Toward improved scientific software productivity on leadership facilities: An Argonne Leadership Computing Facility View

Lisa Childers

Venkatram Vishwanath, Thomas Uram, Lisa Childers, Hal Finkel, Jeff Hammond, Kalyan Kumaran, Paul Messina, Michael E. Papka

Argonne National Laboratory
<venkat, turam, childers, hfinkel, kumaran, messina, papka>@anl.gov

ASCR Workshop on Software Productivity for Extreme-Scale Science
January 2014

Very Brief ALCF Overview

- One half of the leadership class facility supported by the U.S. Department of Energy
- Home of Mira, 10-petaflops IBM Blue Gene/Q
  - 49,152 nodes / 786,432 cores
  - 786 TB of memory
  - 3,145,728 hardware threads
- Host to 40 INCITE, 18 ALCC, and 112 Director's Discretionary computational research projects*
- Home to 996 users*

*Informal tally as of 1/9/14

ALCF-Related Software

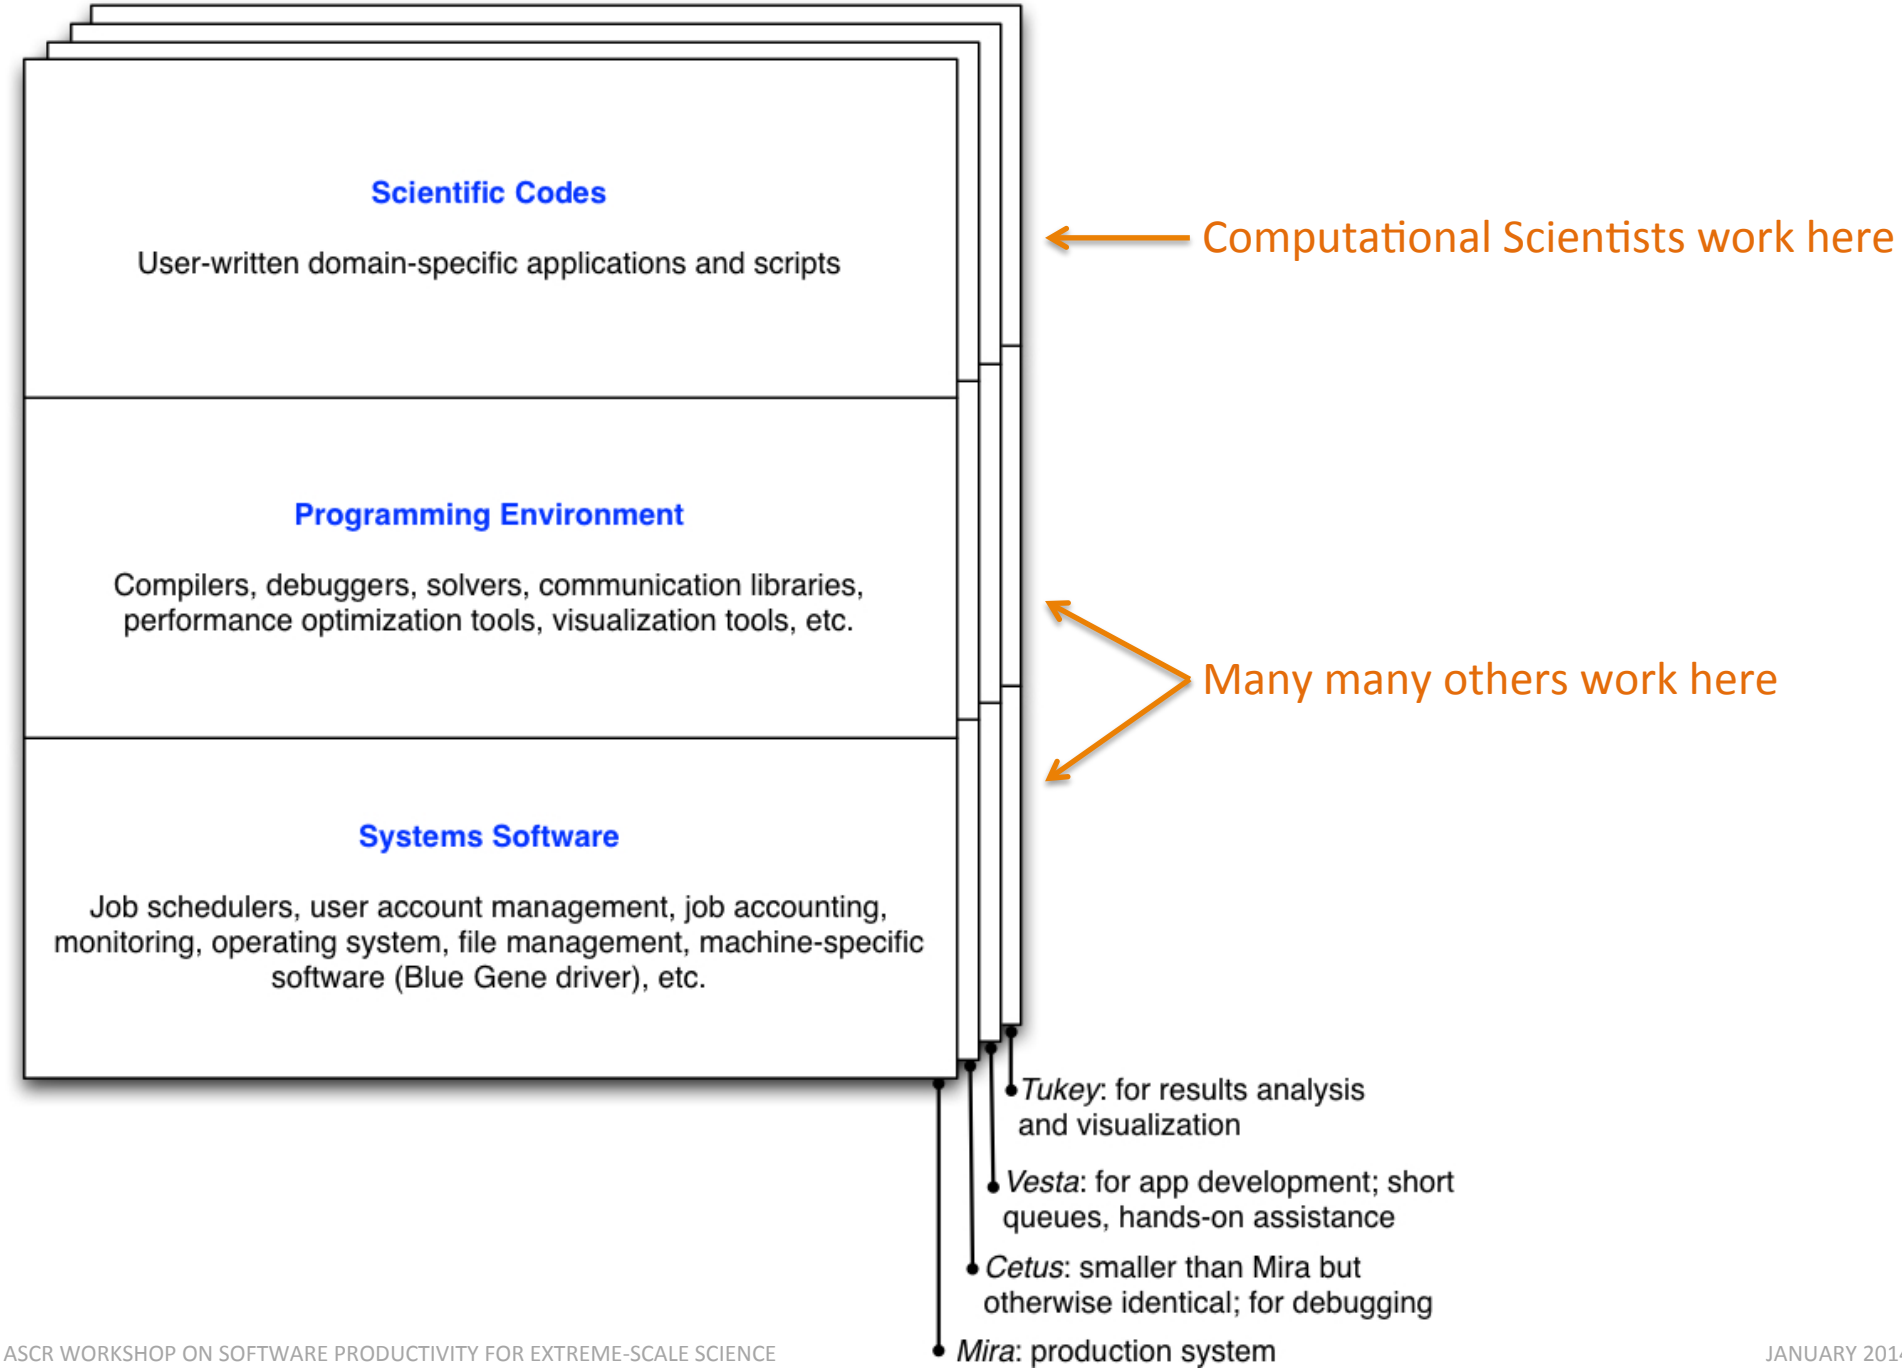

ALCF-Related Software

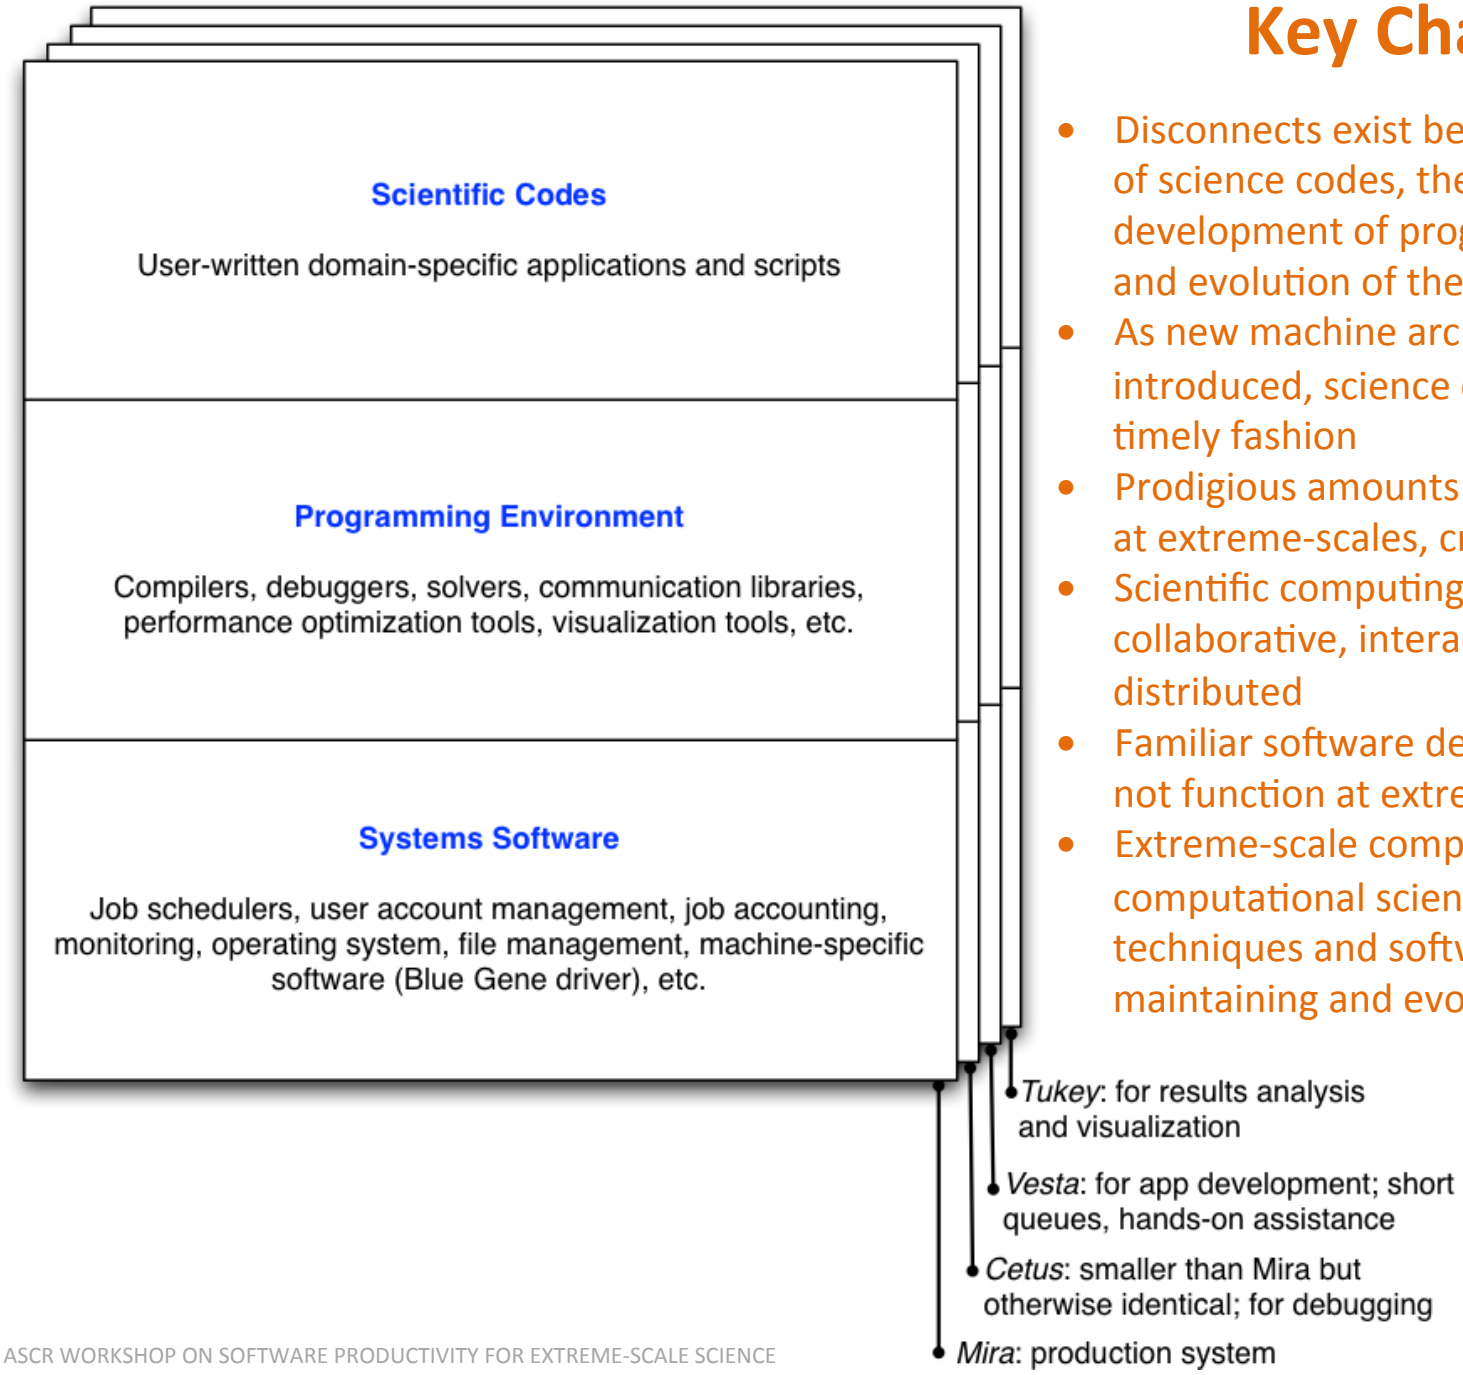

Key Challenges

- Disconnects exist between the development of science codes, the design and development of programming environments, and evolution of the underlying systems
- As new machine architectures are introduced, science codes must adapt in a timely fashion
- Prodigious amounts of data will be produced at extreme-scales, creating its own problem
- Scientific computing is increasingly collaborative, interactive, and geographically distributed
- Familiar software development tools may not function at extreme scale
- Extreme-scale computing may require the computational scientist to learn new coding techniques and software engineering for maintaining and evolving codes

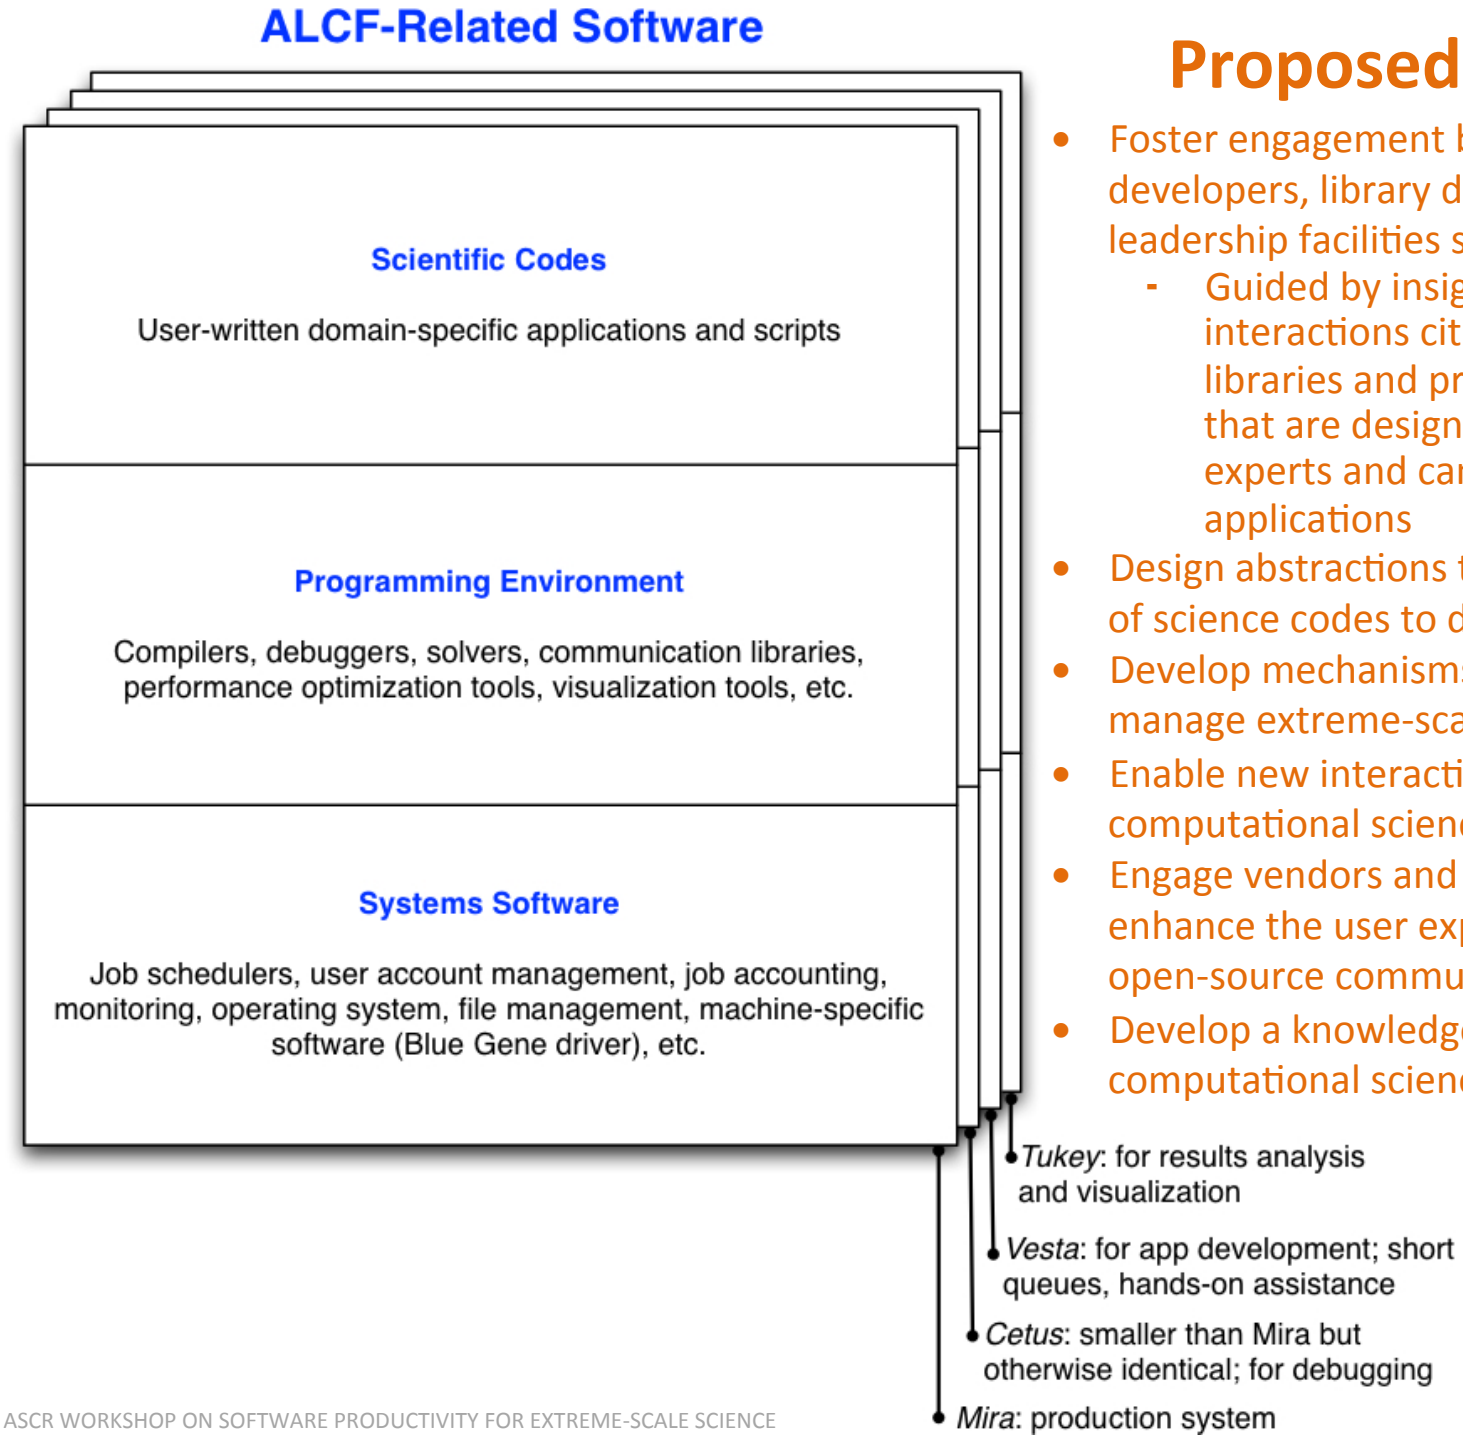

Proposed Approach

- Foster engagement between application developers, library developers, and DOE leadership facilities staff
  - Guided by insights learned through the interactions cited above, invest in libraries and programming frameworks that are designed and implemented by experts and can be used by many applications
- Design abstractions to facilitate adaptation of science codes to disruptive architectures
- Develop mechanisms that help scientists manage extreme-scale data
- Enable new interaction modes for computational science campaigns
- Engage vendors and researchers, and enhance the user experience, by leveraging open-source community software
- Develop a knowledge base for extreme-scale computational science