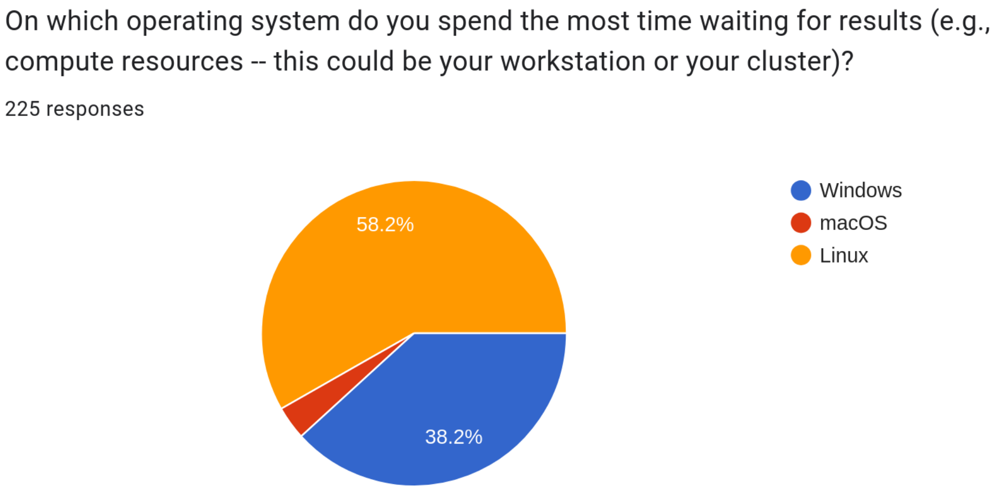

On which operating system do you spend the most time waiting for results (e.g., compute resources -- this could be your workstation or your cluster)?

225 responses

What are we going to do in response?

This indicates to us that Linux in particular needs more targeted optimization, followed by Windows.

Setup and Simulation Environment - 3

Which of the following do you run the MCNP code on?

225 responses

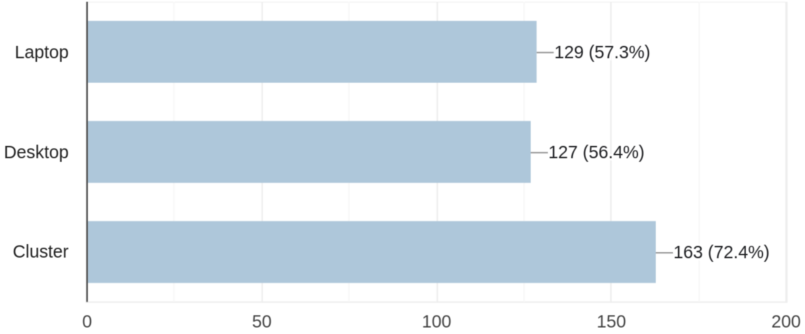

Setup and Simulation Environment - 4

If you were to run your problems on a single typical desktop, what timescales would your simulations usually take?

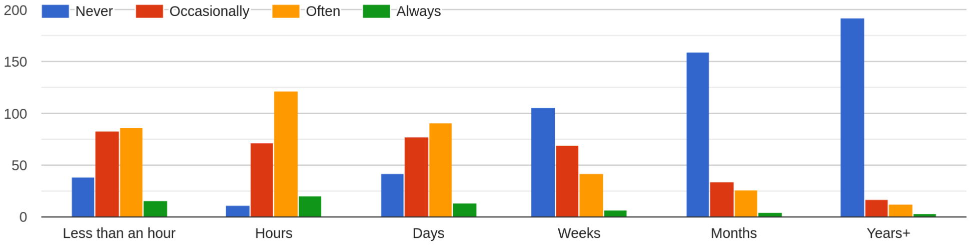

Discussion

Do you build the code?

Why was this question asked?

We wanted to know if there was anything we could do to reduce the number of people who need source code access or make the code more flexible.

Expected answers:

For: Mostly MPI with rare SOURCE / TALLYX users.

Against: No need, foreign national, SQA.

Survey confusion:

A surprising number of people aren't familiar with the terminology "code", and a few unfamiliar with "build".

Options we forgot:

Someone built it for me.

Building the Code: Results

Do you build the code yourself?

225 responses

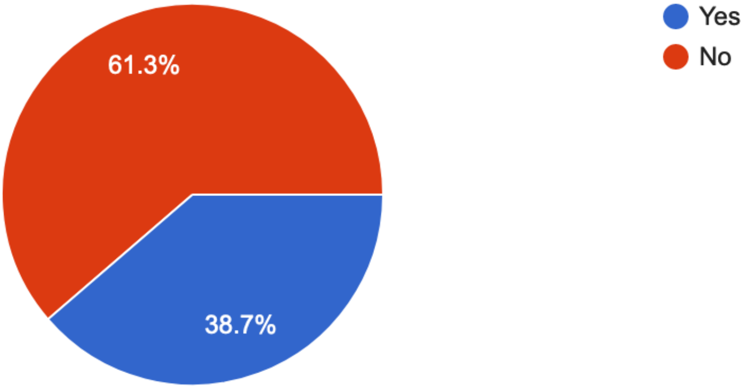

Why?

Why build the code (percent of detailed responses)?

1. Custom Patches (37%)
2. MPI (35%)
3. For fun / knowledge (18%)
4. Unspecified issues/bugfixes (16%)
5. TALLYX/SOURCE (14%)

Why not build the code?

1. Did not need to (56%)
2. Foreign National / No Access (23%)
3. Someone else built for me (16%)
4. SQA Requirements (10%)

What will we do?

MPI

For 6.3, we plan to release “best-effort” MPI builds on all OSs (MPICH and OpenMPI for Linux/macOS, Microsoft-MPI for Windows). Portability is going to be complicated, and this is our first attempt, so take care to read the directions! Expect to still need source for MPI.

TALLYX/SOURCE

For 6.4, we are creating an API to allow users to implement custom sources and tallies without full source access. The interfaces will be more limited in scope, because you will not have access to the entire memory of the code.

Bugfixes

Please let us know about bugfix patches, particularly if they’re applicable to 6.2 or (soon) 6.3! We often cannot use patches as-is, either due to the code moving onward, or the patch having side effects, though.

For other components, source *will* still need to be available.

Discussion

Which variance reduction methods do you use?

Why was this question asked?

Variance reduction greatly complicates the innards of the code. We need to know which ones to prioritize in cleanup.

Expected answers:

Expected almost all features used uniformly.

Variance Reduction Results

Which variance reduction methods do you use?

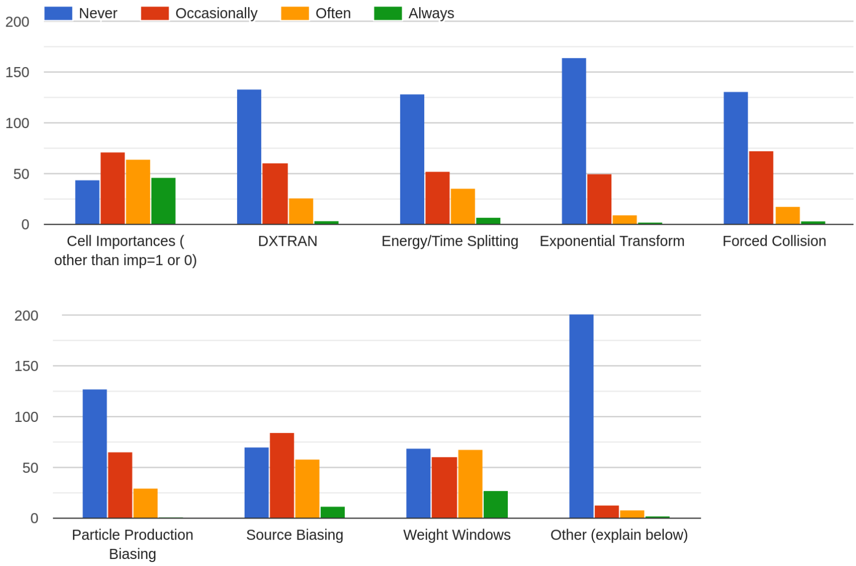

Variance Reduction Results

- Cell importances, weight windows, and source biasing most common
- Then DXTRAN, energy/time splitting
- Exponential transform probably least used variance reduction method
- Majority of “Other” answers was FW-CADIS/ADVANTG.
  - We are considering a PARTISN-driven CADIS implementation.

Discussion

What sort of tallies do you use?

Why was this question asked?

There are many tally structures within MCNP that either incompletely or completely duplicate one another's functionality.

Expected answers:

Expected F tallies to be used more frequently, mesh tallies less frequently

Options we forgot:

Did not add SSW, which was mentioned by a number of users and often used with postprocessing tools.

Tally Results

What sort of tallies do you use?

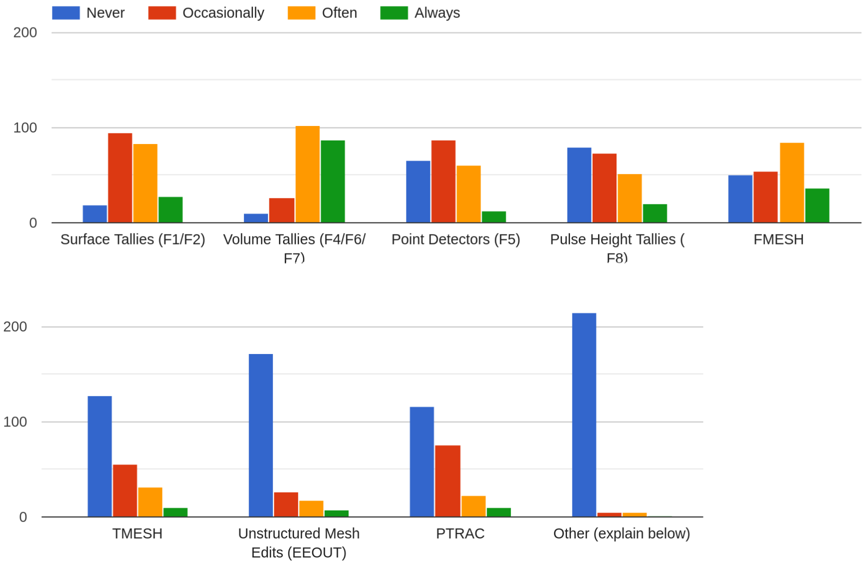

Discussion on Tallies

FMESH and TMESH

In 6.3, FMESH will have an XDMF+HDF5 output format that can be optionally enabled. This allows for visualization in Paraview and similar tools.

In 6.4, it is planned to add all TMESH features to FMESH, and mark TMESH as deprecated. This includes energy deposition, point detector effects, and spherical geometries. Those switching should get a significant performance boost.

EEOUT

In 6.3, EEOUT will also get an XDMF+HDF5 format.

PTRAC

In version 6.3, there is an optional HDF5 PTRAC file that is compatible with the new version of MCNPTools (now on GitHub!)

Discussion

Usage of Less Common Features

Why was this question asked?

For some, it was unclear if they were being used at all. Others, the external codes they link to may no longer be supported. Some features require significant attention, and knowing how many use it is the first step towards getting resources.

Expected answers:

Expected moderate usage of ACT, BURN, custom routines, FT, PERT, and nearly no usage of the others.

Usage of Less Common Features - Results

For the following features, please mark how often you use them:

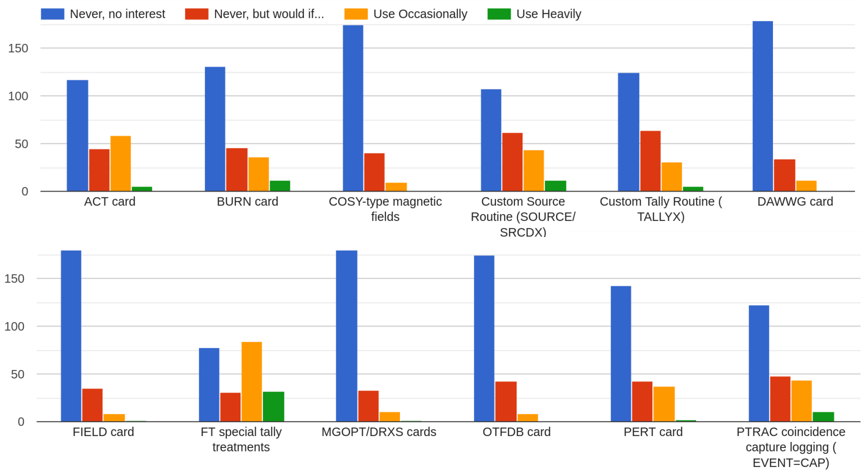

Less Common Features Discussion

- It is clear that ACT and BURN are both popular features.
  - A degree of distrust of the results for both
  - Many people want fixed-source BURN instead of just k-eigenvalue for things like accelerator-driven subcritical reactors
  - Both ACT and BURN are priorities for us, but are not actively funded and validation data is limited
- Very few people use COSY, FIELD, OTFDB, multigroup, or DAWWG for various reasons.
  - DAWWG no longer couples with modern PARTISN
  - OTFDB documentation was cited as being insufficient for use
- Many people comment that this is the first time they've heard of most of these features!

Discussion

Future Features

Why was this question asked?

The ideas listed are things we have been considering for a while, and wanted some sense on if users would like them. In addition, it helps us prioritize.

Which features would you most like to see in a speculative MCNP version 7.0?

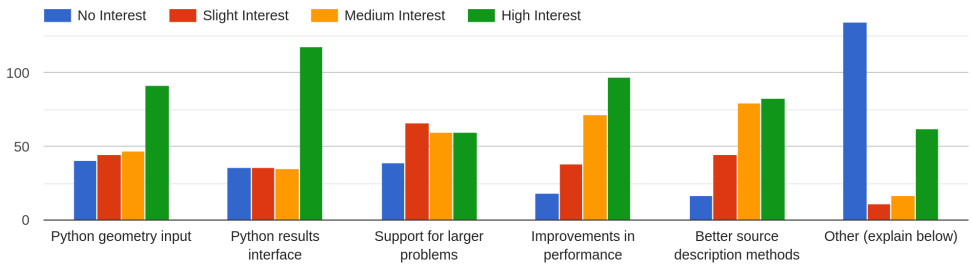

Feature Requests Discussion

- There seems to be enthusiasm for most options.
- Those that want a Python API were extremely excited.

Most common alternate entries (percent of detailed responses):

1. Improved visualization or a GUI (30%)
2. More input and output options, particularly with meshes (12%)
3. ADVANTG/FW-CADIS (6% – many more mentions in variance reduction block)
4. Improved documentation (6%)

Many other features with 1-2 responses. Some people asked for features we already had (such as the Intrinsic Source Generator).

What will we do?

Visualization

6.3 will contain a preview for our new plotter. It uses the same rendering engine, but the buttons and interface have been redone. There are future plans to improve the legibility of the rendered output.

Interactive GUI

An interactive GUI is exceedingly difficult with our resources, particularly if we wish to make it feature complete. There are preliminary thoughts to an API where you provide an incomplete geometry and it renders it, which could form the basis for a GUI in the future.

Mesh IO

6.3 will support an HDF5 file format as well as Abaqus. Due to the very large number of possible formats, we welcome conversion scripts.

Documentation

The documentation is always being improved, and many of the comments will be used to guide efforts.

Discussion

Difficulties with the Code

Why was this question asked?

We wanted to get a sense of if setup time / postprocessing effort exceeded runtime for most users.

Where do you find slowdowns in your MCNP workflow?

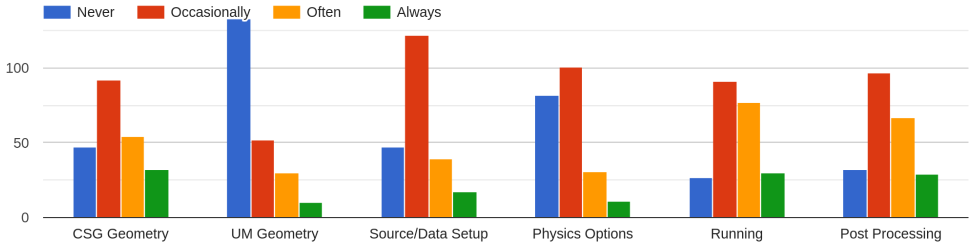

Difficulties Discussion

- It was generally noted that people’s time is more expensive than CPU time on all but the largest problems
- Many comments talked about postprocessing issues
- Code performance was only a concern in a few cases, mainly, lack of ability to run on clusters due to license issues, variance reduction that results in single histories that keep going, and occasional X11 issues.

What will we do with this knowledge?

We will continue to maintain or improve performance, but a big focus for next few versions will be the user interface on both input and output.

Discussion

Opinions on HDF5

Why was this question asked?

We are moving more and more capabilities to HDF5, to provide a “power user” interface to the code (particularly as problem sizes grow into the TBs). Text formats won’t go away, but it’s unclear the best way to provide access to them.

Expected Answers:

About an even split for and against.

We plan to output more results in an HDF5 format in future versions, with it eventually becoming the default in some places. What are your thoughts on this?

225 responses

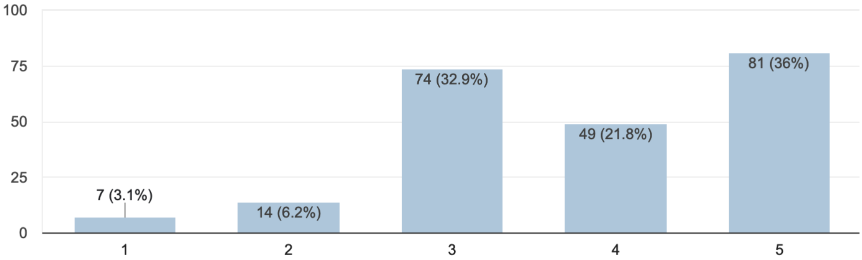

(1 is strongly negative, 5 is strongly positive)

Why not HDF5?

Concerns over classic outputs

Classic outputs will not go away, but we have considered making utilities you'd have to run to get text output.

Documentation

All current HDF5 formats are documented in the 6.3 manual (available on our site now). Let us know if there is sufficient detail. Also, files will be versioned going forward, to allow for better interfaces.

Concerns over data corruption

Implementation already designed to minimize corruption risk. Code would have to crash while writing a dump, not during transport. If further issues found, we can consider switching to SWMR mode.

Some users didn't like Paraview or Python

HDF5 is available in Fortran, C/C++, Julia, Rust, Matlab, etc...

Nobody who gave it "strongly negative" gave a reason, unfortunately.

Discussion

RSICC

There were a non-trivial number of complaints about RSICC, either knowingly or unknowingly:

- Request process is time consuming, and thus impractical for short-term use or student use.
- Cost:
  - Ballooning fees.
  - Repeat fees for updates.
  - Cluster admin license requirement can significantly multiply cost.
- Confusion over availability:
  - Why do users need to request a specific OS and not just get all of them?
  - Unclear if older versions are available.
  - People who could get it at one point no longer able to.
- Disk-based distribution is becoming unusable for many.
- Several people stating that they are going to switch to OpenMC/PHITS/GEANT4 *strictly due to licensing issues*.

Discussion