EXASCALE COMPUTING PROJECT

IDEAS productivity

better scientific software

So your code will see the future.

Enabling a Culture of Developer Productivity and Software Sustainability

Elaine M. Raybourn
Sandia National Laboratories
emraybo@sandia.gov, @elaineraybourn

SIAM CSE19 March 1, 2019

Outline

- The IDEAS-ECP team
- The context in which we work
- Why we believe in enabling a culture of productivity and sustainability
- How we do it
- What resources are available and where to find them
- How you too can get involved

IDEAS-ECP team

https://www.ideas-productivity.org

Michael Heroux (SNL), **Co-Lead PI**, **Director, Software Technology**
Lois Curfman McInnes (ANL), **Co-Lead PI**
David Bernholdt (ORNL), **Institutional PI, Outreach Lead**
Elsa Gonsiorowski (LLNL), **Institutional PI**
Osni Marques (LBNL), **Institutional PI, Webinars Lead**
David Moulton (LANL), **Institutional PI**
Boyana Norris (Univ of Oregon), **Institutional PI**
Elaine Raybourn (SNL) **Institutional PI, PSIP Lead**
Satish Balay (ANL)
Roscoe Bartlett (SNL)
Anshu Dubey (ANL)
Patricia Grubel (LANL)
Rinku Gupta (ANL), **BSSw Editor-in-Chief**
Stephen Hudson (ANL)
Reed Milewicz (SNL)
Mark Miller (LLNL)
Jared O'Neal (ANL)
Barry Smith (ANL)
Greg Watson (ORNL)
Jim Willenbring (SNL), **SDK Lead**
Paul Wolfenbarger (SNL)
Lisa Childers (ALCF)
Rebecca Hartman-Baker (NERSC)
Judy Hill (OLCF)
Hai Ah Nam (LANL), **BSSw Fellows**
Jean Shuler (LLNL)

Computing Facilities Liaisons

I help organizations tell their stories.

- I am a social scientist who deepens understanding by being embedded in different cultures.

- Cultures: DARPA, DoD (Army, SOF, USMC, OSD, Team Orlando), BT Research, FhG FIT, INRIA, DOE Office of Science, Academia, National Labs
- Research: innovation and productivity, immersive learning environments, design of transmedia learning ecosystems, cultural awareness
- Focus on ECP productivity since 2017, transmedia learning since 2010, games, immersive virtual environments, social simulations, and intelligent community systems since 2000
- Passion: ***Seize opportunities that allow us to learn about ourselves and others***
- Favorite question: ***Why not?***

What is the Exascale Computing Project (ECP)?

- As part of the National Strategic Computing initiative, ECP was established to accelerate delivery of a **capable exascale computing system** that integrates hardware and software capability to deliver approximately 50 to 100 times more performance than today's petaflop machines, within a similar size, cost, and power footprint.
- Oh, and did I mention by 2023?

ECP by the Numbers

A seven-year, $1.7 B R&D effort that launched in 2016

Six core DOE National Laboratories: Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Sandia

- Staff from most of the 17 DOE national laboratories take part in the project

Three technical focus areas (Application Development, Software Technology, Hardware and Integration)

More than 100 top-notch R&D teams

- Hundreds of consequential milestones delivered on schedule and within budget since project inception

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

ECP applications target six strategic areas

National security

Stockpile stewardship

Next-generation electromagnetics simulation of hostile environment and virtual flight testing for hypersonic re-entry vehicles

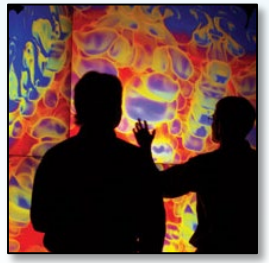

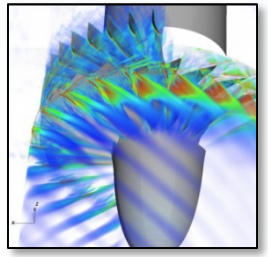

Energy security

Turbine wind plant efficiency

High-efficiency, low-emission combustion engine and gas turbine design

Materials design for extreme environments of nuclear fission and fusion reactors

Design and commercialization of Small Modular Reactors

Subsurface use for carbon capture, petroleum extraction, waste disposal

Scale-up of clean fossil fuel combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Urban planning

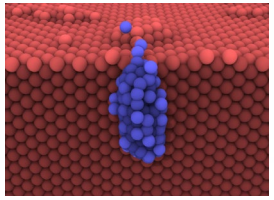

Scientific discovery

Find, predict, and control materials and properties

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Demystify origin of chemical elements

Light source-enabled analysis of protein and molecular structure and design

Whole-device model of magnetically confined fusion plasmas

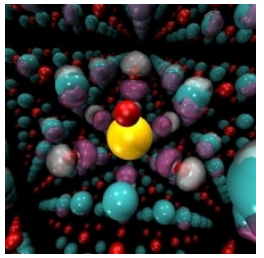

Earth system

Accurate regional impact assessments in Earth system models

Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols

Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

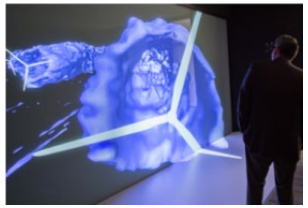

Health care

Accelerate and translate cancer research

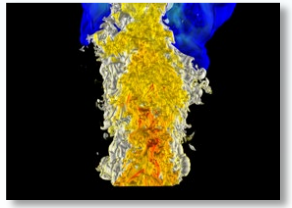

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Many ECP ST products are available (https://e4s.io)

For example…

Programming Models and Runtimes Produc

Legion	http://legion.stanford.edu
ROSE	https://github.com/rose-compiler
Kokkos	https://github.com/kokkos
DARMA	https://github.com/darma-tasking
Global Arrays	http://hpc.pnl.gov/globalarrays/
RAJA	https://github.com/LLNL/RAJA
CHAI	https://github.com/LLNL/CHAI
Umpire	
MPICH	
PaRSEC	
Open MPI	
Intel GEOPM	
LLVM OpenMP compiler	
OpenMP V&V Suite	
BOLT	
UPC++	
GASNet-EX	
Qthreads	

Development Tools (19)

SICM	https://confluence.exascaleproject.org/display/STSS07
QUO	https://github.com/lanl/libquo
Kitsune	https://github.com/lanl/kitsune
SCR	https://github.com/llnl/scr
Caliper	https://github.com/llnl/caliper
mpiFileUtils	https://github.com/hpc/mpifileutils
Gotcha	http://github.com/llnl/gotcha
TriBITS	https://tribits.org
Exascale Code Geneneration Toolkit	
PAPI	http://icl.utk.edu/exa-papi/
CHiLL Autotuning Compiler	
Search using	
	hpctoolkit.org
	www.paradyn.org
	www.cs.uoregon.edu/research/tau
	ft.ornl.gov/research/papyrus
	ft.ornl.gov/research/openarc
	rm.org/
	www.cs.uoregon.edu/research/pdt/home.php

Mathematical Libraries Products (16)

xSDK	https://xsdk.info
hypre	http://www.llnl.gov/casc/hypre
FleCSI	http://www.flecsi.org
MFEM	http://mfem.org/
Kokkoskernels	https://github.com/kokkos/kokkos-kernels/
Trilinos	https://github.com/trilinos/Trilinos
SUNDIALS	https://computation.llnl.gov/projects/sundials
PETSc/TAO	http://www.mcs.anl.gov/petsc
libEnsemble	https://github.com/Libensemble/libensemble
STRUMPACK	http://portal.nersc.gov/project/sparse/strumpack/
SuperLU	http://crd-legacy.lbl.gov/~xiaoye/SuperLU/
ForTrilinos	https://trilinos.github.io/ForTrilinos/
SLATE	http://icl.utk.edu/slate/
MAGMA-sparse	https://bitbucket.org/icl/magma
DTK	https://github.com/ORNL-CEES/DataTransferKit
Tasmanian	http://tasmanian.ornl.gov/

etc…

To achieve capable exascale requires a holistic approach

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable and productive software stack	Hardware technology elements	Integrated exascale supercomputers

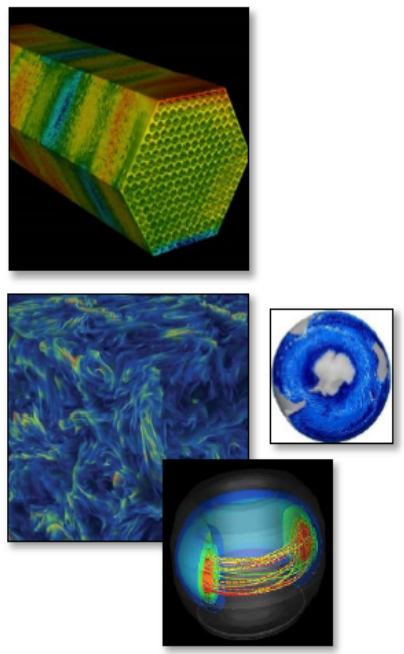

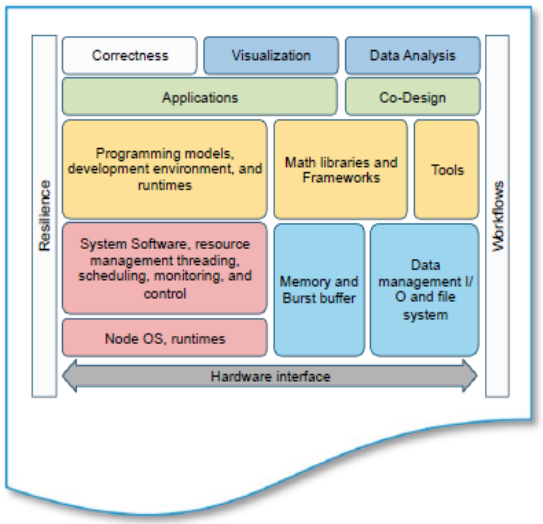

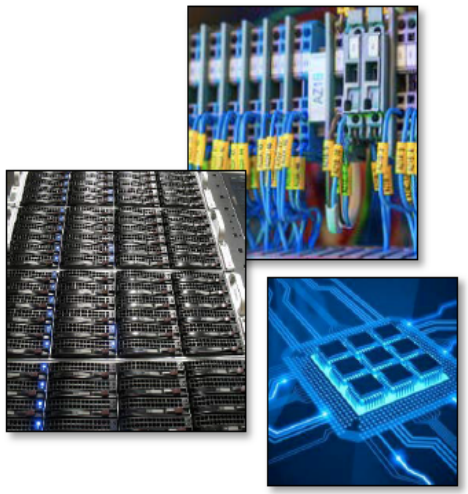

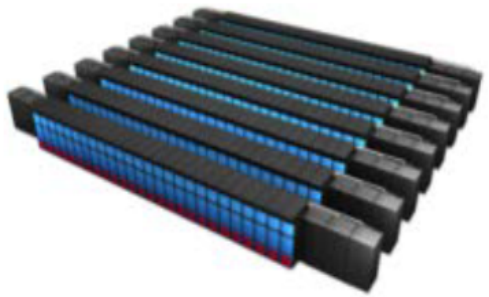

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Why we believe in enabling a culture of developer productivity and software sustainability

Science through computing is only as good as the software that produces it.

CSE Technical Challenges

- All parts of the ecosystem can be under under research
- Requirements change throughout the lifecycle as knowledge grows
- Importance of reproducibility, sustainability
- Verification is complicated
- Real world is messy, so is the software

CSE Social Challenges

- Competing priorities and incentives
- Limited resources
- Perception of “invisible work” with deferred or no benefit
- Need for interdisciplinary interactions
- Boutique operations must scale!

IDEAS-ECP Project

https://www.ideas-productivity.org, https://bssw.io

Co-PIs: Lois McInnes (ANL) & Michael Heroux (SNL)

Goal: Improve Exascale Computing Project (ECP) developer productivity and software sustainability while ensuring continued scientific success.

1 Interviews with Exascale Computing teams

- Applications & Software Technology
- Understand crosscutting productivity challenges, priorities, and opportunities

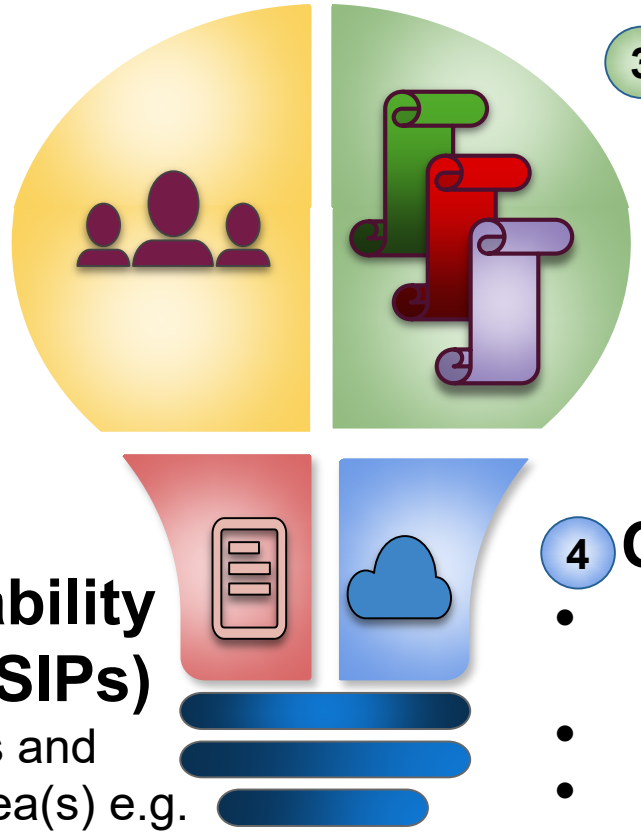

2 Productivity and Sustainability Improvement Planning (PSIPs)

- Work with team to define focus and track progress on particular area(s) e.g. research software engineering

3 Customize, create, and curate methodologies

- Targeting application productivity and sustainability
- Create user stories to convey requirements from interview & PSIPs to determine priorities, plans for work

4 Outreach and training

- In partnership with US Department of Energy facilities
- Documents: WhatIs, HowTo, PSIP policies
- Webinar series and tutorials
- Better Scientific Software site (https://bssw.io)

Interactions with ECP teams have illuminated needs

Process for interviews, synthesis, outreach

ECP Application, Co-Design, and Software Teams:

CANDLE, ExaGraph, Exascale MPI, ExaStar, E3SM-MMF, EXAALT, MARBL, NWChemEx, UnifyCR, QMCPack, and WDMApp

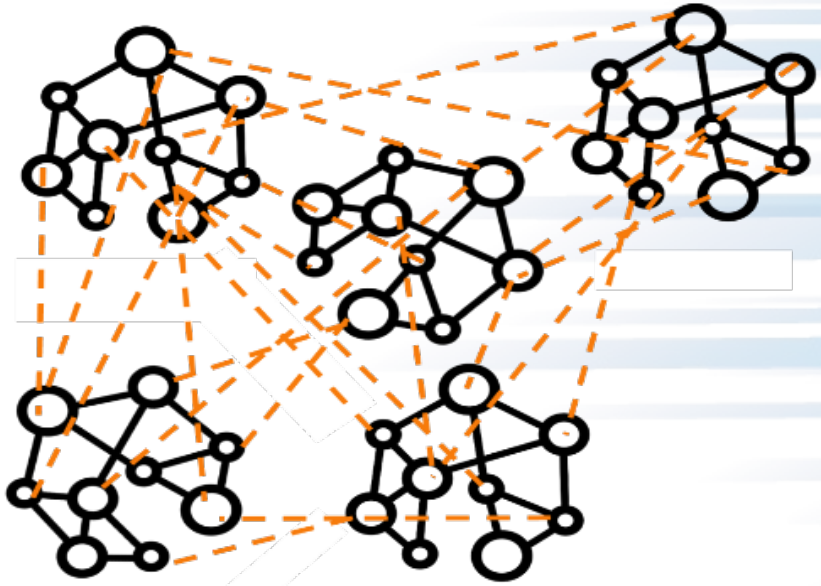

‘teams of teams”

High-priority needs for ECP teams:

Tools that foster productive and sustainable collaboration (through software) for ***aggregate*** ECP science teams

- testing/verification of scientific software
- team onboarding and team member transitions
- intermediate/advanced Git (especially for aggregate teams)
- code reviews for identifying defects
- agile team management,
- agile workflows for scientific software
- use of (interoperable) scientific libraries

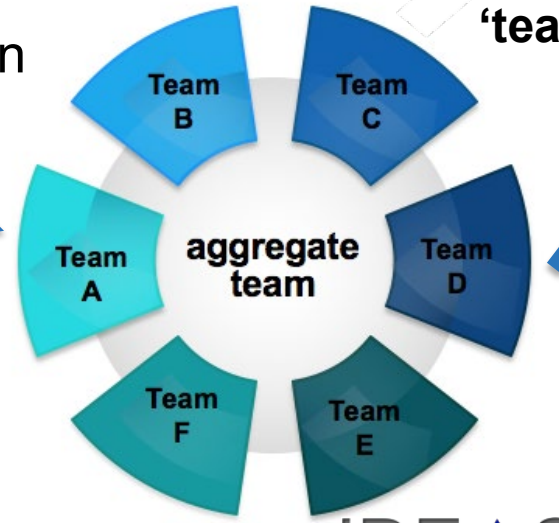

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Productivity and Sustainability Improvement Planning (PSIP) Examples: EXAALT & MPICH

PSIP workflow helps a team create user stories, identify areas for improvement, select a specific area and topic for a single improvement cycle, and then develop those improvements with specific metrics for success.

MPICH PSIP: Onboarding new team members

Practice: Create Centralized Training Resources		
Score (0 – 4)	Description	Tracking
0	Initial Status : No training process in place.	
1	Understand MPICH requirement for developers and typical challenges for new hires	✓
2	Review and gather specific training materials	✓
3	Design "MPICH Training Base" website	✓
4	Solicit feedback, improve, add and prune content to ensure effectiveness	2019

EXAALT PSIP: Continuous integration (CI) testing

BSSw blog article: Adopting Continuous Integration for Long Timescale Materials Simulation, Rick Zamora (Sept 2018)

PSIP Process: Continuous Integration (CI)

Target: Implement and document a basic CI pipeline to act as the foundation for automated build and functionality testing.

0. **Initial Status.** No comprehensive CI framework in place	
1. Develop a minimal docker image, with EXAALT dependencies	✓
2. Implement a minimal 'yml' script for the CI pileine	✓
3. Update EXAALT docker image to leverage CMake, and create a ParSplice-specific image for build testing	✓
4. Generate step-by-step "how-to" Docker-image documentation	✓
5. Extend CI to automate build and functionality testing with both CMake and Boost.	✓

Score (0-5): 4

PSIP Process: Testing

Target: Implement and document practical testing examples for ongoing EXAALT development.

0. **Initial Status.** No comprehensive testing framework in place	✓
1. Add 1–3 example tests using the existing CMake infra-structure (CTest)	✓
2. Add 1–3 example tests using the 'Boost Test' library	✓
3. Integrate the CTest infrastructure with the new Boost tests	✓
4. Integrate the Boost-enabled CTest framework into the CI pipeline	✓
5. **Bonus:** Work with EXAALT team to add more advanced tests to improve code coverage	

Score (0-5): 3

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Workflow for Best Practices Content Development

Modern learning theory:

Build from knowledge base: Elaboration and models

Vast body of SE content from broad community

Learn, adapt, adopt, assimilate

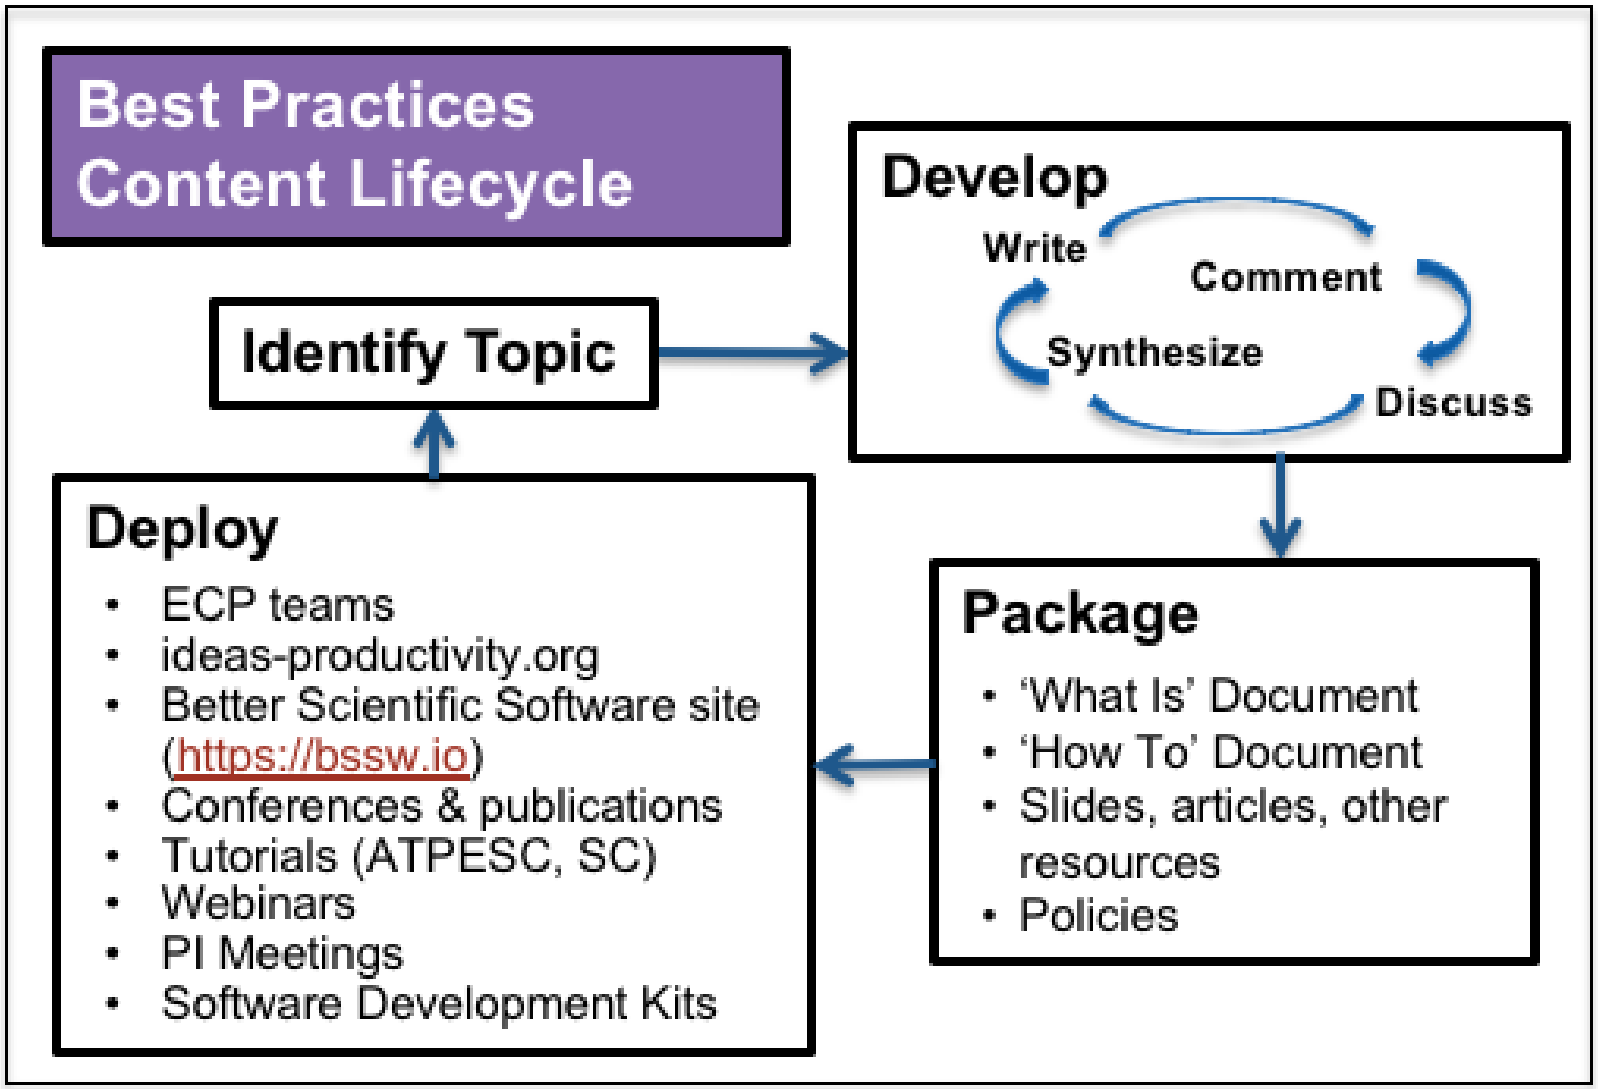

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Enabling a culture of productivity and sustainability by involving others

better scientific software

BSSw Fellowship Program

Recognition & funding to leaders and advocates of high-quality scientific software

Class of 2019

Class of 2018

We will begin accepting applications for the 2020 BSSw Fellowship Program in September, 2019

We are looking for people who are:

- Passionate about scientific software.
- Interested in contributing powerful ideas, tools, methodologies, and more that improve the quality of scientific software.
- Able to use the fellowship to broadly benefit the scientific software community.
- Willing to participate as an alum in subsequent years to guide selection of future fellows and promote better scientific software in their community.

https://bssw.io *So your code will see the future*

Subscribe to BSSw mailing list to be notified about BSSw Fellowship for 2020

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

HPC Best Practices Webinar Series

- 2017-06-07 Python in HPC, Rollin Thomas (NERSC), William Scullin (ALCF) and Matt Belhorn (OLCF)
- 2017-09-13 Barely Sufficient Project Management: A few techniques for improving your scientific software development, Mike Heroux (SNL)
- 2017-11-01 Managing Defects in HPC Software Development, Tom Evans (ORNL)
- 2018-01-17 Bringing Best Practices to a Long-Lived Production Code Charles Ferenbaugh (LANL)
- 2018-04-18 Software Citation Today and Tomorrow, Daniel Katz (NCSA and UIUC)
- 2018-06-13 Popper: Creating Reproducible Computational and Data Science Experimentation Pipelines, Ivo Jimenez (UCSC)
- 2018-07-18 How Open Source Software Supports the Largest Computers on the Planet, Ian Lee (LLNL)
- 2018-08-21 Software Sustainability: Lessons learned from different disciplines, Neil Chue Hong (Software Sustainability Institute, UK)
- 2018-09-19 Modern CMake, Bill Hoffman (Kitware)

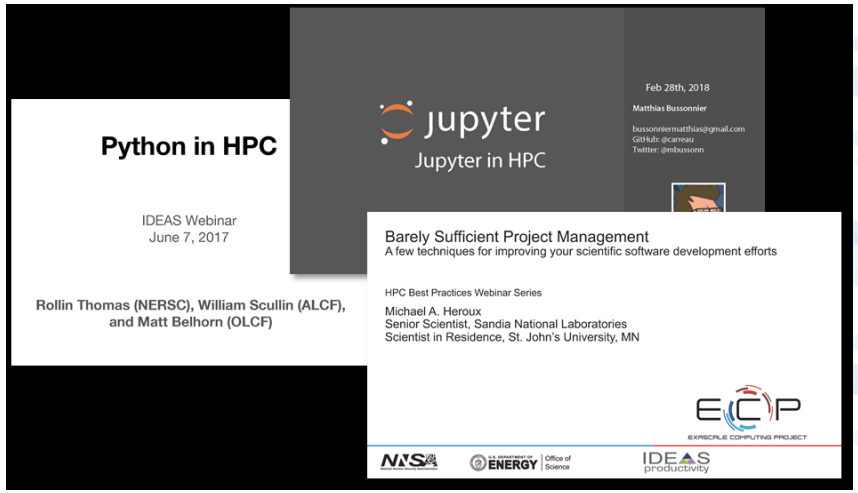

https://www.ideas-productivity.org

IDEAS productivity ECP EXASCALE COMPUTING PROJECT

better scientific software

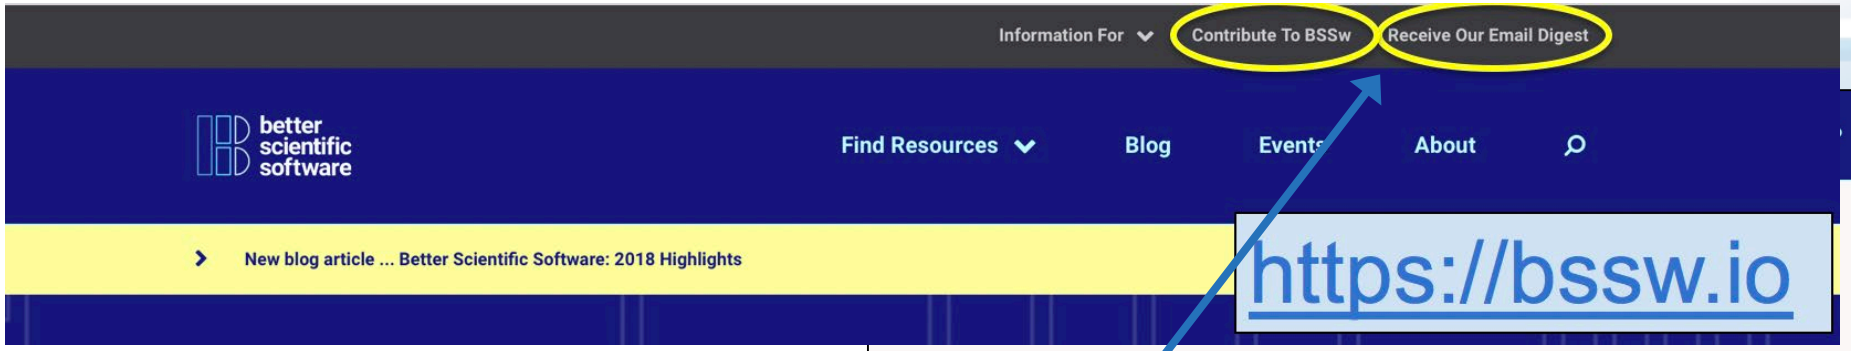

What is BSSw?

Community-based resource for sharing information on practices, techniques, and tools to improve developer productivity and software sustainability for computational science and engineering.

We want and *need* contributions from the community … Join us!

- **Types of content**
  - Informative articles
  - Curated links
    - Highlight other web-based content
  - Events
  - WhatIs, HowTo docs
  - Blog articles

Receive our email digest

Many ECP contributors

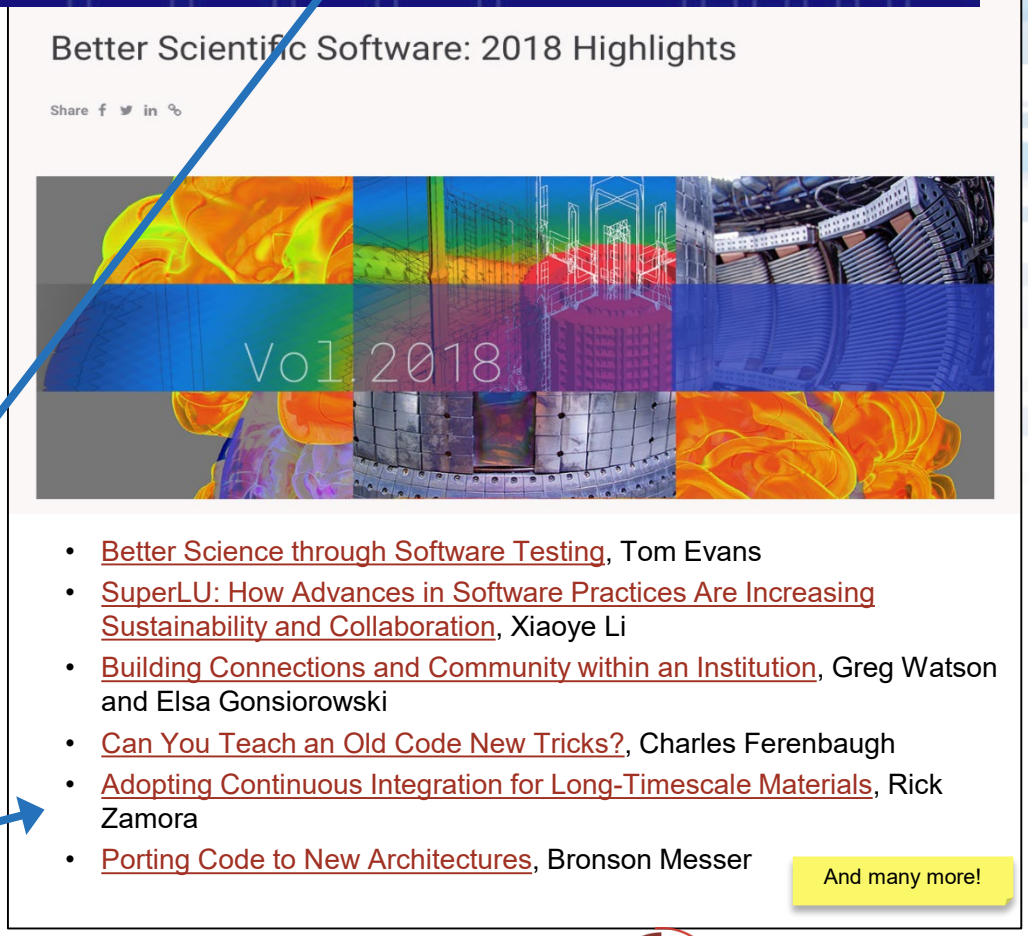

If you like our resources, tell us!

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Science through computing is only as good as the software that produces it.

License, citation and acknowledgements

License and Citation

- Requested citation: Raybourn, E.M. Enabling a Culture of Developer Productivity and Software Sustainability. 2019 SIAM Conference on Computational Science and Engineering, Spokane, WA, March 1, 2019. DOI: https://doi.org/10.6084/m9.figshare.7789604. SAND2019-2224 C.

Acknowledgements

- *Special thanks to the members of IDEAS-ECP.*
- This work was supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (ASCR), and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.
- Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. Images used by permission.