


Mochi Update



Phil



Matthieu



Kevin



Rob




George



Chuck




Pierre




Rob



Shane



Greg



Qing



Andrew




Bob




Brad



Galen



Dana



Jerome

Mochi: What are we trying to accomplish?

We're trying to transform HPC data services from a monoculture to an ecosystem.

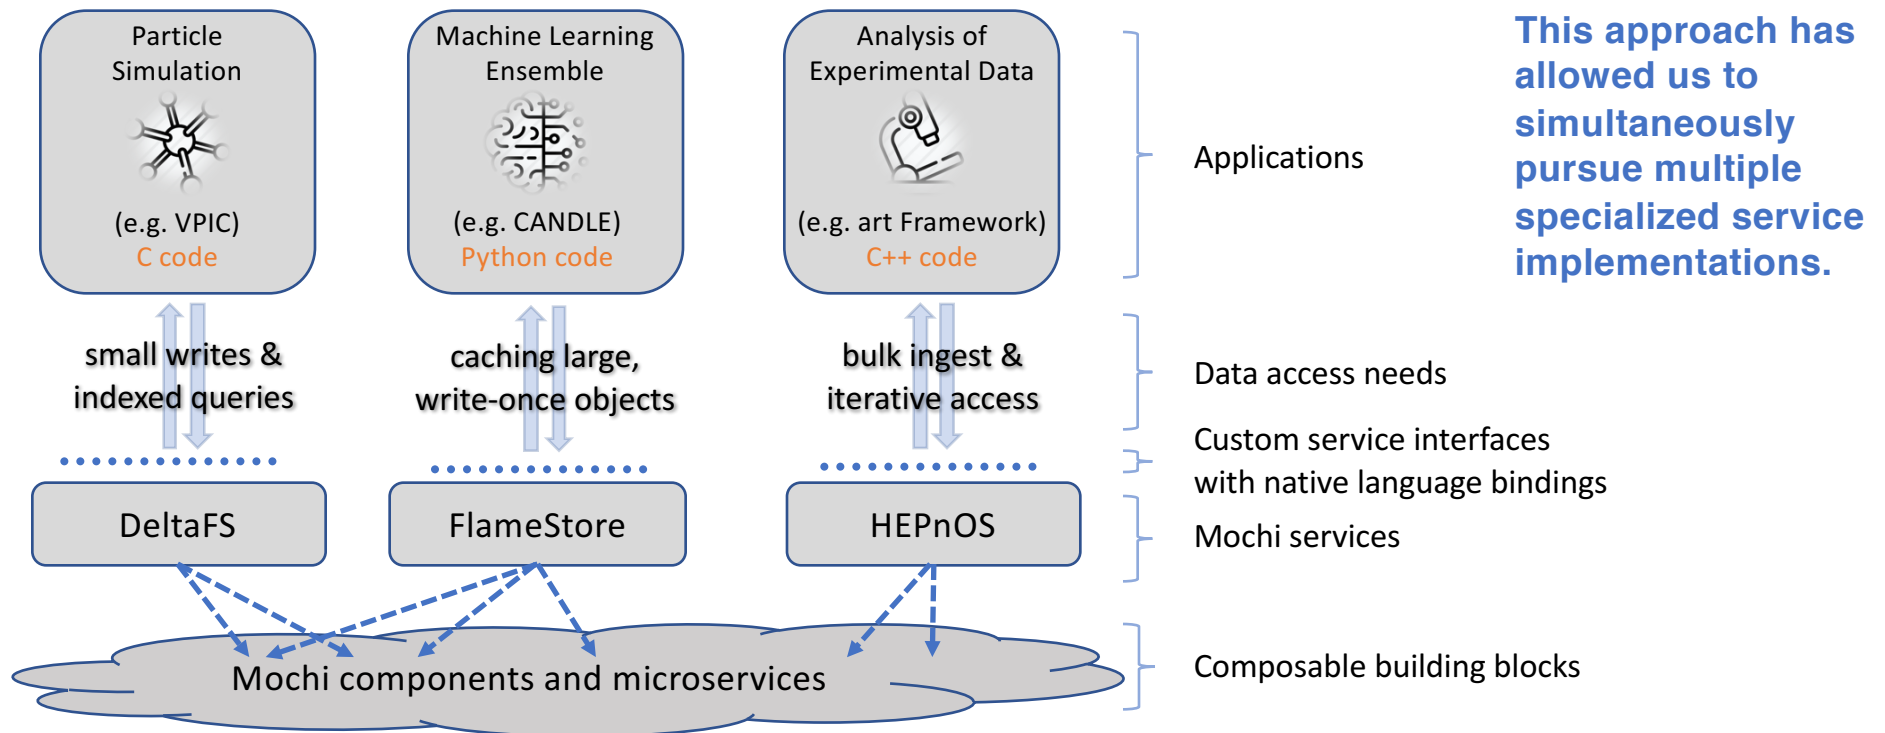
- Redefining how teams design and develop distributed services for use in HPC systems.
- Providing a portable "programming model" for these services.
- Providing a set of core building blocks.
- Demonstrating the methodology and tools with DOE science use cases.

We're trying to foster a community of service developers.

- Developing a set of training materials that will help others employ the tools.
- Making all these building blocks available to the larger community.

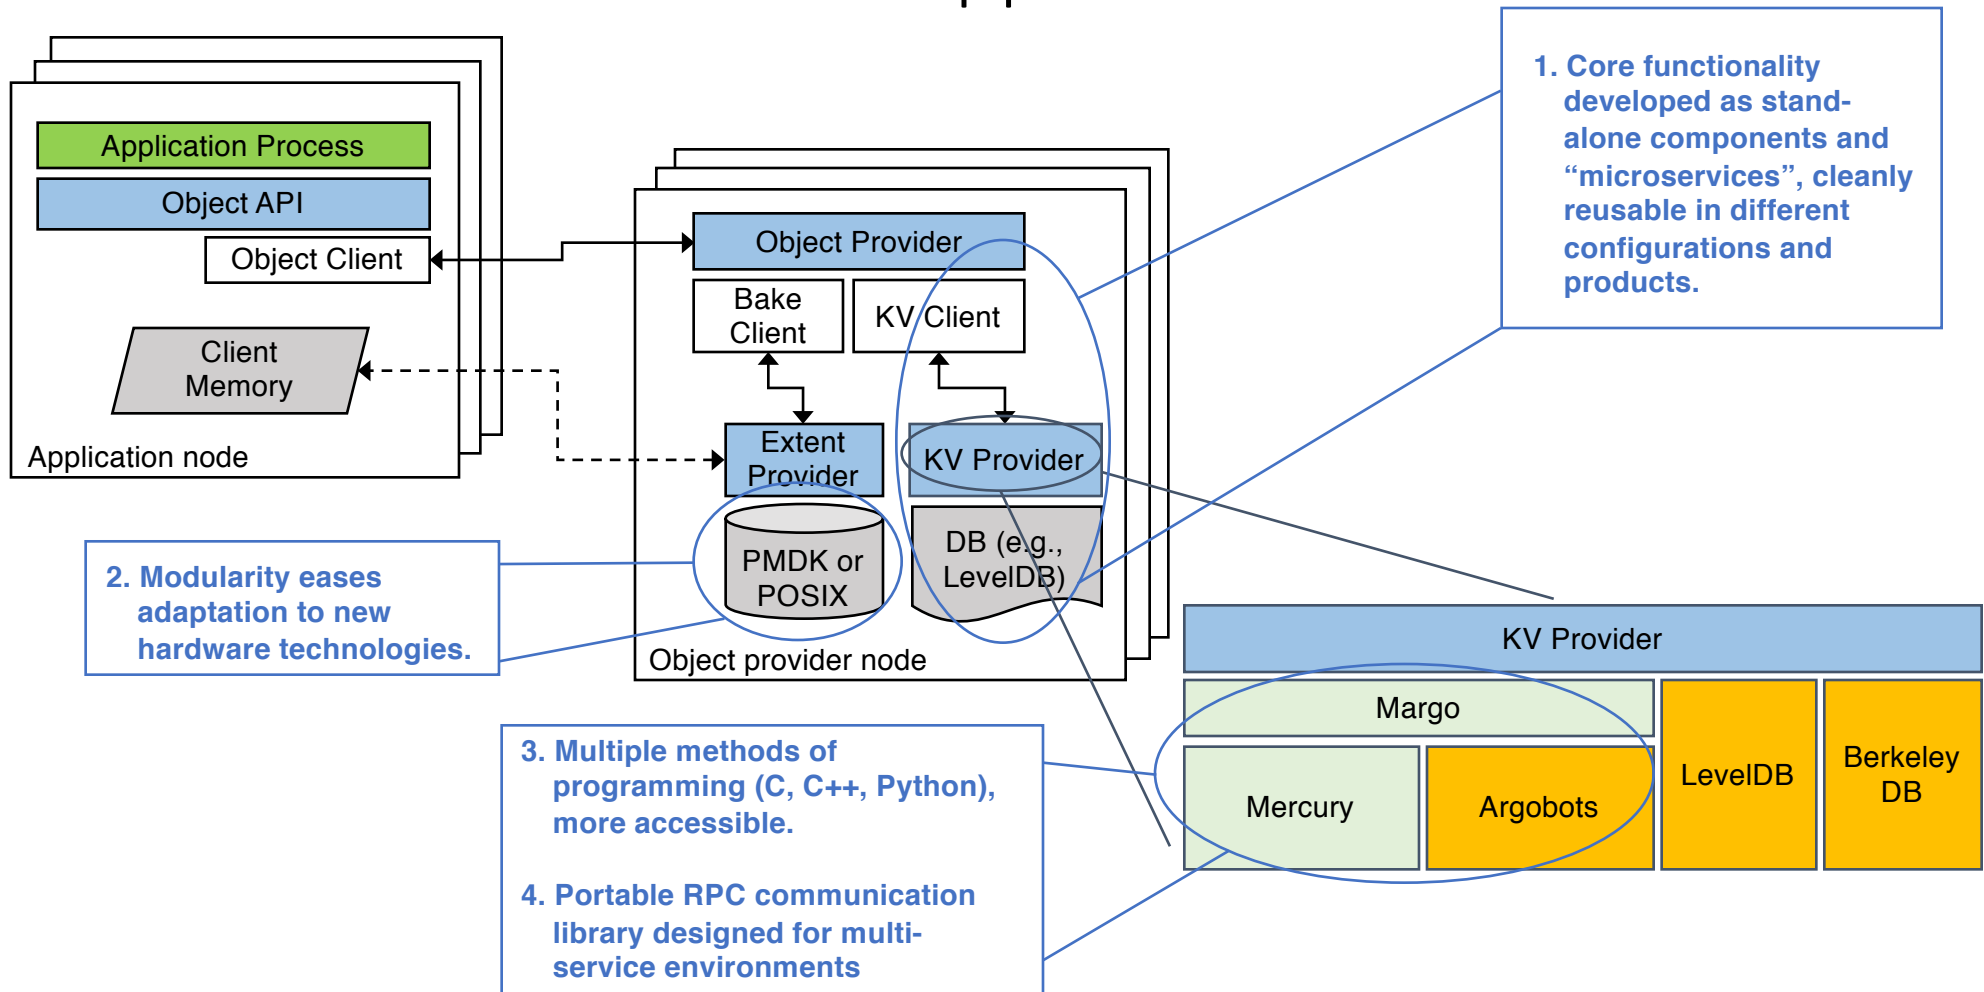
What's new in the Mochi approach?

An ecosystem of services co-existing and reusing functionality



Instead of "one size fits all", Mochi data services present tailored interfaces, semantics, and policies for data access while still leveraging robust building blocks.

What's new in the Mochi approach?



Accelerating HEP Data Analysis on HPC Platforms

Scientific Achievement

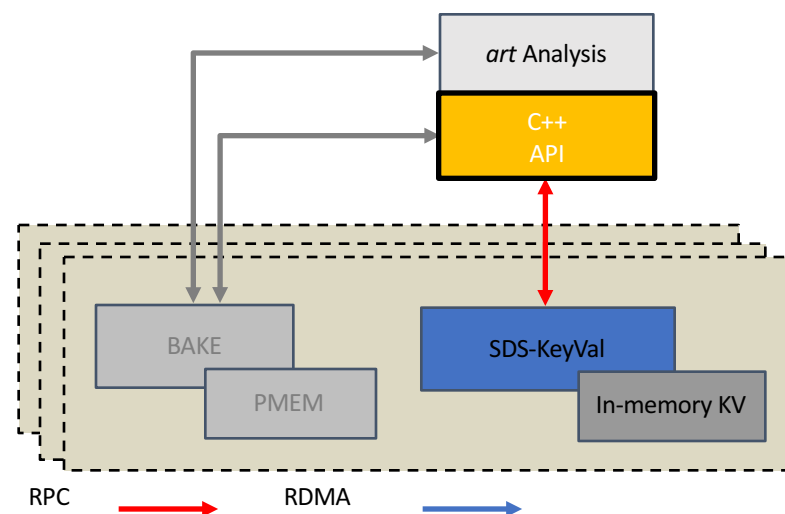
A custom data service for the HEP experiments, built to accelerate analysis on state-of-the-art HPC systems, with initial focus on NOvA neutrino experiment data.

Significance and Impact

Demonstrates HEP data analysis harnessing the ever-increasing power of ASCR (and other) supercomputers to extend the physics capability of HEP experiments.

Research Details

- First demonstrations of capability complete, using ALCF Theta system and approx. 100,000 events from NOvA
- HEPnOS data service hooks directly to *art* analysis tools
- Singularity used to package HEP analysis for execution on Theta
- Next steps focus on scaling to full analysis runs
- HEPnOS event model is amenable to use in other HEP analysis workflows (e.g., ICARUS, DUNE)



The ASCR Mochi project is researching methods for rapid specialization of data services for SC mission needs. In this case, these tools are being used to develop a low latency, in-system store (HEPnOS) for physics event data storage and analysis in order to accelerate HEP analysis workloads. First prototype targets a NOvA analysis workflow, and BAKE component has not been needed so far.

<https://xgitlab.cels.anl.gov/sds/HEPnOS/wikis/home>

Some Status...

- Major Theta performance issue debugged and addressed
- Initial small-scale runs completing
- ...

- Paper repo in place, some initial text and outline
<https://www.overleaf.com/9242663582wmznpwdbmsh>
- Target SC20?