

Python for High Performance and Scientific Computing BOF Report

Abstract:

The Python for High Performance and Scientific Computing BOF is intended to provide current and potential Python users and tool providers in the high performance and scientific computing communities: a forum to talk about their current projects; ask questions of experts and leading vendors; explore methodologies; delve into issues with the language, modules, tools, and libraries; build community; discuss the concerns of the community with advocates and active developers present; and help continue the path forward.

Session Leaders:

William Scullin, Argonne National Laboratory
Andreas Schreiber, German Aerospace Center (DLR)
Andrew Terrel, Continuum Analytics

Outcomes:

The domain PyHPC.org was acquired and a web page for future coordination was created at: <http://pyhpc.org/>. A Google+ group was discussed and is yet to be set up while discussion continues as to PyHPC's status as a formal organization.

Materials from the BoF, workshop, and tutorials are all available at <http://pyhpc.org/>.

Following lightning talks, discussion produced a few action items:

- The need to address serialization of objects in Python 3 to encourage the use of Python with MPI. A PEP is in draft and should be submitted in the first quarter of 2014.
- The issue of file abstractions and I/O was re-raised and a number of solutions were discussed. As the community grows, it was noted that we should be noting regressions in performance in formal bug reports - even if rejected it keeps HPC specific problems on the minds of the maintainers.
- The reuse of elements of the Bohrium benchmarks, Synergia, NSF benchmarks, and GPAW as part of a common benchmark suite, though it may not be started until the second quarter of 2014