EXASCALE COMPUTING PROJECT

ECP-I-PO-RPT_2020_00003

Executive Summary Report
2020 Exascale Computing Project Annual Meeting

Marta García and Todd Munson
Argonne National Laboratory

May 2020

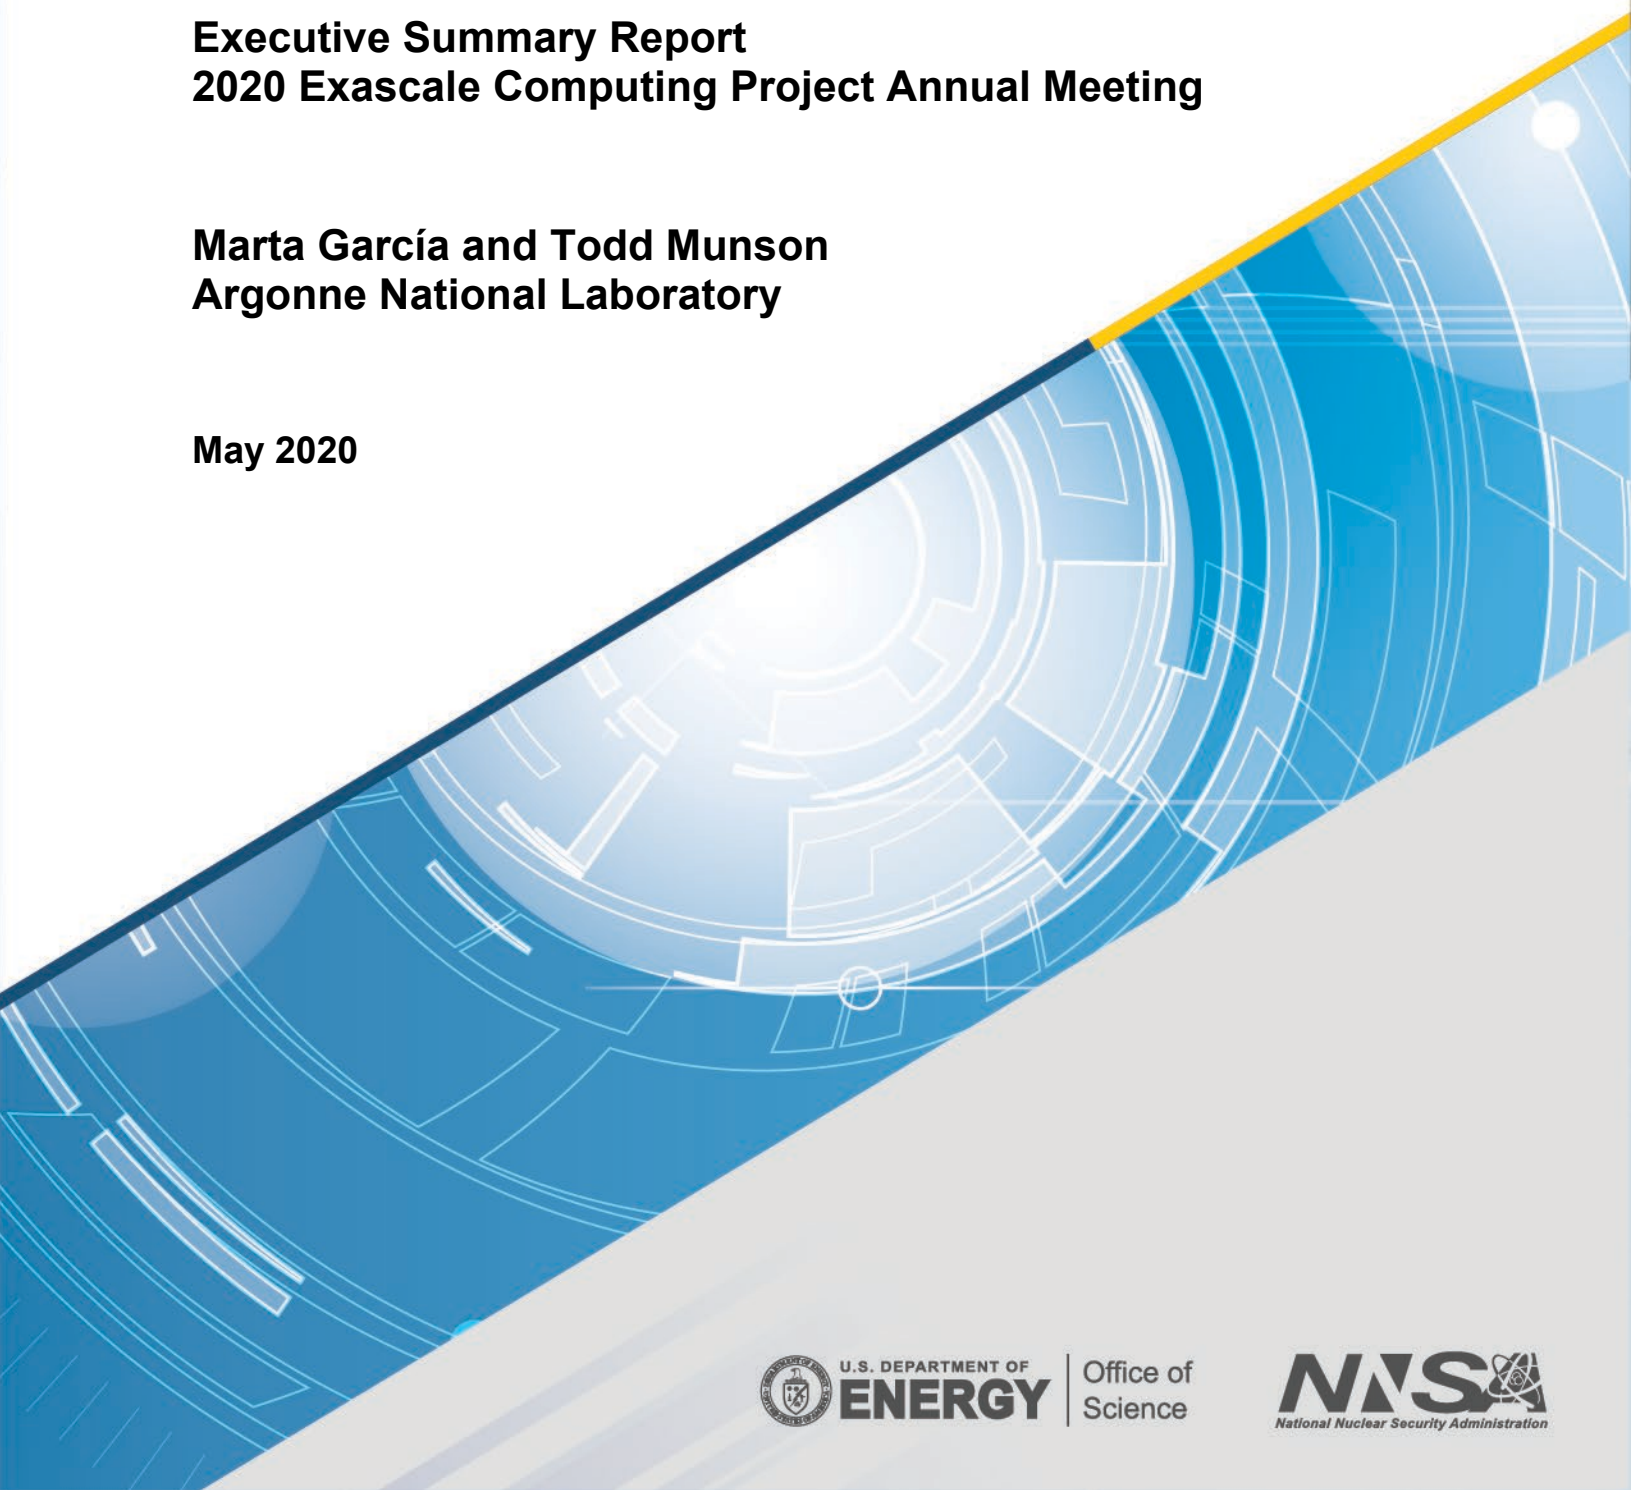

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website www.osti.gov
Website http://www.osti.gov/scitech/

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website http://classic.ntis.gov/

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website http://www.osti.gov/contact.html

ECP-I-PO-RPT_2020_00003

Executive Summary Report
2020 Exascale Computing Project Annual Meeting

Office of Advanced Scientific Computing Research
Office of Science
US Department of Energy

Office of Advanced Simulation and Computing
National Nuclear Security Administration
US Department of Energy

May 2020

Executive Summary Report
2020 Exascale Computing Project Summary Report

APPROVALS

Submitted by:

DocuSigned by:
Marta G.
04AC29C4D0B6427...

DocuSigned by:
Todd Munson
8F8E4B2009C9433...

5/14/2020 5/14/2020

Marta García and Todd Munson
Annual Meeting Program Co-Chairs
Argonne National Laboratory

Date

Concurrence:

DocuSigned by:
Doug Collins
7CE3FF7BD3964B3...

5/14/2020

Doug Collins
Annual Meeting Logistics Chair
Exascale Computing Project

Date

Approval:

DocuSigned by:
Douglas B. Kothe
22832EC21F69437...

5/15/2020

Doug Kothe
Exascale Computing Project Director

Date

CONTENTS

EXECUTIVE SUMMARY

The Exascale Computing Project (ECP) delivers specific applications, software products, and outcomes on DOE computing facilities[1]. Integration across these elements for specific hardware technologies for exascale[2] system instantiations is fundamental to ECP success. The outcome of the ECP is the delivery of a capable exascale computing ecosystem to provide breakthrough solutions addressing our most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security. This outcome is not a matter of ensuring more powerful computing systems. The ECP is designed to create more valuable and rapid insights from a wide variety of applications ("capable"), which requires a much higher level of inherent efficacy in all methods, software tools, and ECP-enabled computing technologies to be acquired by DOE laboratories ("ecosystem").

The ECP annual meeting provides a unique opportunity for the core technical expertise in the United States focused on achieving this next plateau of computational science and computing performance to engage in direct discussions on project execution. Face-to-face gatherings in technical communities like this are common and needed for the exchange of scientific ideas and technical performance. The ECP annual meeting stands apart from other technical conferences and meetings in the computing community as it is uniquely and solely focused on the execution of the ECP and the integration of technical activities leading to the creation of the exascale computing ecosystem for the future. The direct interaction of key critical technical staff, who are leaders in their respective fields, and the resulting give-and-take between software, applications, and hardware and the technical co-design therein, is unique and essential to the effective execution of the ECP.

The first annual meeting was held in Knoxville, Tennessee, January 31 – February 2, 2017 and brought together, for the first time, a diverse collection of researchers from 16 DOE national laboratories as well as university computer and computational science researchers to discuss shared problems and joint solutions for the development of a capable exascale computing ecosystem. These interactions resulted in focused technical plans and an energized community centered on advances for ECP.

The second annual meeting was held in Knoxville, Tennessee, February 5–9, 2018. It included 643 individual thought leaders and performers in application development, software research and deployment, and hardware research and integrators, all of whom are part of the multifaceted, billion dollar HPC community. This meeting provided a platform to discuss and disseminate numerous examples where researchers with common goals and synergistic solutions came together for the first time to deliver tangible results. Additionally, at the 2018 meeting, ECP researchers had the opportunity to digest all US HPC vendor R&D product roadmaps pointing to exascale – not only to learn how their research can play a role, but, more importantly, to influence those roadmaps to ensure successful delivery on DOE applications that will contribute to (if not solve) problems of national interest in national security, science, energy, and health, as well as growing security threats.

The third annual meeting was held in Houston, Texas, January 14–17, 2019. With a 19% increase in the number of registrations (768 people), and the change in location, the third annual meeting was considered the most impactful of the three at the time. The new website provided a better platform for the

[1] Typically referred to as "Facilities" throughout this document.
[2] Exascale is 10^{18} operations per second.

dissemination of the content, the new venue as a meeting hotel instead of a conference center facilitated interactions and discussions after event hours, and the addition of an award-winning mobile event conference app (Whova) transformed dramatically the attendee experience at the event.

This fourth annual meeting was held in Houston, Texas, February 3-7, 2020. This meeting had an increase in the number of attendees for a total of 824 people registered (782 attendees) and included numerous enhancements based on feedback and lessons learned from previous meetings, some of which are listed here:

- Improved quality of the sessions, their material and the whole program.
- Had more industry participation and addition of external collaborators from overseas.
- Published the full agenda earlier to better accommodate attendance and travel plans based on schedule.
- Centralized all sessions in one venue.
- Provided additional hotels and room blocks for the attendees.
- Improved communication with the audience (links, material, directions, notifications, etc.) to go paperless.
- Enhanced side meeting scheduling, management and user experience.
- Made available additional space and tables for impromptu meetings and side discussions.
- Improved IT and A/V solutions for speakers.

In addition, our final survey captured the following points as opportunities for improvement in future meetings: consider a different meeting location that is more pedestrian friendly, reduce talks during working meals to allow more collaboration and informal time, adapt the agenda to acknowledge attendees from different timezones, consider recording some of the tutorials and/or sessions to share broadly with the HPC community, have a larger poster room, provide additional power strips, and improve the WiFi.

1. EXASCALE COMPUTING PROJECT

In 2009, the US Department of Energy (DOE) began examining exascale computing as part of the solution to many of the nation's, and even the world's, most challenging problems. In fiscal year (FY) 2016, a 7-year combined effort by the Office of Science's (SC's) Advanced Scientific Computing Research (ASCR) office and the National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) office was begun to prepare and uplift the high-performance community toward capable exascale platforms, software, and applications.

The mission of the ASCR program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to DOE. A particular challenge of this program is fulfilling the computational and data science potential of emerging computing systems, which is believed to require a technology step change in today's tools and techniques to bring exascale-enabled, high-confidence solutions to the nation's most challenging problems.

Established in 1995, the ASC office supports the NNSA Office of Defense Program's simulation-based predictive science for the stewardship of the US nuclear weapons stockpile. Under ASC, computer simulation capabilities are developed to analyze and predict the performance, safety, and reliability of nuclear weapons and to certify their functionality. Beyond stewardship of today's stockpile, the rapidly growing capability of potential adversaries to defeat US weapons using advanced defensive systems represents a sobering threat. ASC applications must also target performance assessment of current and life-extended weapon systems subject to a wide variety of hostile environments and potential threat scenarios. To execute its mission for the NNSA Stockpile Stewardship Program, ASC oversees the high-performance simulation and computing work of three NNSA laboratories—Los Alamos National Laboratory (LANL), Lawrence Livermore National Laboratory (LLNL), and Sandia National Laboratories (SNL)—as a nationally coordinated program.

The Exascale Computing Project (ECP) delivers specific applications, software products, and outcomes on DOE computing facilities[3]. Integration across these elements for specific hardware technologies for exascale[4] system instantiations is fundamental to ECP success. The outcome of the ECP is the delivery of a capable exascale computing ecosystem to provide breakthrough solutions addressing our most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security. This outcome is not a matter of ensuring more powerful computing systems. The ECP is designed to create more valuable and rapid insights from a wide variety of applications ("capable"), which requires a much higher level of inherent efficacy in all methods, software tools, and ECP-enabled computing technologies to be acquired by DOE laboratories ("ecosystem").

[1] Typically referred to as "Facilities" throughout this document.
[2] Exascale is 10^{18} operations per second.

2. 2020 ECP ANNUAL MEETING

More than 800 HPC experts drawn from every area critical to delivering a comprehensive and capable exascale computing ecosystem were brought together for a week-long meeting to exchange ideas, discuss accomplishments, and set the stage for future collaborative work. These individuals are leaders in their fields, ranging from tools development to applications optimization and from co-design to continuous integration. This one-of-a-kind opportunity afforded participants numerous sessions to learn about cutting-edge research and development (R&D) advances in HPC and to participate in a briefing by six major US vendors on their HPC roadmaps. Among the many experts at the meeting, there were representatives, for example, in the following areas.

- Applications relevant to the DOE mission (e.g., in chemistry and materials, energy, earth and space science, and national security)
- Co-design, data analytics, code optimization, and abstract machine models
- Programming models, development tools, mathematical libraries, data and visualization
- Continuous integration and software delivery
- Hardware evaluation (e.g., for memory technologies, interconnect and system simulation)

The emphasis of the 2020 ECP annual meeting was on the discussion of technical results and performance. The meeting was planned by technical staff within the project, who focused the agenda on the dissemination of critical results to meeting participants, co-design working groups, and direct engagement with vendors on their technology plans. A complete program agenda can be found at https://www.ecpannualmeeting.com/overview.php. Meeting participants shared technical plans and milestone accomplishments—as well as challenges lying ahead—with other technical staff within the ECP community, including the ECP technical focus areas (Application Development, Software Technology, Hardware and Integration), DOE HPC facilities, and US HPC vendors. The ECP operates under the principle of co-design and integration; thus, this gathering of the ECP community provided a venue for the exchange of best practices, lessons learned, tools, and opportunities for collaboration and partnering. Networking was emphasized, and project participants were expected to explore ways to leverage accomplishments and share requirements with other ECP efforts. Key topics covered at the meeting included discussions of future systems, software stack plans, and interactions with facilities.

Attendees at the meeting included ECP-funded project team members, the ECP leadership team, representatives from the six DOE SC and NNSA HPC facilities, technical leads for six US HPC vendors, representatives from the Industry Council and external companies (such as Total, BP and Shell), development (R&D) projects within ECP's scope (ECP's PathForward effort), and DOE sponsors. Figure 1 (left) shows a representation with circles based on the number of people from different categories, with the biggest representation (594 people, 72%) from US Government / Government Contractor. Figure 1 (right) displays the representation from the six core ECP partner DOE laboratories: ORNL, ANL, LLNL, LBNL, LANL, SNL.

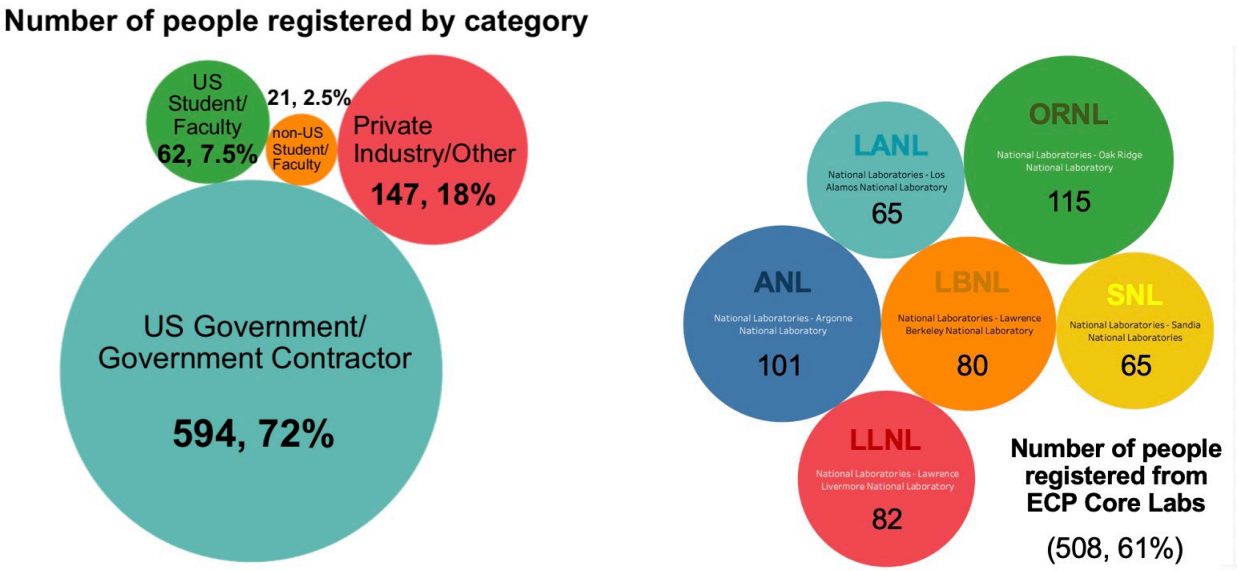

Figure 1. Left: number of people registered by category. Rigth: number of people registered from the six core ECP partner DOE Laboratories.

The size and geographical dispersion of the ECP team precludes, from a cost and efficiency point of view, frequent face-to-face gatherings of all key technical staff for working meetings on technical co-design activities and progress (give-and-take between software, hardware, applications, etc.). The ECP annual meeting affords this opportunity on a once-a-year basis. The annual meetings have reaped critical benefits for the ECP. The first meeting brought together a diverse collection of researchers from DOE national laboratories as well as university computer and computational science researchers to discuss shared problems and joint solutions for the development of a capable exascale computing ecosystem. These interactions resulted in enormous advances for ECP. The fourth meeting, held this year, provided numerous examples in which researchers with common goals and synergistic solutions came together for the first time to deliver tangible results. Further, at this year's meeting ECP researchers had the opportunity to digest all US HPC vendor R&D product roadmaps pointing to exascale—not only to learn how their research can play a role but, more importantly, to influence those roadmaps to ensure successful delivery of DOE applications.

3. COMMITTEE

The organizing committee was composed of skilled members from different national laboratories to assist in the process of planning, managing, and executing the numerous tasks required for the 2020 ECP annual meeting. The representatives of the committee worked together on the organization and planning of the conference to ensure a successful meeting.

This meeting was not a public event. It was open only to individuals funded by ECP (research projects and vendors) and staff at the DOE HPC facilities. Requests to attend by individuals not associated with the ECP were asked to be submitted to ecpam-general@exascaleproject.org.

Invitations were sent by email to the ECP DOE Program Directors, Federal Project Director, Board of Directors[5], Laboratory Operations Task Force (LOTF)[6], DOE Leadership Computing Facility (LCF) Directors, PathForward vendor representatives, Industry Council members, ECP funded staff, as well as other key DOE program office, vendor and HPC collaborators. A sample of the communication sent and the invitation included as an attachment is Figure 2 below.

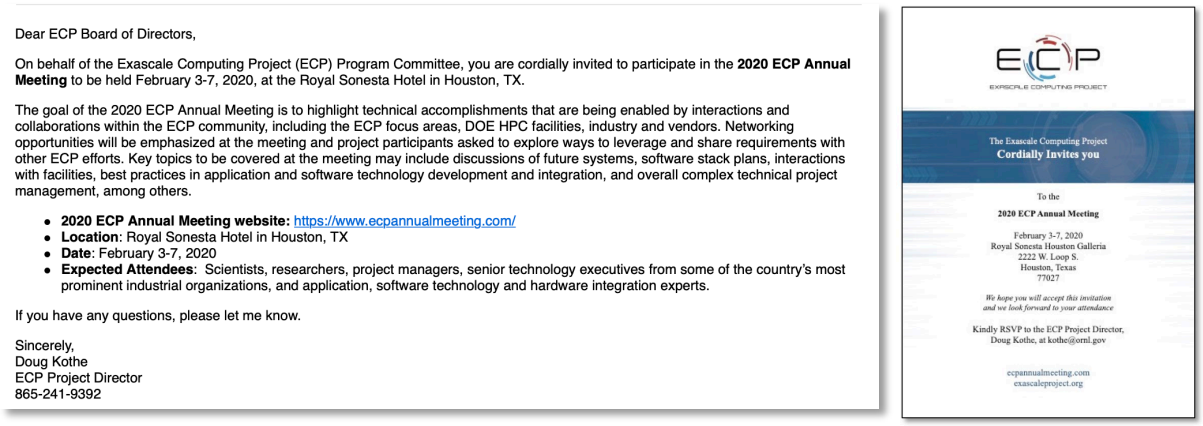

Dear ECP Board of Directors,

On behalf of the Exascale Computing Project (ECP) Program Committee, you are cordially invited to participate in the **2020 ECP Annual Meeting** to be held February 3-7, 2020, at the Royal Sonesta Hotel in Houston, TX.

The goal of the 2020 ECP Annual Meeting is to highlight technical accomplishments that are being enabled by interactions and collaborations within the ECP community, including the ECP focus areas, DOE HPC facilities, industry and vendors. Networking opportunities will be emphasized at the meeting and project participants asked to explore ways to leverage and share requirements with other ECP efforts. Key topics to be covered at the meeting may include discussions of future systems, software stack plans, interactions with facilities, best practices in application and software technology development and integration, and overall complex technical project management, among others.

- **2020 ECP Annual Meeting website:** https://www.ecpannualmeeting.com/
- **Location**: Royal Sonesta Hotel in Houston, TX
- **Date**: February 3-7, 2020
- **Expected Attendees**: Scientists, researchers, project managers, senior technology executives from some of the country's most prominent industrial organizations, and application, software technology and hardware integration experts.

If you have any questions, please let me know.

Sincerely,
Doug Kothe
ECP Project Director
865-241-9392

ECP EXASCALE COMPUTING PROJECT

The Exascale Computing Project
Cordially Invites you

To the
2020 ECP Annual Meeting

February 3-7, 2020
Royal Sonesta Houston Galleria
2222 W. Loop S.
Houston, Texas
77027

We hope you will accept this invitation and we look forward to your attendance

Kindly RSVP to the ECP Project Director,
Doug Kothe, at kothe@ornl.gov

ecpannualmeeting.com
exascaleproject.org

Figure 2. Sample of email communication and invitation sent

Foreign National participants to the meeting were vetted using processes in compliance with US export control laws and regulations. There were 244 non-US citizens registered for the ECP Annual Meeting, with 230 actually attending. Once the registration was completed, an ORNL employee would go on-line to the ORNL Personnel Access System (PAS) and send an invitation for the non-US citizen to complete. Once the non-US citizen submitted the PAS request, the ORNL employee would collect visa and passport information and submit the request along with the required documents to the ORNL Visitors Center for processing; each request needed 7 days to process.

When the approvals came through, the ORNL employee would send out an email to notify the attendee that their badge had been approved and what documents they needed to bring to the registration desk.

On the day of registration, two ORNL employees would check-in the non-US citizens by looking at their visa and passport and comparing it to what the attendee submitted. The attendee would sign the required document and then attend the meeting. At the end of the meeting, one of the ORNL employees would go into the PAS system and mark the attendee as "reported".

The committee opened calls and hosted sessions on key topics and areas to the ECP community including discussions of future systems, software stack plans, and interactions with facilities. In addition, the committee prepared the conference schedule and was responsible for the technical content, including breakouts, tutorials, and poster sessions offered during the week.

[5] Consists of the laboratory directors from the six core ECP partner DOE laboratories.
[6] Composed of associate laboratory director with line management responsibility for high performance computing at their respective ECP core partner DOE laboratories.

See the list of committee members below.

Responsibility	Committee Members	Institution
Keynote	Kathlyn Boudwin	ORNL
Program	Marta García Todd Munson	ANL ANL
Training	Ashley Barker Osni Marques	ORNL LBNL
Posters	Lora Wolfe Elaine Raybourn	ORNL SNL
PathForward Vendors	Bronis de Supinski	LLNL
Industry Council	David Martin Suzy Tichenor	ANL ORNL
Communications	Mike Bernhardt	ORNL
Information Technology	Jonathan Wilson Joe Lawson	ORNL ORNL
Project Office	Doug Collins	ORNL
Logistics	Corliss Thompson Ashton Chatman Lora Wolfe Judy Potok	ORNL ORNL ORNL ORNL

Additional support was provided by volunteers during on-site activities, such as for NDA session check-in and for registration.

Marta García and Todd Munson (ANL) were contacted by Kathlyn Boudin (ECP Project Management Director) at the end of February 2019 to inquire about their availability to take the lead as Program Co-Chairs of the 2020 ECP Annual Meeting. After internal approvals, both Program Co-Chairs officially started planning the 2020 Annual Meeting in mid-May, 2019. The Program Co-Chairs Marta García and Todd Munson dedicated 25% and 10% FTE, respectively, during a period of 10 months (mid-May 2019 – mid-March, 2020). The ECP Project Office Lead (Doug Collins), a member of the logistics Team (Ashton Chatman) and both Program Co-Chairs had their 2020 Annual Meeting Kick-Off meeting on May 23, 2019 for a very initial discussion to start planning the annual meeting. The objective of that preliminary discussion was to begin brainstorming the tasks that are needed and who should be the owners. Figure 3 shows a screenshot of the Gantt chart prepared as a timeline for the annual meeting and maintained by the committee members.

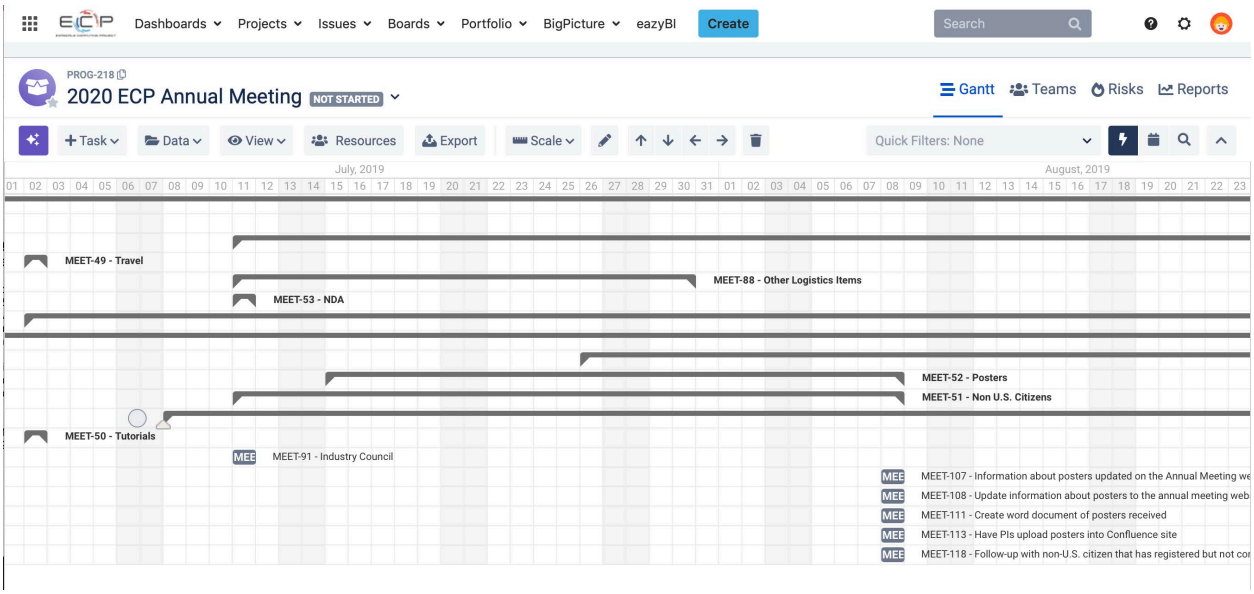

Figure 3. Screenshot of Gantt chart with tasks for the 2020 ECP Annual Meeting.

All team members were assigned a specific role and they worked on their tasks independently communicating by email (if necessary) or through their 2020 ECP Annual Meeting Slack channels. In addition, the committee members met periodically to share updates in one-hour meetings during the following dates:

- 2019: 5/23, 6/27, 7/11, 7/26, 8/8, 9/5, 9/19, 10/3, 10/17, 10/31, 11/14, 12/6, 12/12
- 2020: 1/9, 1/23, 2/2

The committee shared a working space in the ECP Confluence website to post documents, updates and relevant information related to the planning and execution of the annual meeting.

One of the first tasks that the committee worked on was the call for sessions that opened on July 26, 2019. After gathering almost 90 proposed sessions of different types: Birds of a Feather (BoF), Breakout, Panel, Plenary or Tutorials, the Program Co-Chairs evaluated all the submissions (see APPENDIX A, APPENDIX B and APPENDIX C) and started planning the agenda with the objective of maximizing the number of sessions, contributing authors, available seats, and number of rooms that constituted the technical program. The Program Co-Chairs established several conflict resolution policies that constituted the conflict resolution framework used to support the 2020 ECP Annual Meeting agenda strategies:

- **No presence** takes precedence, e.g. if a speaker arrives after the meeting has started or has to leave before the meeting ends.
- **Length of the session** takes precedence, e.g. the length of the sessions was different (45-min, 60-min, 75-min, etc.) and sessions were distributed depending on the length of the slot (equal or higher than the session length).
- **Conflict with other sessions** takes precedence, e.g. if the same speaker has submitted several sessions we cannot schedule her/him at the same time.

- **Conflict with same topic** takes precedence, e.g. several Spack or Kokkos or CMake sessions at the same time.
- **Balance of AD/ST/HI** takes precedence, e.g. distribute the submissions of focus area topics to avoid having too many AD or ST or HI sessions at the same time.
- **Preferred day** takes precedence, e.g. author's preference in the presentation dates were considered and we tried to accommodate them.
- **Number of authors** takes precedence, e.g. if the number of authors was low, there was more flexibility to move the session to another potential slot.
- **NDA and PathForward content (or not)** takes precedence, e.g. PathForward sessions were scheduled on a particular day: Monday of the meeting week.
- **Expected audience vs room size** takes precedence, e.g. sessions with expected low or high attendance were scheduled in particular rooms.
- **Kind of session** takes precedence, e.g. Keynote vs Plenary vs Tutorial vs Panel vs Breakout vs BoF.

A sample of the representation and graphs (2020 ECP Annual Meeting *Compute Board*) used to prepare the agenda is presented below (a) as well as the "human-oriented" representation of the agenda in PDF format (b) that was shared with the general public.

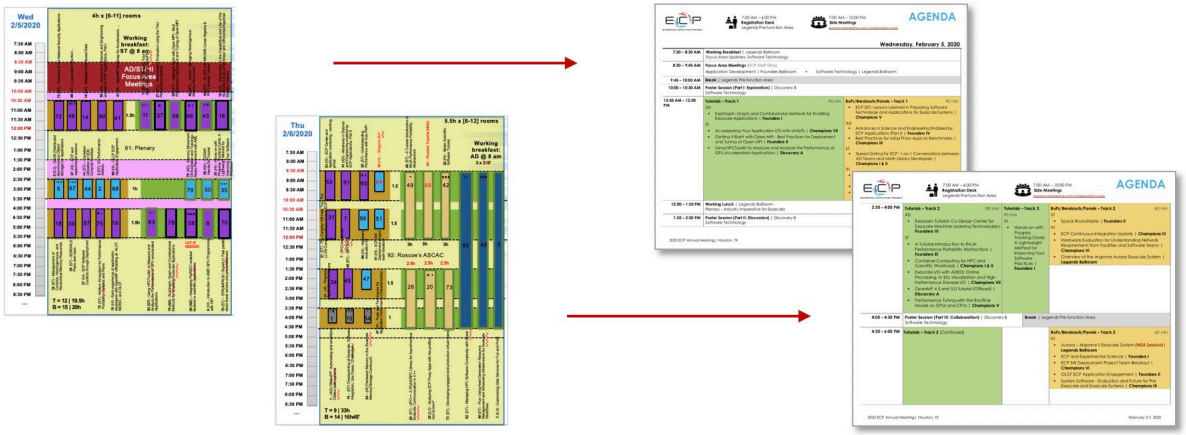

a) Graphs with conflict resolution framework **b) Final Agenda**

Figure 4. (a) Sample of diagrams used to prepare agenda with information about sessions, (b) and screenshot of the respective final agenda.

The 2020 ECP annual meeting website opened on October 31, 2019 with the full agenda, the call for posters for the poster sessions and the link to the registration portal. During the months of November and December, the committee members focused all their effort on outreach communications for sessions (see APPENDIX D) and poster presenters (see APPENDIX E) and general expected audience and logistics tasks, e.g.: coordination of NDA sessions, support for questions, approvals for non-ECP members, shuttle services, side meetings, preparation of the posters sessions, etc. The month of January was instrumental to the completion of tasks essential to the meeting and the beginning of the preparation of this report and the corresponding slide deck (see APPENDIX F).

4. HIGHLIGHTS

4.1 KEYNOTE

The keynote speaker was determined based on input from the program office, who provided a ranked list of possible speakers. Following the appropriate protocol, each was asked in order if they were willing to serve as the keynote speaker. The first speaker on the list, Dr. Chris Fall, accepted the invitation. He was offered two times to speak, Monday night or Tuesday morning, eventually settling on Tuesday morning for his session. The delay in the choice of time had ripple effects on the program and the program committee assembled two programs to accommodate either choice. In the future, we recommend only one date/time be provided to the candidate keynote speakers to remove uncertainty in scheduling the remainder of the program. The following is the description of the keynote delivered at the meeting.

- **Date/Time**: Tuesday, February 4, 2020 @ 8:30–10:00 AM
- **Speakers**: Dr. Chris Fall (Director, DOE Office of Science)

Dr. Chris Fall serves as Director of the Department of Energy's Office of Science, the lead federal agency supporting fundamental scientific research for energy and the nation's largest supporter of basic research in the physical sciences. He oversees the office's two principal thrusts: direct support of scientific research and development, and construction, and operation of unique, open-access scientific user facilities that are made available to external researchers. The Office of Science also is responsible for stewardship of 10 of the department's 17 national laboratories.

Before joining the Office of Science, Fall served as a Senior Advisor to the Undersecretary for Energy and as Acting Director of DOE's Advanced Research Projects Agency–Energy (ARPA-E). Fall came to DOE from the Office of Naval Research (ONR), where he served for more than 7 years in a variety of roles including Acting Chief Scientist and Lead for the Research Directorate, Deputy Director of Research, Director of the International Liaison Office, and the ONR Innovation Fellow. While on loan from ONR, Fall served for 3 years in the White House Office of Science and Technology Policy as Assistant Director for Defense Programs and then as Acting Lead for the National Security and International Affairs Division. Before government service, Fall was a faculty member at the University of Illinois at Chicago, and he completed postdoctoral fellowships at the University of California at Davis Institute for Theoretical Dynamics and the New York University Center for Neural Science.

Fall earned a Ph.D. in Neuroscience and a B.S. in Mechanical Engineering from the University of Virginia. He also holds an MBA from Northwestern University's Kellogg School of Management.

4.2 PLENARY TALKS

Experts in the community were invited to address the meeting participants in plenary sessions. Plenary sessions focused on themes of interest to the community at large. As part of the call for sessions, speakers could submit ideas for plenary talks. The program committee also contributed topics that they considered timely and relevant to the ECP community. After the topics were solicited, the program committee consulted with the ECP leadership and then decided on the speakers and topics. We deliberately chose to place the application-related plenary on Tuesday after the Keynote so that the DOE officials in attendance could see the impact ECP is having on applications. The following are descriptions of plenaries that were presented at the meeting.

– *AI for Science* –

- **Date/Time**: Monday, February 3, 2020 @ 6:30–7:30 PM
- **Speaker**: Rick Stevens (Argonne National Laboratory / The University of Chicago)

In this talk, Rick described an emerging DOE initiative to advance the concept of Artificial Intelligence (AI) aimed at addressing challenge problems in science. This initiative is called "AI for Science." The basic concept is threefold: (1) to identify those scientific problems where existing AI and machine learning methods can have an immediate impact (and organize teams and efforts to realize that impact); (2) to identify areas of where new AI methods are needed to meet the unique needs of science research (frame the problems, develop test cases, and outline work needed to make progress); and (3) to develop the means to automate scientific experiments, observations, and data generation to accelerate the overall scientific enterprise. Science offers plenty of hard problems to motivate and drive AI research, from complex multimodal data analysis, to integration of symbolic and data-intensive methods, to coupling large-scale simulation and machine learning to drive improved training to control and accelerate simulations. A major sub-theme is the idea of working toward the automation of scientific discovery through integration of machine learning (active learning and reinforcement learning) with simulation and automated high-throughput experimental laboratories. Rick provided some examples of projects that are under way and laid out a set of long-term driver problems.

– *EQSIM: Transforming Earthquake Hazard and Risk Assessment Through Exascale Simulations* –

- **Date/Time**: Tuesday, February 4, 2020 @ 12:30–1:30 PM
- **Speakers**: David McCallen (Lawrence Berkeley National Laboratory)

Large earthquakes present a significant risk around the world and are a major issue across the DOE mission space ranging from the safety of DOE's own inventory of one-of-a-kind mission critical facilities to all major US energy systems (electric/gas distribution systems, renewable energy production facilities, nuclear power plants etc.). Beyond the DOE enterprise, addressing earthquake risk is a major worldwide societal challenge for virtually every element of the built environment including transportation, health, data/commerce, and all urban infrastructure. The tremendous developments occurring in high performance computing will transform earthquake hazard and risk assessments. As computational power increases, the reliance on simplifying idealizations, approximations, and legacy empirical methods can diminish, and attention can be focused on dealing with the fundamental physics uncertainties in earthquake processes. With the advent of exascale computations, regional-scale ground motion simulations are becoming computationally feasible at frequencies of engineering interest, and simulation models that connect the domains of seismology, geotechnical, and structural engineering are within grasp. The ECP EQSIM Application Development project is focused on creating an unprecedented computational toolset and workflow for earthquake hazard and risk assessment. EQSIM is building an end-to-end, fault-to-structure capability to simulate from the initial fault rupture to surface ground motions (earthquake hazard) and ultimately to infrastructure response (earthquake risk) with the ultimate goal of removing computational limitations as a barrier to scientific exploration and understanding of earthquake phenomenology. This presentation described the EQSIM framework and highlighted the computational advancements achieved to date, including key advanced algorithm development and transition to GPU-based platforms. In addition, the enabling interdependencies with ECP S&T software developments to fully develop the multidisciplinary fault-to-structure workflow were described.

– *Industry Imperatives for Exascale* –

- **Date/Time**: Wednesday, February 5, 2020 @ 12:30–1:30 PM
- **Organizers:**
  - David Martin (Argonne National Laboratory)
  - Suzy Tichenor (Oak Ridge National Laboratory)

- **Moderator:**
  - Brunon (Dave) Kepczynski (Chief Information Officer of General Electric Research)
- **Panelists:**
  - Prith Banerjee (Chief Technology Officer, ANSYS)
  - Frank Ham (Chief Executive Officer, Cascade Technologies)
  - Bill Nitzberg (Chief Technology Officer of PBS Works at Altair Engineering)

Independent software vendor (ISV) executives on the ECP Industry Council discussed how they are preparing their codes for emerging exascale systems as well as how they are preparing their customers. Because most industrial users do not write their own software, industry's ability to use exascale systems will be tied closely to ISV codes. Moderated by ECP Industry Council chair Brunon (Dave) Kepczynski, the panelists discussed their strategy, hurdles they are facing, and tools and techniques from ECP they think will help them.

– *ASCAC Transition of Exascale Computing Project Report* –

- **Date/Time**: Thursday, February 6, 2020 @ 12:30–1:30 PM
- **Organizers**: Prof. Roscoe Giles (Boston University)

The ASCR Advisory Committee (ASCAC) has been charged by the DOE Office of Science to identify elements of ECP that should be transitioned into ASCR research or other SC/ASCR initiatives. The ASCAC Transition Subcommittee was convened to study the situation and develop a detailed suggested response to this charge. The report provides guiding strategies and approaches that will be key to ensuring future US leadership across the full range of disciplines stewarded by ASCR. This talk provided an update on the activities of the Transition Subcommittee and the status and content of the report provided to ASCAC.

A summary of the plenary sessions can be found in APPENDIX A.

4.3 BOF/BREAKOUT/PANEL SESSIONS

For the 2020 ECP annual meeting, members of the ECP community were asked to propose, organize, and lead technical sessions on topics on the critical path to their performance goals, within ECP's focus area scope, and relevant to roadmaps for DOE HPC facilities and US HPC vendors. The process was to send email to the entire ECP team to solicit sessions with a deadline. The team members would then enter their ideas into an online form that asked for the details of the proposed session. The program team noticed that some of the focus areas had not supplied as many sessions as desired, so they worked with the leadership of those focus areas to encourage their team members to submit additional sessions. Once all of the information was gathered, including the speaker information, restrictions on dates, and conflicts with other sessions, the program committee assembled a program to satisfy the constraints, including room assignments (see conflict resolution policies in Section 3). The assignment to rooms was based on historical attendance of similar sessions from previous annual meetings and survey data from participants on which sessions they wanted to attend. For this meeting, two programs were generated as there was a delay in determining the date of the keynote. Once the date of the keynote was determined, the corresponding program was made available.

The program committee made an early decision to place all of the PathForward NDA sessions on the first day of the annual meeting to eliminate conflicts with the rest of the program. These NDA sessions are heavily attended by the ECP community and were offered sequentially, so that participants could attend each one and without having to repeat sessions. This decision made for a long day for the participants and a need to manage the logistics of getting everone into and out of the assigned rooms.

The final program included a diverse series of 53 BoF, Breakout, and Panel sessions 45, 60, 75, and 90 minutes in length were held on Tuesday, Wednesday, and Thursday of the annual meeting. The sessions promoted knowledge transfer relevant to R&D to accelerate delivery of a capable exascale computing ecosystem. Participants were encouraged to attend sessions to learn more about these important topics and areas in the ECP community. Speakers discussed accomplishments, requirements sharing, and opportunities for interactions. This community-driven set of events was an exciting opportunity for thought-leaders to share and exchange advances in the field and identify opportunities for closer coordination of activities.

The following is a complete list of the 53 BoF/Breakout/Panel Sessions (in alphabetical order).

- Advances in Science and Engineering Enabled by ECP Applications (Part I)
- Advances in Science and Engineering Enabled by ECP Applications (Part II)
- AMReX and AMReX-Based Applications
- App Integration at the Facilities: Lessons Learned So Far
- Aurora – Argonne's Exascale System (NDA Session)
- Best Practices for Using Proxy Apps as Benchmarks
- CEED: High-Order Methods, Applications, and Performance for Exascale
- Checkpointing at Exascale: Software Integration, Use Cases, Challenges
- CMakePP: Automating and Simplifying CMake Build Systems
- Collaboration Opportunities with the BSSw Fellows
- Container Utilization at DOE Compute Facilities
- Cultivating Software Sustainability and Productivity through BSSw.io
- DAOS (Distributed Asynchronous Object Storage) for Applications
- Early Experience of Application Developers with OpenMP Offloading at ALCF, NERSC, and OLCF
- ECP 201: Lessons Learned in Preparing Software Technology and Applications for Exascale Systems
- ECP and Experimental Science
- ECP Center and Application Monitoring – Working Group
- ECP Continuous Integration Update
- ECP Software Curation through Deployment
- ECP SW Deployment Project Team Breakout
- Hardware Evaluation for Understanding Network Requirements from Facilities and Software Teams
- I/O Performance Addicts
- Intel Developer Feedback Session
- Intelligent Distributed Data Movement for Exascale Computing
- Kokkos Ecosystem for Performance Portability: Updates and Plans
- Lossy Data Reduction/Compression for ECP Applications
- Management of Complex, Heterogeneous, and Hierarchical Memory Resources
- MPI BoF
- Multiprecision Focus Effort in the ECP Math Library Ecosystem
- OLCF ECP Application Engagement
- OpenMP Roadmap for Accelerators across DOE Pre-Exascale/Exascale Machines
- Overview of the Argonne Aurora Exascale System
- Overview of the National Security Applications
- PathForward Update – AMD (NDA Session)
- PathForward Update – Cray (NDA Session)

- PathForward Update – HPE (NDA Session)
- PathForward Update – IBM (NDA Session)
- PathForward Update – Intel (NDA Session)
- PathForward Update – NVIDIA (NDA Session)
- Performance Portability at the Intersection of the Exascale Computing Project and the DOE Centers of Excellence
- Perlmutter - a Waypoint for ECP Teams
- Persistent Memory in the Exascale Memory/Storage Continuum
- Preparing Applications for Aurora
- Slingshot BoF (NDA Session)
- Spack Roundtable
- Speed Dating for ECP: 1-on-1 Conversations between AD Teams and Math Library Developers
- SUNDIALS Plans and Uses
- SYCL Programming Model for Aurora
- System Software - Evaluation and Future for Pre-Exascale and Exascale Systems
- The Compiler Will Help You! The LLVM Compiler Infrastructure in ECP
- The Scope and Role of DevOps in HPC
- Tools and Techniques for Application Performance Characterization
- Understanding Performance with Exa-PAPI

Details of all the BoF/breakout/panel sessions can be found in APPENDIX B, and the list of session authors is in Appendix D.

4.4 TUTORIAL SESSIONS

Training is a critical element of the ECP. It is used to disseminate and transfer application development knowledge, lessons learned, and best practices across application teams and to provide training on key ECP technologies. For applications to take full advantage of exascale hardware and software, a robust developer training program is necessary to keep participants abreast of emerging technologies.

Members of the ECP community and facilities were solicited for topics to include in the tutorials program of the 2020 ECP annual meeting. The ECP agenda included a series of thirty-five 60, 90, 150, 180, and 330 minute tutorials offered Tuesday through Thursday at the annual meeting covering a diverse set of topics. More in-depth than the BoF/breakout/panel sessions, the tutorials afforded participants a chance to learn more about and to experiment in various areas including algorithms and methods, high performing libraries, memory and storage hierarchies, on-node parallelism and vectorization, application portability techniques, and software engineering design principles and best practices.

The following is a complete list of the 35 tutorial sessions (in alphabetical order).

- A Tutorial Introduction to RAJA Performance Portability Abstractions
- Accelerating Your Application I/O with UnifyFS
- Achieving High Performance I/O with HDF5
- Analyzing ECP Proxy Apps with the Profiling Tool Score-P
- Application-Driven Fault-Tolerance for High Performance Distributed Computing
- Better Scientific Software Tutorial
- CANDLE Hands-on Tutorial
- Co-Design Center for Particle Applications (CoPA) Summary
- CODAR Hands-on Tutorial

- Container Computing for HPC and Scientific Workloads
- Customizing Data Services for Fun and Profit
- Developing Engaged and Productive Virtual Teams
- ECP Continuous Integration Startup Tutorial
- ExaGraph: Graph and Combinatorial Methods for Enabling Exascale Applications
- ExaLearn Tutorial: Co-Design Center for Exascale Machine Learning Technologies
- Exascale I/O with ADIOS: Online Processing, in Situ Visualization, and High-Performance Storage I/O
- Flux: Using Next-Generation Resource Management and Scheduling Infrastructure for Exascale Workflows
- Frontier Tutorial
- Getting It Right with Open MPI - Best Practices for Deployment and Tuning of Open MPI
- Hands-on with Progress Tracking Cards: A Lightweight Method for Improving Your Software Practices
- In Situ Visualization and Analysis with Ascent
- Integrating PaRSEC-Enabled Libraries in Scientific Applications
- Introduction to the Capabilities and Use of the SUNDIALS Suite of Nonlinear and Differential/Algebraic Equation Solvers
- Managing HPC Software Complexity with Spack
- Managing Power Efficiency of HPC Applications with Variorum and GEOPM
- Memory Movement Orchestration and Topology-Aware Placement with AML
- Modern CMake Tutorial
- OpenMP 4.5 and 5.0 Tutorial (Offload)
- Performance Evaluation Using the TAU Performance System
- Performance Tuning with the Roofline Model on GPUs and CPUs
- SLATE and MAGMA Linear Algebra & FFT Libraries Tutorial
- STRUMPACK / SuperLU: Fast Parallel Direct Linear Solvers and Preconditioners
- Umpire: Managing Heterogenous Memory Resources
- UPC++: A PGAS/RPC Library for Asynchronous Exascale Communication in C++
- Using HPCToolkit to Measure and Analyze the Performance of GPU-Accelerated Applications

Details of all the tutorials can be found in APPENDIX C, and the list of session authors is in APPENDIX D.

4.5 POSTER SESSIONS

The 2020 poster session was organized to engender collaboration and information sharing across five ECP functional areas: Application Development, Software Technologies, Hardware and Integration, DOE HPC Facilities, and Industry. The co-chairs solicited both virtual (PDF) and physical posters, along with three key terms describing the poster content from each team. The PDF versions of the poster were available via an interactive digital display throughout the meeting. The key terms coupled with poster content were used to organize the physical layout of the poster session in order to maximize attendee interaction. The co-chairs leveraged theories of social networks and proximity-based face-to-face communication to enhance co-located interaction during the poster session. For example, posters of application development teams leveraging specific software technologies and/or facilities were placed in physical proximity of the others to facilitate collaboration and idea sharing among teams. The posters were on display during two poster sessions that were held on Tuesday and Wednesday of the annual meeting. Each poster session was divided into three parts: Part I – Exploration (30 minutes), Part II – Discussion (60 minutes), and Part III – Collaboration (30 minutes).

A total of 123 posters were presented throughout the annual meeting, representing 100% participation. In a few instances, some teams presented more than one poster—indicative of progress made on activities that merited further elaboration. Poster session attendance was excellent, and the poster sessions provided a popular avenue for networking among the project teams, facility representatives, vendors, and members of the ECP Industry Council and Industry partners. Lessons learned from this new approach will be incorporated into future poster session planning.

Feedback from participants of ECP's previous annual meetings indicated that the poster sessions in general present very valuable opportunities to receive a high-level overview of all of the R&D activities being executed within the ECP. Finally, since each poster was presented by one or more of the project team members, interested attendees explored the potential use and integration of the technologies and solutions presented and other collaboration opportunities.

The following is a complete list of the 123 posters (in alphabetical order).

APPLICATION DEVELOPMENT (38)

- AMReX: A Software Framework for Block-Structured AMR Applications
- Authentication and Management of Exascale Computing Resources
- CANDLE: Exascale Deep Learning Enabled Precision Medicine for Cancer
- CEED: Center for Efficient Exascale Discretizations
- CODAR Technologies: Online Data Analysis and Reduction Motifs
- CoPA: Co-Design Center for Particle Applications
- Data Analytics at the Exascale for Free Electron Lasers
- E3SM-MMF on Summit: Early Science and Computational Results
- Enabling GAMESS for Extreme Computing in Chemistry and Materials
- EQSIM: High Performance, Multidisciplinary Simulations for Regional Scale Earthquake Hazard and Risk Assessments
- ExaBiome: Exascale Solutions for Microbiome Analysis
- ExaGraph: Graph and Combinatorial Methods for Enabling Exascale Applications
- ExaLearn
- EXascale Atomistics for Accuracy, Length and Time
- Exascale Subsurface Simulator of Coupled Flow, Transport, Reactions and Mechanics
- ExaSGD – Optimization grid dynamics at exascale
- ExaSky Project: Computing the Sky at Extreme Scales
- ExaSMR: Monte Carlo Neutronics and CFD Simulation of Small Modular Reactors
- ExaStar: Exascale simulations of stellar explosions
- ExaWind: Exascale Predictive Wind Plant Flow Physics Modeling
- Grid QCD code: performance portability
- Hybrid Multi-grid/Deflation Algorithms for Lattice QCD
- Implementing the LQCD Wilson Dslash Operator using Kokkos and SYCL
- LANL Ristra: next-generation multi-physics
- MFiX-Exa: Adapting CFD-DEM to Exascale Computing
- Mixed Precision Optimization (MPO)
- NWChemEx - Moving Computational Chemistry to the Exascale
- Pele - Transforming Combustion Science and Technology with Exascale Simulation
- Performance & technology improvements for code-coupling in WDMApp
- Process Simulation at the Fidelity of the Microstructure
- Progress toward Performance Portability in QMCPACK

- PROTEAS-TUNE Overview
- Scalable Stochastic Transmission Expansion: A Use Case for ExaSGD
- The Hidden Mystery Behind Proxy Applications
- The MARBL Multi-physics Code
- Use of CUDA Profiling Tools Interface (CUPTI) for Profiling Asynchronous GPU Activity
- User-Directed Loop Transformations
- WarpX: Exascale modeling of plasma particle accelerators

SOFTWARE TECHNOLOGY (62)

- Accelerating Your Application I/O with UnifyFS
- ADIOS Framework for Scientific Data on Exascale Systems
- AID: Advanced Infrastructure for Debugging
- ALExa: Accelerated Libraries for Exascale
- ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis
- Arbalest: Dynamic Data Mapping Issue Detector for OpenMP Applications
- ATDM Programming Models and Runtimes
- BeeFlow Container Workflow Engine
- CHAI: Complementing RAJA with Heterogeneous Memory Support
- Clacc: OpenACC Support for Clang and LLVM
- Clover: Computational Libraries for Exascale
- Concurrency in LLVM
- Data Libraries and Services Enabling Exascale Science
- Delivering efficient parallel I/O with HDF5
- Distributed Performance-Portable Stencil Computation with Bricks
- DTE - Distributed Tasking for Exascale
- DTE PaRSEC enabled libraries
- E4S - The Extreme-scale Scientific Software Stack
- ECP Math Library Multiprecision Effort
- Exascale OpenMP Runtime
- FFT and Solver Libraries for Exascale: FFTX and SpectralPack
- Flang
- Flux: Next-Generation Resource Management and Scheduling Software Framework
- GASNet-EX: RMA and Active Message Communication for Exascale Programming Models
- Ginkgo - A Sparse Linear Solver Library for Manycore Architectures
- heFFTe: Highly Efficient FFTs for Exascale
- hypre: High Performance Solvers and Preconditioners
- Kokkos Core - Performance, Portability and Productivity
- Kokkos Kernels: Performance Portable Kernels for Sparse/Dense Linear Algebra, Graph and Machine Learning Kernels
- Memory management in Argo
- Metall: A Memory Allocator for Persistent Memory
- MFEM: A Modular Finite Elements Library
- MPICH: A High-Performance MPI Implementation
- OMPSan: Dataflow Analysis for OpenMP Code Sanitization
- Open MPI for Exascale (OMPI-X)
- OpenMP in LLVM
- Papyrus: Parallel Aggregate Persistent Storage

- Performance Measurement and Analysis using the TAU Performance System
- PLASMA Library on CORAL Systems and Beyond
- Power and Resource Management in Argo
- PROTEAS-TUNE: Autotuning Interface and Application to OpenMP Pragmas and Loop Transformation Pragmas Using YTOPT
- RAJA: Portable, Single-source Loop Execution for C++ Applications
- Recent PETSc/TAO Enhancements for Exascale
- Simplified Interface to Complex Memory (SICM)
- SLATE: Scalable Linear Algebra Targeting Exascale
- SNL ATDM Software Ecosystem Operating Systems and On-Node Runtime
- SOLLVE OpenMP Validation and Verification
- SOLLVE: Scaling OpenMP with LLVM for Exascale Performance and Portability
- Spack for users, developers, and HPC administrators
- STRUMPACK & ButterflyPACK: Fast solvers for sparse and dense linear algebra
- SUNDIALS: Suite of Nonlinear and Differential/Algebraic Equation Solvers
- SuperLU sparse direct solver
- SZ: Exascale Lossy Compression for Scientific Data – a software of ECP VeloC-SZ project
- The Cinema Project
- The Exascale Performance API with Modern C++
- The Extreme-scale Scientific Software Development Kit (xSDK) for ECP
- The Legion Programming System
- Umpire: An Application-Focused API for Memory Management
- UPC++: Asynchronous RMA and RPC Communication for Exascale Applications
- VeloC: Very Low Overhead Chekpoint Restart
- VTK-m: Updating HPC Visualization Software for Exascale-Era Processors
- ZFP: Compressed Multidimensional Arrays

HARDWARE & INTEGRATION (12)

- BSSw.io: Community resources for improving scientific software productivity and sustainability
- ECP Continuous Integration – Federated cross-site testing
- ECP Training Updates
- Hardware Evaluation
- IDEAS-ECP: Advancing Scientific Productivity through Better Scientific Software
- Lightweight Software Quality Practice – PSIP: Productivity Sustainability Improvement Planning
- NESAP: Preparing NERSC Apps for Exascale
- Popper 2.0: A Container-based Workflow Engine for Testing Complex Applications and Validating Scientific Claims
- Rate Your Project's Software Practices!
- Software Curation – Deployment at Facility
- Supercontainers
- The Argonne Training Program on Extreme Scale Computing (ATPESC)

FACILITIES (7)

- Argonne Leadership Computing Facility
- Aurora, Argonne's Exascale SystemArgonne
- Exploring Deep Learning for Science Benchmarks on DOE Supercomputers

- Perlmutter: A Pre-Exascale System for Science
- Power consumption in Pre-Exascale Systems: A case study on NERSC Cori
- Slate – Container Orchestration at OLCF
- Supporting experimental science at Exascale with the NERSC Superfacility Project

INDUSTRY (4)

- An Integrated Communication & Acceleration Framework with COPA
- Enhancing OFI (OpenFabric) networking interface to support acceleration capabilities in Configurable Network Protocol Accelerator (COPA)
- Goodyear's Exascale Challenge Problems
- The Hartree Centre – Overview & Exascale Highlights

A list of poster presenters can be found in APPENDIX E.

The next pages include thumbnails and titles of some of the posters on display during the poster sessions following the same five categories above mentioned. A more extended list is available at https://ecpannualmeeting.com/poster-interface23894732314e23hreu823rd/.

APPLICATION DEVELOPMENT

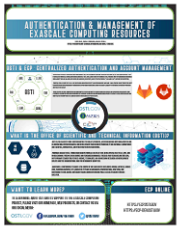

Authentication and Management of Exascale Computing Resources

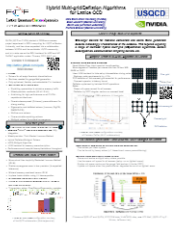

Hybrid Multi-grid/Deflation Algorithms for Lattice QCD

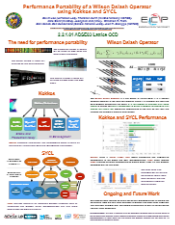

Implementing the LQCD Wilson Dslash Operator using Kokkos and SYCL

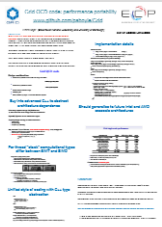

Grid QCD code: performance portability

NWChemEx - Moving Computational Chemistry to the Exascale

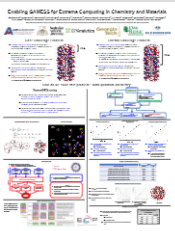

Enabling GAMESS for Extreme Computing in Chemistry and Materials

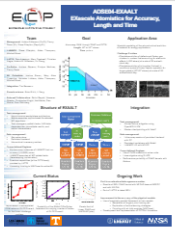

EXascale Atomistics for Accuracy, Length and Time

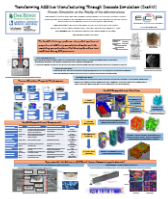

Process Simulation at the Fidelity of the Microstructure

Progress towards Performance Portability in QMCPACK

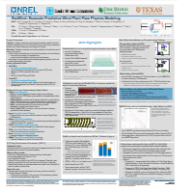

ExaWind: Exascale Predictive Wind Plant Flow Physics Modeling

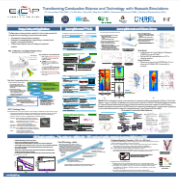

Pele - Transforming Combustion Science and Technology with Exascale Simulation

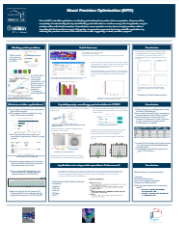

Mixed Precision Optimization (MPO)

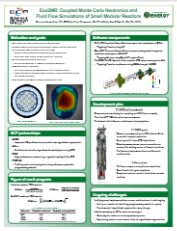

ExaSMR: Monte Carlo Neutronics and CFD Simulation of Small Modular Reactors

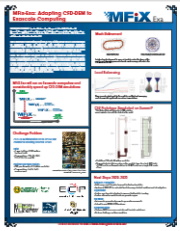

MFiX-Exa: Adapting CFD-DEM to Exascale Computing

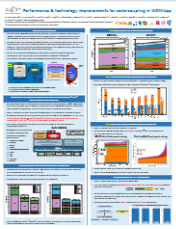

Performance & technology improvements for code-coupling in WDMApp

SOFTWARE TECHNOLOGY

E4S - The Extreme-scale Scientific Software Stack

MPICH: A High-Performance MPI Implementation

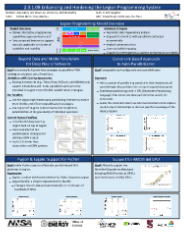

The Legion Programming System

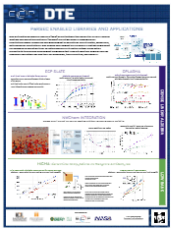

DTE PaRSEC enabled libraries

DTE - Distributed Tasking for Exascale

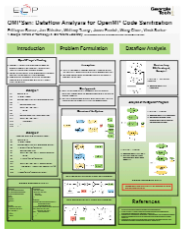

OMPSan: Dataflow Analysis for OpenMP Code Sanitization

SOLLVE OpenMP Validation and Verification

UPC++: Asynchronous RMA and RPC Communication for Exascale Applications

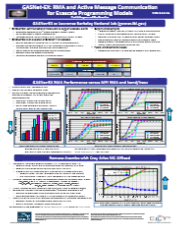

GASNet-EX: RMA and Active Message Communication for Exascale Programming Models

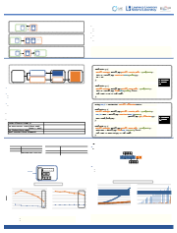

Metall: A Memory Allocator for Persistent Memory

Open MPI for Exascale (OMPI-X)

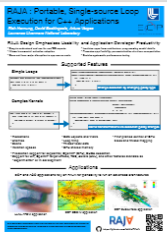

RAJA: Portable, Single-source Loop Execution for C++ Applications

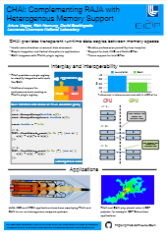

CHAI: Complementing RAJA with Heterogeneous Memory Support

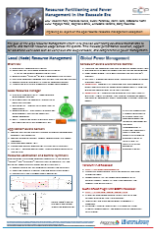

Power and Resource Management in Argo

Memory management in Argo

HARDWARE & INTEGRATION

Hardware Evaluation

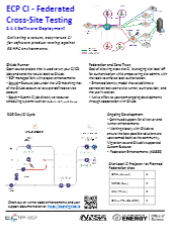

ECP Continuous Integration - Federated cross-site testing

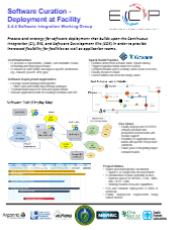

Software Curation - Deployment at Facility

ECP Training Updates

The Argonne Training Program on Extreme Scale Computing (ATPESC)

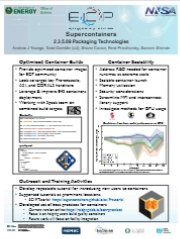

Supercontainers

Popper 2.0: A Container-based Workflow Engine for Testing Complex Applications and Validating Scientific Claims

BSSw.io: Community resources for improving scientific software productivity and sustainability

Lightweight Software Quality Practice – PSIP: Productivity Sustainability Improvement Planning

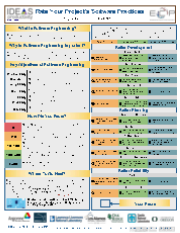

Rate Your Project's Software Practices!

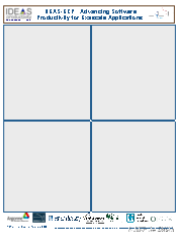

IDEAS-ECP: Advancing Scientific Productivity through Better Scientific Software

FACILITIES

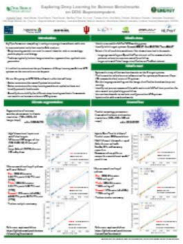

Exploring Deep Learning for Science Benchmarks on DOE Supercomputers

Perlmutter: A Pre-Exascale System for Science

Supporting experimental science at Exascale with the NERSC Superfacility Project

Slate - Container Orchestration at OLCF

Power consumption in Pre-Exascale Systems: A case study on NERSC Cori

INDUSTRY

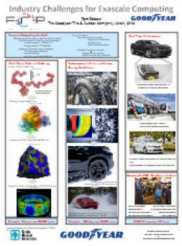

Goodyear's Exascale Challenge Problems

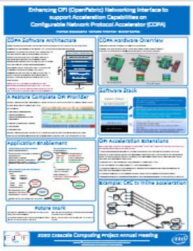

Enhancing OFI (OpenFabric) networking interface to support acceleration capabilities in Configurable Network Protocol Accelerator (COPA)

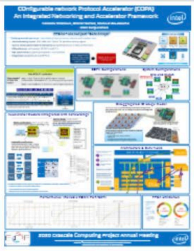

An Integrated Communication & Acceleration Framework with COPA

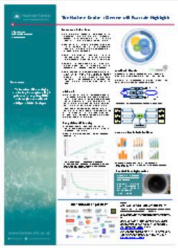

The Hartree Centre - Overview & Exascale Highlights

5. CONCLUSION

The ECP annual meeting is a complex undertaking that includes location selection, logistics, IT (website and event software), a solicitation of session, tutorials, and posters, creating the program, preparing and executing contracts, preparation of surveys, reports and its respective analysis, and, last but not least, ensuring a successful hosting of the event. The team worked together to handle the complexity and the moving parts. As much as possible, starting early and reducing uncertainties are important.

A key uncertainty this meeting was the date and time of the keynote speaker, which held up the generation and publishing of the program. Adding the content on the meeting website of more than 90 session and other events with more than 330 speakers was also a herculean task. There is a large effort required for NDA sessions, both before and during the meeting, so limiting the NDA sessions in the future would be prudent. The badges also seem to always arrive late to the meeting, so we should send them earlier if possible. After the meeting, it was decided to make the session and tutorial slides and poster content publicly available, which required the submitters to go through an extra step of determining if they wanted their content made available and actively checking a box to indicate that the content has been through their review and release process. For future meetings, having this information provided during session and poster submission and when gathering the content would be preferable. To conclude, data gathering, data curation and analysis of all the material in order to prepare this report and the corresponding executive summary slides is certainly a time-consuming task not to be neglected.

All in all, the meeting was successful and we thank the team for their expertise, diligence, and effort to make it run smoothly.

6. ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two DOE organizations—the Office of Science and the National Nuclear Security Administration—responsible for the planning and preparation of a capable exascale ecosystem—including software, applications, hardware, advanced system engineering, and early testbed platforms—to support the nation's exascale computing imperative.

APPENDIX A. DETAILS OF PLENARY SESSIONS

The 2020 ECP annual meeting offered an ECP Welcome plenary session on Monday, February 2, 2020, in which ECP Leadership updated the community as to the state of ECP. In addition, the program included four plenary sessions over the course of the week covering topics of interest to the audience.

See list of plenary sessions below.

- AI for Science
- ASCAC Transition of Exascale Computing Project Report
- EQSIM: Transforming Earthquake Hazard and Risk Assessment through Exascale Simulations
- Industry Imperative for Exascale

TITLE **AI for Science**

FOCUS AREA ---
TYPE Plenary
LEVEL N/A

DAY Monday, February 3, 2020
TIME 6:30 PM – 7:30 PM
DURATION 60 min

AUTHORS
Rick Stevens (Argonne National Laboratory / The University of Chicago) stevens@anl.gov

ABSTRACT
In this talk, I will describe an emerging DOE initiative to advance the concept of Artificial Intelligence (AI) aimed at addressing challenge problems in science. We call this initiative "AI for Science". The basic concept is threefold: (1) to identify those scientific problems where existing AI and machine learning methods can have an immediate impact (and organize teams and efforts to realize that impact); (2) identify areas of where new AI methods are needed to meet the unique needs of science research (frame the problems, develop test cases, and outline work needed to make progress); and (3) to develop the means to automate scientific experiments, observations, and data generation to accelerate the overall scientific enterprise. Science offers plenty of hard problems to motivate and drive AI research, from complex multimodal data analysis, to integration of symbolic and data intensive methods, to coupling large-scale simulation and machine learning to drive improved training to control and accelerate simulations. A major sub-theme is the idea of working toward the automation of scientific discovery through integration of machine learning (active learning and reinforcement learning) with simulation and automated high-throughput experimental laboratories. I will provide some examples of projects underway and lay out a set of long-term driver problems.

TITLE **ASCAC Transition of Exascale Computing Project Report**

FOCUS AREA ---
TYPE Plenary
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 12:30 PM - 1:30 PM central
DURATION 60 min

AUTHORS
Roscoe Giles (Boston University) roscoe@bu.edu

ABSTRACT
The ASCR Advisory Committee (ASCAC) is charged by the DOE Office of Science to identify elements of ECP that should be transitioned into ASCR research or other SC/ASCR initiatives. The ASCAC "Transition Subcommittee" was convened to study the situation and develop a detailed suggested response to this charge. The report will provide guiding strategies and approaches that will be key to ensuring future U.S. leadership across the full range of disciplines stewarded by ASCR. In this talk, I will provide an update on the activities of the Transition Subcommittee and the status and content of the report provided to ASCAC.

TITLE **EQSIM – Transforming Earthquake Hazard and Risk Assessment Through Exascale Simulations**

FOCUS AREA **AD**
TYPE Plenary
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 12:30 PM - 1:30 PM central
DURATION 60 min

AUTHORS
David McCallen (Lawrence Berkeley National Laboratory) dbmccallen@lbl.gov
Anders Petersson (Lawrence Livermore National Laboratory) andersp@llnl.gov
Arthur Rodgers (Lawrence Livermore National Laboratory) rodgers7@llnl.gov
Marmun Miah (Lawrence Berkeley National Laboratory) mmiah@lbl.gov
Arben Pitarka (Lawrence Livermore National Laboratory) pitarka1@llnl.gov
Houjun Tang (Lawrence Berkeley National Laboratory) htang4@lbl.gov

ABSTRACT
Large earthquakes present a significant risk around the world and are a major issue across the DOE mission space ranging from the safety of DOE's own inventory of one-of-a-kind mission critical facilities to all major US energy systems (electric/gas distribution systems, renewable energy production facilities, nuclear power plants etc.). Beyond the DOE enterprise, addressing earthquake risk is a major world-wide societal challenge for virtually every element of the built environment including transportation, health, data/commerce and all urban infrastructure. The tremendous developments occurring in high performance computing will transform earthquake hazard and risk assessments. As computational power increases, the reliance on simplifying idealizations, approximations and legacy empirical methods can diminish, and attention can be focused on dealing with the fundamental physics uncertainties in earthquake processes. With the advent of Exascale computations, regional-scale ground motion simulations are becoming computationally feasible at frequencies of engineering interest, and simulation models that connect the domains of seismology, geotechnical and structural engineering are within grasp. The ECP EQSIM Application Development project is focused on creating an unprecedented computational toolset and workflow for earthquake hazard and risk assessment. EQSIM is building an end-to-end, fault-to-structure capability to simulate from the initial fault rupture to surface ground motions (earthquake hazard) and ultimately to infrastructure response (earthquake risk) with the ultimate goal of removing computational limitations as a barrier to scientific exploration and understanding of earthquake phenomenology. This presentation will describe the EQSIM framework and highlight the computational advancements achieved to-date, including key advanced algorithm development and transition to GPU-based platforms. In addition, the enabling interdependencies with ECP S&T software developments to fully develop the multidisciplinary fault-to-structure workflow will be described.

TITLE **Industry Imperative for Exascale**

FOCUS AREA ---
TYPE Plenary
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 12:30 PM - 1:30 PM central
DURATION 60 min

AUTHORS
Organizers
David Martin (Argonne National Laboratory) dem@alcf.anl.gov
Suzy Tichenor (Oak Ridge National Laboratory) tichenorsp@ornl.gov

Moderator
Brunon (Dave) Kepczynski (General Electric) Brunon.Kepczynski@ge.com

Panelists
Srikanth Mahalingam (Altair Engineering, Incorporated) sam@altair.com
Frank Ham (Cascade Technologies, Incorporated) ham@cascadetechnologies.com

ABSTRACT
Independent software vendor (ISV) executives on the ECP Industry Council will discuss how they are preparing their codes for emerging exascale systems as well as how they are preparing their customers. Since most industrial users do not write their own software, industry's ability to use exascale systems will be tied closely to ISV codes. Moderated by ECP Industry Council chair Brunon (Dave) Kepczynski, the panelists will discuss their strategy, hurdles they are facing, and tools and techniques from ECP they think will help them.

APPENDIX B. DETAILS OF BOF/BREAKOUT/PANEL SESSIONS

The 2020 ECP annual meeting offered a diverse series of 45, 60, and 90 minute Birds of a Feather (BoF)/ breakout/panel sessions that were held on Tuesday, Wednesday, and Thursday of the annual meeting. Participants were encouraged to attend breakouts to learn more about important topics and areas in the ECP community. Presenters discussed accomplishments, requirements sharing, and opportunities for interactions.

See the list of BoF/Breakout/Panel sessions below.

- Advances in Science and Engineering Enabled by ECP Applications (Part I)
- Advances in Science and Engineering Enabled by ECP Applications (Part II)
- AMReX and AMReX-Based Applications
- App Integration at the Facilities: Lessons Learned so Far
- Aurora – Argonne's Exascale System (NDA Session)
- Best Practices for Using Proxy Apps as Benchmarks
- CEED: High-Order Methods, Applications, and Performance for Exascale
- Checkpointing at Exascale: Software Integration, Use Cases, Challenges
- CMakePP: Automating and Simplifying CMake Build Systems
- Collaboration Opportunities with the BSSw Fellows
- Container Utilization at DOE Compute Facilities
- Cultivating Software Sustainability and Productivity through BSSw.io
- DAOS (Distributed Asynchronous Object Storage) for Applications
- Early Experience of Application Developers with OpenMP Offloading at ALCF, NERSC, and OLCF
- ECP 201: Lessons Learned in Preparing Software Technology and Applications for Exascale Systems
- ECP and Experimental Science
- ECP Center and Application Monitoring - Working Group
- ECP Continuous Integration Update
- ECP Software Curation through Deployment
- ECP SW Deployment Project Team Breakout
- Hardware Evaluation for Understanding Network Requirements from Facilities and Software Teams
- I/O Performance Addicts
- Intel Developer Feedback Session
- Intelligent Distributed Data Movement for Exascale Computing
- Kokkos Ecosystem for Performance Portability: Updates and Plans
- Lossy Data Reduction/Compression for ECP Applications
- Management of Complex, Heterogeneous, and Hierarchical Memory Resources
- MPI BoF
- Multiprecision Focus Effort in the ECP Math Library Ecosystem
- OLCF ECP Application Engagement
- OpenMP Roadmap for Accelerators across DOE Pre-Exascale/Exascale Machines
- Overview of the Argonne Aurora Exascale System
- Overview of the National Security Applications
- PathForward Update – AMD (NDA Session)
- PathForward Update – Cray (NDA Session)

- PathForward Update – HPE (NDA Session)
- PathForward Update – IBM (NDA Session)
- PathForward Update – Intel (NDA Session)
- PathForward Update – NVIDIA (NDA Session)
- Performance Portability at the Intersection of the Exascale Computing Project and the DOE Centers of Excellence
- Perlmutter - a Waypoint for ECP Teams
- Persistent Memory in the Exascale Memory/Storage Continuum
- Preparing Applications for Aurora
- Slingshot BoF (NDA Session)
- Spack Roundtable
- Speed Dating for ECP: 1-on-1 Conversations between AD Teams and Math Library Developers
- SUNDIALS Plans and Uses
- SYCL Programming Model for Aurora
- System Software - Evaluation and Future for Pre-Exascale and Exascale Systems
- The Compiler Will Help You! The LLVM Compiler Infrastructure in ECP
- The Scope and Role of DevOps in HPC
- Tools and Techniques for Application Performance Characterization
- Understanding Performance with Exa-PAPI

TITLE **Advances in Science and Engineering Enabled by ECP Applications (Part I)**

FOCUS AREA **AD**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS

Chair of Session
Thomas Evans (Oak Ridge National Laboratory) evanstm@ornl.gov

Speakers
Jean-Luc Vay (Lawrence Berkeley National Laboratory) jlvay@lbl.gov
Amitava Bhattacharjee (Princeton Plasma Physics Laboratory) amitava@pppl.gov
Jackie Chen (Sandia National Laboratories) jhchen@sandia.gov
Madhava Syamlal (National Energy Technology Laboratory) madhava.syamlal@netl.doe.gov

ABSTRACT

Four ECP Application Development projects temporarily turn their focus away from code development and optimization to highlight recent scientific and engineering accomplishments in their respective domains:

- Jean-Luc Vay: WarpX: Exascale Modeling of Advanced Particle Accelerators
- Amitava Bhattacharjee: WDMApp: High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas
- Jackie Chen: Combustion-Pele: Transforming Combustion Science & Technology with Exascale Simulations
- Madhava Syamlal: MPIX-Exa: Perf Prediction of Multiphase Energy Conversion Device

TITLE **Advances in Science and Engineering Enabled by ECP Applications (Part II)**

FOCUS AREA **AD**
TYPE Breakout
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 8:30 AM - 10:00 AM central
DURATION 90 min

AUTHORS
Chair of Session
William Hart (Sandia National Laboratories) wehart@sandia.gov

Speakers
Nicholas Sauter (Lawrence Berkeley National Laboratory) nksauter@lbl.gov
John Gounley (Oak Ridge National Laboratory) gounleyjp@ornl.gov
Shrirang Abhyankar (Pacific Northwest National Laboratory) shri@pnnl.gov
Lenny Oliker (Lawrence Berkeley National Laboratory) loliker@lbl.gov

ABSTRACT
Four ECP Application Data Analytics and Optimization projects temporarily turn their focus away from code development and optimization to highlight recent scientific and engineering accomplishments in their respective domains:

- **Nicholas Sauter**: The potential for new proteing crystallographic science with exascale computing
- **John Gounley**: CANDLE and deep learning for cancer surveillance
- **Shrirang Abhyankar**: Optimizing atochastic grid dynamics at exascale
- **Lenny Oliker**: Exabiome: large scale genomic and proteomic analysis

TITLE **AMReX and AMReX-Based Applications**

FOCUS AREA **AD**
TYPE Breakout
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Ann Almgren (Lawrence Berkeley National Laboratory) asalmgren@lbl.gov
John Bell (Lawrence Berkeley National Laboratory) jbbell@lbl.gov
Kevin Gott (Lawrence Berkeley National Laboratory) kngott@lbl.gov
Andrew Myers (Lawrence Berkeley National Laboratory) atmyers@lbl.gov

ABSTRACT
AMReX is a software framework -- supported by the AMReX Co-Design Center -- that supports structured mesh algorithms for the exascale; in particular, those using block-structured adaptive mesh refinement (AMR). In addition to the conventional representation of field variables on a mesh, AMReX supports particle data and embedded boundary (cut cell) representations of complex geometries.

All of the functionality in AMReX is available on NVIDIA GPUs as well as CPUs, with support for AMD GPUs underway. AMReX emphasizes both performance and portability, making it as easy as possible for the application teams to use different architectures without sacrificing algorithmic flexibility and specialized performance optimizations. Key features of the GPU implementation in AMReX include

- Default to managed memory, but make other options available.
- Allow applications to use desired offloading techniques (e.g. CUDA, OpenMP, OpenACC) with user-friendly APIs for offloading
- Switch between CPU and GPU with a compile-time flag

In this breakout session we will cover the latest developments in AMReX as well as the current AMReX-based applications. We will focus on what the applications have found to be the most effective strategies for optimal performance, and discuss plans for how to take best advantage of upcoming architectures.

DESCRIPTION
This session targets primarily the six AD projects that currently include codes based on AMReX, but is open to any other interested AD or ST projects as well as interested vendors. The level will be intermediate -- we are anticipating mostly experienced users of AMReX.

The primary goals are:

1) to share information about future plans -- for both new capability and new architectures -- with application teams,
2) to share best practices between AMReX and the apps,
3) to hear from apps about their concerns as well as current and future needs.

AMReX is a software framework -- supported by the AMReX Co-Design Center -- that supports structured mesh algorithms for the exascale; in particular, those using block-structured adaptive mesh

refinement (AMR). In addition to the conventional representation of field variables on a mesh, AMReX supports particle data and embedded boundary (cut cell) representations of complex geometries.

All of the functionality in AMReX is available on NVIDIA GPUs as well as CPUs, with support for AMD GPUs underway. AMReX emphasizes both performance and portability, making it as easy as possible for the application teams to use different architectures without sacrificing algorithmic flexibility and specialized performance optimizations. Key features of the GPU implementation in AMReX include

- Default to managed memory, but make other options available.
- Allow applications to use desired offloading techniques (e.g. CUDA, OpenMP, OpenACC) with user-friendly APIs for offloading
- Switch between CPU and GPU with a compile-time flag

In this breakout session we will cover the latest developments in AMReX as well as the current AMReX-based applications. We will focus on what the applications have found to be the most effective strategies for optimal performance, and discuss plans for how to take best advantage of upcoming architectures.

TITLE **App Integration at the Facilities: Lessons Learned so Far**

FOCUS AREA **HI**
TYPE BoF
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Deborah Bard (Lawrence Berkeley National Laboratory) djbard@lbl.gov
Judith Hill (Oak Ridge National Laboratory) hilljc@ornl.gov
Katherine Riley (Argonne National Laboratory) riley@alcf.anl.gov

ABSTRACT
ECP app integration at the facilities is progressing rapidly, with engagements formed and facility teams already working closely with AD teams. This BOF is an opportunity for the facility app integration staff to discuss their experiences in working with their target architectures, sharing best practices, lessons learned and identifying new areas for collaboration.

DESCRIPTION
In this BOF, facility app integration staff will have the opportunity to share knowledge and best practices regarding app integration with their respective architectures. The format will involve a series of 5-min talks from selected facility staff, outlining some key technical achievements and lessons learned. This will be followed by general discussion of facility-focussed technical topics, in which the participants will share their experiences on topics including software, compiler integration and system management.

The target audience is the facility staff working on ECP apps, although we expect that this BOF will also be of great interest to the AD teams that are working with the facility experts.

TITLE **Aurora – Argonne's Exascale System (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 4:30 PM - 5:30 PM central
DURATION 60 min

AUTHORS
Susan Coghlan (Argonne National Laboratory) smc@anl.gov
Mark Fahey (Argonne National Laboratory) mfahey@anl.gov
Kevin Harms (Argonne National Laboratory) harms@alcf.anl.gov
Chris Knight (Argonne National Laboratory) knightc@anl.gov
Kalyan Kumaran (Argonne National Laboratory) kumaran@anl.gov
Ti Leggett (Argonne National Laboratory) tleggett@anl.gov
Mike Papka (Argonne National Laboratory) papka@anl.gov
Scott Parker (Argonne National Laboratory) sparker@anl.gov
Jini Ramprakash (Argonne National Laboratory) jini@alcf.anl.gov
Katherine Riley (Argonne National Laboratory) riley@alcf.anl.gov
Venkat Vishwanath (Argonne National Laboratory) venkat@anl.gov
Tim Williams (Argonne National Laboratory) tjwilliams@anl.gov

ABSTRACT
This session will provide an in depth hardware overview of the Aurora exascale system to be delivered to the Argonne Leadership Computing Facility in 2021. The session will include RSNDA material and attendees will be limited to those on the A21 CITR. Pre-registration for the session is required in order to provide an opportunity to verify that attendees are covered by an appropriate RSNDA and are on the A21 CITR.

DESCRIPTION
Talks will focus on the hardware features of Aurora including in-depth presentations on Intel Xeon, Xe architectures and Cray Slingshot network. Also included in the presentation will be performance estimates for a number of applications and benchmarks in the simulation, data and learning space.

TITLE **Best Practices for Using Proxy Apps as Benchmarks**

FOCUS AREA **AD**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
David Richards (Lawrence Livermore National Laboratory) richards12@llnl.gov
Jeanine Cook (Sandia National Laboratories) jeacook@sandia.gov
Hal Finkel (Argonne National Laboratory) hfinkel@anl.gov
Oscar Hernandez (Oak Ridge National Laboratory) oscar@ornl.gov

ABSTRACT
Proxy applications have many uses in software development and hardware/software co-design. Because most proxies are easy to build, run, and understand, they are especially appealing for use in benchmark suites and studies. This session will examine the role of proxies as benchmarks and discuss what is necessary for a proxy application to function as an effective benchmark. We will present performance data that demonstrates differences between good and bad benchmarks. We will also discuss what facilities are looking for when they assemble benchmark suites for use in procurements as well as case studies for what has worked well and what has been problematic. System vendors will provide their perspective on how they use our benchmark suites and what practices they view as most (or least) effective. Most importantly, we will request feedback from attendees on how to improve the usefulness of the ECP Proxy Apps for benchmarking.

DESCRIPTION
Because of their simplicity and easy availability, proxy apps are routinely used in benchmarking activities. Proxy apps are used as benchmarks in numerous contexts including formal benchmark suites associated with DOE system procurements, performance tests in published papers, and in technology evaluations from hardware vendors. Unfortunately, proxies are not always well chosen or even suitable for all of the benchmark studies in which they are used. The purpose of this session is to highlight best practices that should apply to specifying benchmarks involving proxy apps and expose errors that can occur when proxies are used carelessly.

The session will cover the following topics:

- All benchmarks are proxies, but not all proxies are good benchmarks.
- Evaluating the fidelity of proxies as a model for the parent application, including how fidelity can vary according to the problem specification.
- How to specify problems & parameters to make useful benchmark problems, and how this can go wrong.
- How facilities assemble benchmark suites and the considerations for what is and is not included.
- How vendors view our benchmark suite including what is and is not effective.

The overall goal is twofold:

- To help proxy authors understand how they can write & document proxies to increase the chances that benchmarking activities are useful and meaningful.

- To help proxy users understand how to select and configure proxy apps as benchmarks to maximize the value of the information obtained.

TITLE **CEED: High-Order Methods, Applications, and Performance for Exascale**

FOCUS AREA **AD**
TYPE Breakout
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 4:00 PM - 5:30 PM central
DURATION 90 min

AUTHORS
Tzanio Kolev (Lawrence Livermore National Laboratory) tzanio@llnl.gov
Jed Brown (University of Colorado Boulder) jed@jedbrown.org
Paul Fischer (University of Illinois at Urbana-Champaign) fischerp@illinois.edu
Robert Rieben (Lawrence Livermore National Laboratory) rieben1@llnl.gov
Misun Min (Argonne National Laboratory) mmin@mcs.anl.gov
Veselin Dobrev (Lawrence Livermore National Laboratory) dobrev1@llnl.gov
Tim Warburton (Virginia Polytechnic Institute and State University) tim.warburton@gmail.com
Stan Tomov (The University of Tennessee, Knoxville) tomov@icl.utk.edu
Mark Shephard (Rensselaer Polytechnic Institute) shephard@rpi.edu

ABSTRACT
The session will combine an update on the progress in the CEED co-design center with one-on-one discussions between CEED experts and applications interested in massively parallel high-order methods. The progress update (first half of the session) will discuss recent algorithmic and application developments in CEED-developed software (MFEM, Nek, libCEED, libParanumal, BPs), with a particular focus on GPU acceleration. During the second half of the session, attendees will have the chance to schedule one-on-one meetings with CEED researchers to focus on their particular application needs. We are interested in getting feedback from the whole ECP community, so please join us if you'd like to learn more about high-order methods, or have questions for the CEED team. For more details, see our website, https://ceed.exascaleproject.org, and GitHub repositories, https://github.com/ceed.

TITLE **Checkpointing at Exascale: Software Integration, Use Cases, Challenges**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 3:30 PM - 4:15 PM central
DURATION 45 min

AUTHORS
Bogdan Nicolae (Argonne National Laboratory) bnicolae@anl.gov
Franck Cappello (Argonne National Laboratory) cappello@anl.gov
Kathryn Mohror (Lawrence Livermore National Laboratory) mohror1@llnl.gov

ABSTRACT
Efficient checkpoint-restart is a capability that many ECP applications need. However, they rely on custom solutions that are burdened by many issues: (1) performance and scalability of the underlying parallel file system where the checkpoints need to be saved in a reliable fashion; (2) resource-efficiency concerns (e.g., extra memory overhead) needed to achieve high performance and scalability; (3) complexity of interacting with a heterogeneous storage hierarchy on future Exascale machines (e.g., many vendor-specific external storage options like burst buffers with custom APIs) that leads to loss of productivity. The ECP VeloC project aims to solve these issues through a specialized middleware that aims to deliver high performance and scalability with minimal overhead on resource utilization, while exposing a simple API that hides the complexity of interacting with a heterogeneous storage layer.

Over the last year, several ECP application groups (HACC, LatticeQCD, EXAALT, NWChemEx) had been experimenting with VeloC in various checkpoint-restart scenarios tailored to their specific application. This breakout session proposes to share these experiences from an application-level point of view with the rest of the ECP community, in the hope of raising awareness of what benefits can be expected from VeloC and what is the best way to integrate with it.

We plan to include a series of talks given by several application teams. Topics will include two aspects. First, it will discuss best practices to integrate with VeloC (memory-based mode vs. file based-mode, how to use checkpoints for both resilience and other purposes such as analytics and visualization, how to develop a plugin-able framework to separate checkpointing from the main code). Second, it will discuss the use of VeloC both for resilience and beyond: analytics and visualization, suspend-resume to span across multiple reservation and pre-emption of batch jobs, productive checkpointing for applications that need to revisit and/or return to previous states datasets.

TITLE **CMakePP: Automating and Simplifying CMake Build Systems**

FOCUS AREA **AD**
TYPE BoF
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 3:30 PM - 4:15 PM central
DURATION 45 min

AUTHORS
Ryan Richard (Ames Laboratory) rrichard@ameslab.gov
Theresa Windus (Ames Laboratory / Iowa State University) twindus@iastate.edu

ABSTRACT
CMake has become the de facto standard for writing build systems for C++ programs; however, CMake's generality can make it difficult to write robust and reliable build systems. We have created a suite of CMake modules we call CMakePP which focus on automation and boilerplate reduction. Using CMakePP: dependencies can be downloaded, built, and installed with as little as their URLs; libraries and executables can be added simply by providing the path to the directory containing the source files; and packaging files are automatically generated for installed targets.

DESCRIPTION
C++ lacks a standardized build system and package manager, which in turn makes it difficult to create reusable C++ libraries and packages. Using CMake it is possible to create robust, reusable C++ packages, but doing so is hindered by CMake's steep learning curve. Our own experience with CMake shows that most of these considerations are boilerplate. CMakePP wraps this boilerplate in user-friendly functions that require minimal input. This session is targeted at application developers who want to use CMake as their project's build system, but do not want to spend the hours required to learn the intricacies of the CMake language. The goal of this session is to introduce CMakePP to a wider audience and gauge community interest in the CMakePP project.

TITLE **Collaboration Opportunities with the BSSw Fellows**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Hai Ah Nam (Los Alamos National Laboratory) hnam@lanl.gov
Rene Gassmoeller (University of California, Davis) rgassmoeller@ucdavis.edu
Ignacio Laguna Peralta (Lawrence Livermore National Laboratory) ilaguna@llnl.gov
Tanu Malik (DePaul University) tanu@cdm.depaul.edu
Kyle Niemeyer (Oregon State University) Kyle.Niemeyer@oregonstate.edu

ABSTRACT
Better Scientific Software (BSSw) does not happen without learning and implementing BETTER – Planning, Development, Performance, Reliability, Collaboration, and Skills.

Meet the 2019 BSSw Fellows and discover how their work will help your scientific software project be BETTER by:

1) Learning best practices for developing research software [Better Skills],

2) Reducing technical debt through reproducible containers [Better Development],

3) Navigating the social challenges in the evolution of a scientific software project [Better Collaboration], and

4) Improving the reliability through analyzing and debugging floating-point software [Better Reliability].

In this breakout, you will hear from four practitioners and leaders of better scientific software practices, the 2019 BSSw Fellowship Awardees, with time for engagement to discuss your project needs.

DESCRIPTION

Topic details

BSSw Fellows are chosen annually to develop a funded activity that promotes better scientific software, such as organizing a workshop, preparing a tutorial, or creating content to engage the scientific software community. Their expertise is diverse and they are often leaders in their own community; however, no pre-existing connections are required to DOE and the Exascale Computing Project community.

2019 Fellows are:

- **Rene Gassmoeller** (University of California, Davis):
  Rene has significant experience as a scientific software maintainer for the geodynamics community. He is using his expertise and experience to create guides and tutorials about managing the community around a scientific software project, including navigating the social aspects.

- **Ignacio Laguna Peralta** (Lawrence Livermore National Laboratory):
  Ignacio has significant experience in programming models and systems, and has been working on tools for the Sierra HPC system at LLNL, #2 on the Top500. He has created tutorials on tools to debug and improve the reliability of scientific applications, with emphasis on floating-point software defects what are difficult to debug.

- **Tanu Malik** (DePaul University):
  Tanu received a 2019 NSF Career award for her work on computational reproducibility and leads a group focused on developing methods and system for improving reproducible science in computational and data science disciplines. She is teaching how to reduce technical debt through using container technology and contributing to Trusted Tools.

- **Kyle Niemeyer** (Oregon State University):
  Kyle leads a research group focused on computational fluids modeling and teaches a 10-week course on research software development for more trustworthy, reusable and sustainable scientific software. He is condensing his course for intense hands-on bootcamps suitable for the ECP community.

Fellows are chosen for their potential impact on the ECP and broader community toward improving scientific software and overall productivity. During the 90-minute breakout session, each fellow will present their BSSw Fellowship work in a 15-minute presentation. The remaining time will be used to facilitate discussion and interaction between the Fellows and ECP community members in smaller groups.

Goals
The intent of this session is to create an opportunity for ECP projects and members to learn about better scientific software practices and engage with the 2019 BSSw Fellows to foster potential collaborations and integrate the BSSw Fellows into the ECP community. This will broaden the impact of the Fellows' work and provide ECP community members with additional resources to improve their scientific software development practices. BSSw Fellows will benefit from learning about DOE & ECP project needs to better target their work.

Target audience
All ECP annual meeting members can benefit from this session. To ensure engagement, we will send invitations to identified ECP project members we believe can benefit from the work performed by the 2019 BSSw Fellows.

TITLE **Container Utilization at DOE Compute Facilities**

FOCUS AREA **ST**
TYPE Panel
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Andrew Younge (Sandia National Laboratories) ajyoung@sandia.gov
Shane Canon (Lawerence Berkeley National Laboratory) scanon@lbl.gov
Thomas Mendoza (Lawerence Livermore National Laboratory) mendoza33@llnl.gov
Jason Kincl (Oak Ridge National Laboratory) kincljc@ornl.gov
Anthony Agelastos (Sandia National Laboratories) amagela@sandia.gov
David Montoya (Trenza Synergy) dmont@trenzasynergy.com

ABSTRACT
Containerized application and services have gained adoption within the DOE computing community. Many of the ASCR and NNSA facilities offer solutions and services to run containzerized HPC applications and edge services. However, the level of support and nature of the offerings could vary greatly at Exascale. This panel will bring together representatives from across the DOE computing facilities to discuss their current portfolio, future plans, adoption and gaps for user-defined container and services computing. The panel will solicit questions from the audience to hear from the community about challenges, successes and feedback.

DESCRIPTION
Container computing has revolutionized the way groups are developing, sharing, and running software and services. This has initially been led by growth with Docker, which has provided an ecosystem of tools to enable container based computing. This paradigm shift has since made in-roads in the HPC community enabled by container runtimes like Podman, Singularity, Shifter, and Charliecloud, which allow end-users to run containers in environments where standard Docker tools would not be feasible. Orchestration systems such as Kubernetes also allow for quick, scalable service deployments in conjunction with HPC. While this adoption is growing, the larger ECP community is likely to have several questions regarding containers at Exascale. Furthermore, there still exist several undefined questions related to implementing containerization at various DOE facilities that the community address for successful deployments for of containerized HPC applications and services.

This BOF will bring together key experts in the field of containerization and HPC to describe and share their experiences, tools, and knowledge and engage in interactive discussions with the audience on this topic. Specific focus will given to facilities plans to integrate containers into future Exascale platforms. Given the recent increase in interest surrounding containers, we expect to attract a wide group of ECP participants, ranging from facilities and system administrators, to ST project leads and even DOE application developers.

This panel session will outlay an initial overview of containers and their planned integration at the various DOE E6 laboratories. First, we will introduce the Supercontainers ECP ST project and some of the planned integration activities. Next, we will introduce the panel members individually, who will give a short lightning talk describing a quick overview of each representative site's current and planned usage models for containers. This will potentially describe the current state of practice at each HPC site, the

application of HPC container runtimes, and the potential plans and use cases for user-based service container computing.

Following site introductions, the panel session will assemble from the invited participants. This panel will be steered by a moderator who will drive an interactive discussion on containers, and services computing. We envision the bulk of the Panel session to be devoted to lively, constructive, and useful discussion between the panel and audience. The moderator will initially seed discussion topics, but then transition to show-of-hand survey questions and other tools to encourage the audience to participate in the discussion. Specifically, some key questions the panel will discuss include:

- How a given facility supports (or will support) user-generated HPC container images.
- How ECP and individual facilities can leverage current container efforts.
- Future plans for providing containerized services to support user HPC computation.
- What challenges still exist for utilizing containers on future DOE Exascale systems.

We expect this discussion will enable the wider audience to not only become more familiarized with container capabilities, but also help identify key target integration points that exist across ECPapplications, SDKs, and various Exascale facilities.

TITLE **Cultivating Software Sustainability and Productivity through BSSw.io**

FOCUS AREA **HI**
TYPE BoF
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Rinku Gupta (Argonne National Laboratory) rgupta@anl.gov
BSSw Editorial Board (Oak Ridge National Laboratory)
Selected BSSw contributors from the ECP community (Lawrence Livermore National Laboratory)

ABSTRACT
Software – the foundation of discovery in computational science and engineering – faces increasing complexity in computational models and computer architectures. As the ECP community work toward scientific advances, we also need to continually improve software practices in order to more effectively collaborate across teams in AD, ST, HI, and DOE computing facilities. Moreover, we need to share information specifically targeted for high-performance scientific software.

The Better Scientific Software site (https://bssw.io) provides a central hub to share information on practices, techniques, experiences, and tools to improve developer productivity and software sustainability, as key aspects of increasing overall scientific productivity. BSSw.io features ever-expanding resources on scientific software planning, development, performance, reliability, collaboration, and skills, including introductory WhatIs and HowTo information, curated content (brief articles that highlight other web-based materials), and original articles on science teams' experiences with software issues. A GitHub backend enables content development using a collaborative, open workflow; content can also be contributed with an easy-to-use Google form.

In this BOF session, the ECP community will share experiences on using and contributing to BSSw.io. An objective is open discussion to understand pressing needs, challenges, and potential roadblocks for software sustainability and productivity, in order to determine priorities for future directions in site content, usage, and contributions. The long-term vision for BSSw is to serve as an international community-driven and community-managed resource, with content and editorial processes provided by volunteers. We welcome everyone to participate.

DESCRIPTION
The ECP focuses on advancement of computational science and engineering through extreme-scale computing. At the core of efforts in both applications and software technologies is the need to improve developer productivity, positively impacting product quality, development time, and staffing resources, and software sustainability, reducing the cost of maintaining, sustaining, and evolving software capabilities in the future. The IDEAS-ECP project is expressly addressing the challenges facing the ECP teams by focusing on improving how we conduct our software efforts.

The Better Scientific Software (BSSw) community is an international group of researchers, practitioners, and stakeholders who are dedicated to curating, creating, and disseminating information that leads to improved software for the ECP community as well as the broader CSE community. Central to this effort is the BSSw.io website, which aims to be community-based resource with content and editorial processes in place. The BSSw.io website serves as a central hub for sharing information on practices, techniques,

experiences, and tools to improve developer productivity and software sustainability. The goals of the site are to raise awareness of (1) the importance of good software practices to scientific productivity and to the quality and reliability of scientific results and (2) the increasing challenges facing CSE software developers as high-end computing heads to extreme scales.

Historically, opportunities for ECP software developers (and the broader CSE community) to exchange information and experiences have been limited; BSSw.io provides a space to support this kind of sharing. The BSSw portal provides easy access to resources and training materials provided by contributors from the ECP, DOE, and the broader high-performance computational science communities. The site content spans a range of topics including introductory WhatIs and HowTo documents that define terminology and basic processes for improving software productivity and sustainability, curated content - brief articles that highlight other web-based materials, describing why the scientific software community might find them of value, original content in the form of blogs, success stories, event highlights and original experience-based articles.

In the proposed BOF session, we hope to educate people on how to use as well as contribute to BSSw as a resource for advancing productivity and software sustainability within the ECP community. Since BSSw.io is created for the community, we want to solicit ECP community involvement for this resource. The session will incorporate a mix of BOF speakers, including people from the ECP community who have contributed to the site who will share their experiences as well as speak about their perspectives on the site. We will share our experiences and lessons learned about the status of software sustainability and productivity in ECP. One of our objectives is to encourage an open discussion with the ECP community members to understand their needs, challenges and potential roadblocks for software sustainability/productivity and determine the nature of BSSw resources that can help alleviate this pain.

TITLE **DAOS (Distributed Asynchronous Object Storage) for Applications**

FOCUS AREA **HI**
TYPE BoF
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 8:30 AM - 10:00 AM central
DURATION 90 min

AUTHORS
Johann Lombardi (Intel Corp.) mohamad.chaarawi@intel.com
Kevin Harms (Argpnne National Laboratory) harms@alcf.anl.gov
Mohamad Chaarawi (Intel Corp.) johann.lombardi@intel.com

ABSTRACT
DAOS (Distributed Asynchronous Object Store) is an open-source software-defined multi-tenant scale-out object store designed from the ground up for massively distributed Non Volatile Memory (NVM). DAOS will be the primary high-performance storage system on Aurora. This storage system, while being compatible with existing I/O models such as POSIX, MPI-IO and HDF5, will offer advanced features to improve I/O performance and workflows of applications. The session will focus on DAOS from an applications perspective providing information on several key features of DAOS, how to prepare applications and what to expect performance and functionality wise. Current results and a demonstration with existing applications and benchmarks will be presented.

DESCRIPTION
DAOS is a next generation open-source scale-out storage engine designed from the ground up for persistent memory and NVMe storage offering extremely high bandwidth and IOPS. DAOS resides completely in user space with full OS bypass making it a lightweight storage stack presenting a shift from traditional storage stacks that are designed for rotating storage media. DAOS servers maintain their metadata on persistent memory, with bulk data going straight to NVMe SSDs. In addition, small I/Os will be absorbed on the persistent memory before being aggregated and then migrated to the larger capacity flash storage. DAOS uses the Persistent Memory Development Kit (PMDK) to provide transactional access to the persistent memory and the Storage Performance Development Kit (SPDK) for user space I/O to NVMe devices.

DAOS presents a rich, scalable storage interface unconstrained by traditional POSIX limitations. In this new storage paradigm, POSIX is no longer the foundation for new data models. Instead, the POSIX interface is built as a library on top of the DAOS backend API, like any other I/O middleware. DAOS will be directly integrated with several data formats and middleware libraries (e.g. HDF5, MPI-IO, SEGY) and frameworks (e.g. Spark). For smooth migration of applications, DAOS supports a POSIX namespace encapsulation with relaxed compliance, a ROMIO backend for MPI-IO support and HDF5 support via special DAOS VOL connector.

This session will provide an overview of DAOS and its data model, and then review each of the existing I/O models and the current state of affairs regarding what is supported. Some results on prototype hardware using the DAOS API and other middleware libraries on top of DAOS will be presented. The BOF will focus on the key aspects of DAOS from an application perspective (data model, features for enhancing application workflows, performance, etc.) and provide an opportunity for application developers to engage with questions and their use cases and provide feedback on the current design. A

walk-through of how to use DAOS through a demo and a discussion of possible future methods to evaluate DAOS will be provided.

TITLE **Early Experience of Application Developers with OpenMP Offloading at ALCF, NERSC, and OLCF**

FOCUS AREA **HI**
TYPE BoF
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 4:00 PM - 5:30 PM central
DURATION 90 min

AUTHORS
Yasaman Ghadar (Argonne National Laboratory) ghadar@anl.gov
Johannes Doerfert (Argonne National Laboratory) jdoerfert@anl.gov
Jae Hyuk Kwack (Argonne National Laboratory) jkwack@anl.gov
Jose Monsalve Diaz (Argonne National Laboratory) jmonsalvediaz@anl.gov
Brian Homerding (Argonne National Laboratory) bhomerding@anl.gov
Tim Williams (Argonne National Laboratory) tjwilliams@anl.gov
Ray Loy (Argonne National Laboratory) rloy@alcf.anl.gov
Ye Luo (Argonne National Laboratory) yeluo@anl.gov
Colleen Bertoni (Argonne National Laboratory) bertoni@anl.gov
John Tramm (Argonne National Laboratory) jtramm@anl.gov
Brian Mackie-Mason (Argonne National Laboratory) bmackiemason@anl.gov
Jack Deslippe (Lawrence Berkeley National Laboratory) jrdeslippe@lbl.gov
Christopher Daley (Lawrence Berkeley National Laboratory) csdaley@lbl.gov
Rahul Gayatri (Lawrence Berkeley National Laboratory) rgayatri@lbl.gov
Thorsten Kurth (Lawrence Berkeley National Laboratory) tkurth@lbl.gov
Charlene Yang (Lawrence Berkeley National Laboratory) CJYang@lbl.gov
Brian Friesen (Lawrence Berkeley National Laboratory) bfriesen@lbl.gov
Doug Doerfler (Lawrence Berkeley National Laboratory) dwdoerf@lbl.gov
Pieter Ghysels (Lawrence Berkeley National Laboratory) pghysels@lbl.gov
Hadia Ahmed (Lawrence Berkeley National Laboratory) hahmed@lbl.gov
Oscar Hernandez (Oak Ridge National Laboratory) oscar@ornl.gov
Reuben Budiardja (Oak Ridge National Laboratory) reubendb@ornl.gov
Bronson Messer (Oak Ridge National Laboratory) bronson@ornl.gov
Gustav Jansen (Oak Ridge National Laboratory) jansengr@ornl.gov
Matthew Norman (Oak Ridge National Laboratory) normanmr@ornl.gov
Austin Harris (Oak Ridge National Laboratory) harrisja@ornl.gov
Piotr Luszczek (The University of Tennessee, Knoxville) luszczek@icl.utk.edu
Barbara Chapman (Stony Brook University) barbara.chapman@stonybrook.edu
Vivek Kale (Brookhaven National Laboratory) vkale@bnl.gov

ABSTRACT
The next generation of supercomputers at ALCF, NERSC, and OLCF will consist of heterogeneous CPU+GPU systems from three different hardware vendors. To target these systems, OpenMP is one possible programming model that can be used to take advantage of the massive parallelism available on GPUs. OpenMP is a portable programming model which will be supported at all three computing facilities, with the ability to run on GPUs from all three vendors. In this session, developers who are using OpenMP to target GPUs currently in their applications (spanning a variety of domains and including Fortran and C/C++) will give short talks about how they are using OpenMP, any lessons learned or best practices discovered, and any feedback they have for any of the current OpenMP implementations. From

these case studies, attendees can learn which features of the OpenMP specification current developers are using, any specific strategies they employed, and best practices developed. We are also looking for feedback on important OpenMP 5.0 features applications people plan to use or are using on current implementations.

DESCRIPTION

The plan for this session is to contain many short talks from a variety of developers (across ALCF, NERSC, and OLCF) about how they are using OpenMP in their applications. Since OpenMP is a programming model which will be portable across the upcoming systems at ALCF, NERSC, and OLCF, this topic is relevant to ECP application developers to give examples of how others are using OpenMP to target some or all of the upcoming exascale systems and to provide lessons learned and application best practices.

The goals of this session are: 1) to provide concrete examples of applications which are using OpenMP (in particular offloading, since ALCF, OLCF, and NERSC are all getting systems with GPUs) so that application developers who are interested in using OpenMP can learn which features of the OpenMP specification current developers are using, any specific strategies they employed, and best practices developed, and 2) allow developers to voice any feedback to vendors about the current OpenMP implementations. The target audiences are application developers who are curious about who is using OpenMP and how successfully, and vendors/tools developers to see how application developers are using OpenMP. We plan to include examples of application codes from a variety of domains and include both Fortran, and C/C++ codes.

TITLE **ECP 201: Lessons Learned in Preparing Software Technology and Applications for Exascale Systems**

FOCUS AREA ---
TYPE Breakout
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS

Suzy Tichenor (Oak Ridge National Laboratory) tichenorsp@ornl.gov
David Martin (Argonne National Laboratory) dem@alcf.anl.gov
Michael Heroux (Sandia National Laboratories) maherou@sandia.gov
Andrew Siegel (Argonne National Laboratory) siegela@mcs.anl.gov

ABSTRACT

With several years of progress behind us, AD and ST teams continue to push forward to develop applications and software technology needed for a robust exascale ecosystem. Even if an approach is successful, the idea of developing, building, testing, and deploying verified code on many different platforms in anticipation of Exascale systems is daunting. Leaders of the Software Technology and Application Development focus areas will discuss lessons learned from the trenches as teams have made progress in creating and adapting applications and software technology to be ready when exascale systems are installed.

DESCRIPTION

Target audience
Industry attendees and others that are interested in how ECP projects are approaching the architectures and programming environments of future exascale systems.

This breakout is intended to be a strategic/crosscutting set of talks on lessons learned across ST and AD projects that industrial attendees/developers can take away and use to guide their own internal efforts to prepare for exascale. Since the session will only be an hour and a half, it will naturally be at a somewhat higher level.

TITLE **ECP Center and Application Monitoring - Working Group**

FOCUS AREA **HI**
TYPE BoF
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 1:30 PM - 3:00 PM central
DURATION 90 min

AUTHORS
Terry Jones (Oak Ridge National Laboratory) trjones@ornl.gov
Oscar Hernandez (Oak Ridge National Laboratory) oscar@ornl.gov
David Montoya (Trenza Synergy) dmont@trenzasynergy.com
Jim Brandt (Sandia National Laboratory) brandt@sandia.gov
Jeanine Cook (Sandia National Laboratories) jeacook@sandia.gov
Matt Legrande (Lawrence Livermore National Laboratory) legendre1@llnl.gov
Michael Mason (Los Alamos National Laboratory) mmason@lanl.gov

ABSTRACT
Today, DOE computing centers track a wealth of information on the health, usage, and efficiency of our machines, workflows and programming environments. The information is critical to understanding observed performance and the early diagnosis of any issue. But would-be users of this data face significant challenges: without a clear game-plan the cost of the infrastructure to collect the data, store the data, and analyze the data makes many potential uses prohibitive. With many new potential sources of data in both the A21 exascale systems and CORAL2 exascale systems, what possible solutions provided by the various ECP centers can be used to collectively address existing challenges and emerging needs/wants to support exascale environments? We have pulled together a working group to address and will review current efforts across centers and discuss potential focus areas.

DESCRIPTION

Topic

There are many monitoring efforts across the DOE HPC centers that have evolved because of the need to collect data on performance and issues occurring on the large HPC machines and supporting environments. This includes a wealth of information on the health, usage, and efficiency of our machines, workflows and programming environments. The information is critical to understanding observed performance and the early diagnosis of any issue and optimize usage of these critical resources. Within the exascale ecosystem we are focusing on providing the software tools for applications to utilize to increase their ability to fully scale and effectively utilize the resources in the increasingly complex node architectures and complex runtimes. Monitoring and analytic frameworks allow for a window into this complexity and can provide information to better navigate and harness the potential power of these systems.

Goal and Target audience

This area is seen as a gap area in the ECP software ecosystem with no ST project focused on integration to support this need. As mentioned, the facilities are investing in this area but are addressing independently. We have established a working group to pull this gap area together, increase communication across sites and ECP projects and help establish a road map for going forward.

Target audience

Facilities, application teams, run time software teams.

TITLE **ECP Continuous Integration Update**
FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS

Ryan Adamson (Oak Ridge National Laboratory) adamsonrm@ornl.gov
David Montoya (Trenza Synergy) dmont@trenzasynergy.com
Kyle Shaver (Argonne National Laboratory) kshaver@anl.gov
Mario Melara (Lawrence Berkeley National Laboratory) mamelara@lbl.gov
Tiffany Connors (Lawrence Berkeley National Laboratory) tconnors@lbl.gov
David Nicholaeff (New Mexico Consortium) dnicholaeff@newmexicoconsortium.org
Paul Bryant (New Mexico Consortium) pbryant@newmexicoconsortium.org
Jeff Neel (Argonne Natinal Laboratory) jneel@anl.gov

ABSTRACT

Continuous Integration (CI) has become an industry standard, utilized by teams to better facilitate testing and integration of services and applications by using a robust, repeatable, and scalable process. The ECP CI project has been integrating features to the GitLab CI product and working with the DOE HPC Centers to bring this capability to ECP software teams and allow for testing across sites.

This past year has included adding features to add key features key to HPC centers such as userid-based job submission and ability to utilize HPC batch submission interfaces such as Slurm and LSF. In addition, security features for cross-site user authentication and control have been added to support job submission across sites. The HPC Centers have also identified and developed processes for CI access to many HPC resources. Our goal is to allow teams to integrate new workflows into their testing and deployment strategies.

Join the members of the CI team as we provide an overview and update of the ongoing efforts, a guide for project integration, and an open Q&A with the goal of providing a clear path forward for ECP teams interested in getting started.

DESCRIPTION

The ECP CI effort has been working for over a year in adding needed security features and added capability to the GitLab CI products to allow ECP ST and AD teams to test their software across E6 Facility HPC machines. This has been in development and not available to software teams until recently. This session is to provide an update and getting started guidelines to ECP users that are interested in getting started with ECP CI. This does not go into the detail that the CI tutorial session will go into but provide users with an update and status.

Topics

- ECP CI Project status & team
- Outline ongoing / open development efforts
- How the ECP CI ecosystem works
- Mirroring repository at OSTI
- Authentication and running at sites
- Connecting to training resources

Goals

- Update on project status and what resources are available
- Basic understanding of GitLab CI and how to get started
- Path forward and contacts provided to help get started after the event has concluded

Target Audience

- ECP AD/ST project teams interested in leveraging the ECP CI infrastructure through OSTI's GitLab instance.
- Anyone with an interest in leveraging GitLab CI either at their facility
- Interested in using ECP CI infrastructure
- Curious about CI in general

TITLE **ECP Software Curation through Deployment**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Frank Willmore (Argonne National Laboratory) willmore@anl.gov
Todd Gamblin (Lawrence Livermore National Laboratory) tgamblin@llnl.gov
Sameer Shende (University of Oregon) sameer@cs.uoregon.edu
Jim Willenbring (Sandia National Laboratories) jmwille@sandia.gov
Paul Bryant (New Mexico Consortium) pbryant@newmexicoconsortium.org
David Montoya (Trenza Synergy) dmont@trenzasynergy.com

ABSTRACT
Software deployment to facilities presents a unique set of challenges, particularly as exascale systems are one-of-a-kind and highly experimental. These challenges include heterogeneous compute environments, installation and maintenance of coherent software versions, cross-compilation, containerization, varied build systems, and ever-evolving software stacks. Curation of well-defined versions plays a crucial role both for the reproducibility of scientific outputs as well as for the focused maintenance of deployments to HPC facilities. In this session we will describe the layers of software curation through to deployment that will include; the SDK and how it establishes a community policy and interoperability for a segment of software, the Extreme-scale Scientific Software Stack (E4S) and how it serves as a release manifest of software product deployment recipes and packaged release in various container runtimes, and the software build pipelines both through the spack ecosystem and the facility software deployment infrastructures. This is expected to provide an enhanced deployment environment for facility software teams, application teams and support integration of container deployments. We invite submissions of both success stories and ongoing challenges that face both developers and operations teams as we adopt an ever-more DevOps approach to HPC.

DESCRIPTION

Topic
The ECP software ecosystem is complex with constant change occurring to support new additions to machine architectures and their supporting software stacks. In addition, software dependencies and versions are always being revised based on application testing and performance results. The software curation process is evolving out of the ECP project to better support those needs. Curation of well-defined versions plays a crucial role both for the reproducibility of scientific outputs as well as for the focused maintenance of deployments to HPC facilities.

Goal
In this session we will describe the layers of software curation through to deployment that will include; the SDK and how it establishes a community policy and interoperability for a segment of software, the E4S and how it serves as a release manifest of software product deployment recipes and packaged release in various container runtimes, and the software build pipelines both through the spack ecosystem and the facility software deployment infrastructures. This is expected to provide an enhanced deployment environment for facility software teams, application teams and support integration of container deployments.

Target Audience

- ECP Application teams
- ECP ST teams
- Facility infrastructure teams

TITLE **ECP SW Deployment Project Team Breakout**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 4:30 PM - 5:30 PM central
DURATION 60 min

AUTHORS
Ryan Adamson (Oak Ridge National Laboratory) adamsonrm@ornl.gov
David Montoya (Trenza Synergy) dmont@trenzasynergy.com
Kyle Shaver (Argonne National Laboratory) kshaver@anl.gov
Matt Belhorn (Oak Ridge National Laboratory) belhornmp@ornl.gov
Mario Melara (Lawrence Berkeley National Laboratory) mamelara@lbl.gov
Tiffany Connors (Lawrence Berkeley National Laboratory) tconnors@lbl.gov
Anthony Agelastos (Sandia National Laboratories) amagela@sandia.gov
Jennifer Green (Los Alamos National Laboratory) jgreen@lanl.gov

ABSTRACT
This is a breakout session for the ECP Software Deployment project, 2.4.4. The project team and partners will work through a status of current thrust areas some of which include CI implementation and project on-boarding, CI test resources, software deployment pipelines, software curation process, and facility plans. Plans will be worked through for project milestone targets and we will work through site plans and needs. The focus of the session will be transition from capability development that took up FY19 and early FY20 and the move to deployment in Continuous Integration (CI) and in the software deployment process that will feed off the software curation layers that are being developed in concert with the ECP ST software ecosystem and packaging projects. We expect to have some of our partners and stakeholders attend which include SDK, E4S, Spack, and Facility centered individuals.

DESCRIPTION
This is a breakout session for the ECP Software Deployment project, 2.4.4. The project team and partners will work through a status of current thrust areas some of which include CI implementation and project on-boarding, CI test resources, software deployment pipelines, software curation process, and facility plans.

Topics:

- Project team update - structure, etc.
- CI Thrust areas
  - CI framework and ecosystem
  - Project on-boarding
  - CI test resources
  - GitLab CI support issues
  - Facility plans
- Software Deployment thrust areas
  - SW Curation process
  - Build pipelines
  - Facility plans

Goals:

- Face-to-Face for team as it is distributed across the DOE facility sites.
- Update and plans and resourcing needs for milestones

Target audience:

- Project team under 2.4.4
- ECP ST partners from SDK, E4S, Spack
- Facility partners

TITLE **Hardware Evaluation for Understanding Network Requirements from Facilities and Software Teams**

FOCUS AREA **HI**
TYPE BoF
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Jeremiah Wilke (Sandia National Laboratories) jjwilke@sandia.gov
Scott Hemmert (Sandia National Laboratories) kshemme@sandia.gov
Misbah Mubarak (Argonne National Laboratory) mmubarak@anl.gov
Taylor Groves (Lawrence Berkeley National Laboratory) tgroves@lbl.gov
Scott Pakin (Los Alamos National Laboratory) pakin@lanl.gov

ABSTRACT
Building an exascale system is a complicated problem that requires both selecting the individual components and integrating them together into a system. System-level design concerns are often best-addressed through modeling and simulation to understand scaling problems before systems are procured and configured. Even for existing systems, simulation can extract detailed metrics and statistics that would be difficult or impossible to extract from a test bed. Here we discuss system-level design problems, primarily in the interconnect, being addressed in the hardware evaluation working groups. These challenges range from topology, routing, congestion management, QoS, network protocols, collective algorithms, or even programming models. We also identify ways to improve requirements gathering between hardware evaluation, facilities, and software technology teams.

DESCRIPTION
The proposed topic covers work from the Hardware Evaluation (HE) team within ECP, particularly the interconnects group. The target audience is broad and is intended to include application developers, facilities, and vendors. The goal of the session is both generally informing the ECP community about the hardware evaluation project while also receiving feedback and requirements from application, runtime, and vendor teams. The session will include short panel presentations (25 minutes) describing the tools used while also identifying gaps where more input is needed from facilities and apps. The remaining time will be devoted to Q&A and audience feedback on the challenges identified in the panel presentations. The interactive session will clarify to a broad audience what simulation can and can not provide while also seeking feedback from the audience on how to better connect simulation-based analysis with facilities and application requirements.

Specific tools to be discussed will be CODES, SST, and Conceptual. The presentations should illustrate the high-level design tradeoffs and basic usage of the tools. Results from recent ECP milestones exploring topology and routing, quality-of-service, and performance analysis will be presented.

TITLE **I/O Performance Addicts**

FOCUS AREA **ST**
TYPE BoF
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Shane Snyder (Argonne National Laboratory) ssnyder@mcs.anl.gov
Philip Carns (Argonne National Laboratory) carns@mcs.anl.gov
Rob Latham (Argonne National Laboratory) robl@mcs.anl.gov
Rob Ross (Argonne National Laboratory) rross@mcs.anl.gov

ABSTRACT
Achieving high I/O performance on ECP platforms can be challenging. Darshan (https://www.mcs.anl.gov/research/projects/darshan/) is an I/O characterization tool that is installed by default on most platforms, providing an easy way to gain insight into I/O behavior. In this BoF we will demonstrate how to find Darshan I/O characterization output on ECP platforms and how to interpret this data to understand the I/O behavior of application codes. Through examples we will highlight good and bad behavior and suggest remedies for common mistakes. We'll also discuss advanced features of Darshan that can be used to further dig into performance issues.

DESCRIPTION
We present a few "big" I/O tutorials each year, including typically material at SC and ATPESC. Basically what we're proposing to do here is extract key materials from the darshan components of those full-day talks and specialize this material to cover exactly what ECP teams would need to know to use this data on current platforms.

TITLE **Intel Developer Feedback Session**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 5:30 PM - 6:30 PM central
DURATION 60 min

AUTHORS

Facilitator
Jeff Hammond (Intel Corporation) jeff.r.hammond@intel.com

Technical Experts
Greg Stoner (Intel Corporation) gstoner@intel.com
Xinmin Tian (Intel Corporation) xinmin.tian@intel.com
John Pennycook (Intel Corporation) john.pennycook@intel.com

ABSTRACT

Intel invites ECP annual meeting registrants to a *Developer Feedback* session on Tuesday, February 4, from 5:30 pm - 6:30 pm central. We want to ensure the programming models and runtime environment (including oneAPI) planned for Argonne's Aurora meet the needs of ECP developers.

We will have several technical experts present to host the session, hear your feedback, and answer any questions you may have about our plans.

The session will be hosted in the Legends Ballroom. This event is open to the broad ECP meeting attendance. We look forward to seeing you there!

If you have any questions, please email **Sharon Hall** at sharon.j.hall@intel.com.

TITLE **Intelligent Distributed Data Movement for Exascale Computing**

FOCUS AREA **AD**
TYPE BoF
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Alex Sim (Lawrence Berkeley National Laboratory) asim@lbl.gov
Chin Guok (Lawrence Berkeley National Laboratory) chin@es.net
Ian Foster (Argonne National Laboratory) foster@anl.gov
Salman Habib (Argonne National Laboratory) habib@anl.gov

ABSTRACT
Applications that use exascale computing must frequently move data to and from the compute resources, for example to ingest data from scientific instruments, perform remote analysis, or transfer data to remote storage. With the exponential growth in data sets and requirements for integrated workflows exploiting distributed resources, the seamless movement of data between geographically diverse facilities is becoming essential to the efficiency of end-to-end workflows requiring exascale computing. The goal of this BoF is to examine use cases and present some solutions for intelligent data movement to support exascale computing. A large portion of the BoF will be dedicated to discussing the data flow requirements of large scale science applications that utilize exascale computing as part of a larger workflow.

DESCRIPTION
The focus of this BoF is to understand the data movement requirements (e.g., traffic volume, dataset characteristics (e.g., many small files, a few very large files), transfer characteristics (e.g., deadline driven), frequency, rate, locations, etc) of ECP projects and distributed workflows requiring Exascale compute resources. These requirements will help feed into the design and implementation of networks such as the DOE science mission network ESnet (www.es.net), as well as higher-level data movement tools and services. The target audience for this session are domain scientists, network engineers, and data movement researchers and administrators.

The session will be organized as roughly 30 minutes of presentations on current technologies (such as ESnet and Globus capabilities) and sample ECP data transfer use cases (e.g., such as moving experiment data to a remotely located HPC for analysis, or repositioning data when storage space at one location is exhausted), followed by 60 mins for discussions from session audiences on their data movement requirements, transfer workloads characteristics, etc.

The BOF organizers combine expertise in high-speed network infrastructure (Guok), data movement tools and services (Sim, Foster), and ECP applications (Habib, Foster, Sim).

TITLE **Kokkos Ecosystem for Performance Portability: Updates and Plans**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 4:00 PM - 5:30 PM central
DURATION 90 min

AUTHORS
Christian Trott (Sandia National Laboratories) crtrott@sandia.gov
Siva Rajamanickam (Sandia National Laboratories) srajama@sandia.gov

ABSTRACT
Kokkos is an ecosystem for performance portability with three distinct components - Kokkos Core programming model, Kokkos Kernels, and Kokkos Tools. Kokkos allows writing the applications in a portable manner with programming abstractions that translate well to performance. Kokkos has been adopted into several ECP applications. This session will update the ECP applications, ST teams, and codesign centers new features being planned in the Kokkos ecosystem that are targeted towards the Exascale systems in mind. This session will also cover the key new features released as part of Kokkos 3.0 release anticipated to be in FY20 Q1. We will also seek user feedback at the end of the session.

DESCRIPTION
The session will cover the three components of the Kokkos ecosystem. The agenda will include 20 minutes for each of the three components and a 30 minute interactive session to seek feedback from users.

Considering there are several ECP applications, co-design centers and software technology projects using some form of node level programming model, this session is targeted towards the intermediate user who has some experience with programming heterogeneous architectures. We will not assume prior Kokkos experience.

The first 60 minutes will cover expected new features targeting Exascale architectures and just released features in Kokkos 3.0. The first part will lay the foundation for the 30 minute discussion. We will seek feedback from both users of Kokkos and users of other programming models on these new features in an interactive session.

TITLE **Lossy Data Reduction/Compression for ECP Applications**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 4:00 PM - 5:30 PM central
DURATION 90 min

AUTHORS

Franck Cappello (Argonne National Laboratory) cappello@anl.gov
Peter Lindstrom (Lawrence Livermore National Laboratory) pl@llnl.gov

ABSTRACT

Extreme-scale scientific simulations and experiments on scientific instruments are already generating more data that can be communicated, stored, and analyzed. Lossy reduction/compression of scientific data is a necessity to allow for the classic usages of visualization, to dramatically accelerate I/O, and to reduce the data footprint on storage. Recent results have shown that lossy data reduction/compression can also be used to accelerate checkpoint/restart, reduce the memory footprint, and even accelerate computation. However, lossy reduction/compression needs to be performed wisely in order to keep the information that matters for the users. The current approach of data decimation (dropping a specified percentage of the produced data) drastically tampers with the data analytics performed on the reduced data. Instead, users can rely on advanced lossy compressors based on prediction, transforms, high-order SVD or other techniques.

The goal of this breakout session is to provide information about the latest developments of the compression technologies, the diversity of usages, and the variety of applications of lossy compression. We will also discuss the challenges faced by users when they add lossy compression in their pipeline. Specifically, we will cover three main topics related to lossy compression for scientific data:

(i) ECP compression software, including SZ, ZFP, MGARD and others;
(ii) the ECP use cases, covering different usages; and
(iii) methods to evaluate, compare, and assess the impacts of lossy compression.

The ECP compression technologies will be presented by their authors. We expect that some ECP use cases will be presented by the users.

DESCRIPTION

Background

Extreme-scale scientific simulations and experiments on scientific instruments are already generating more data that can be communicated, stored, and analyzed. Lossy reduction/compression of scientific data is a necessity to allow for the classic usages of visualization, to dramatically accelerate I/O, and to reduce the data footprint on storage. Recent results have shown that lossy data reduction/compression can also be used to accelerate checkpoint/restart, reduce the memory footprint, and even accelerate computation. However, lossy reduction/compression needs to be performed wisely in order to keep the information that matters for the users. The current approach of data decimation (dropping a specified percentage of the produced data) drastically tampers with the data analytics performed on the reduced data. Instead, users can rely on advanced lossy compressors based on prediction, transforms, high-order SVD or other techniques.

The ECP VeloC/SZ and Alpine/ZFP software technology projects are developing lossy compression technologies specifically tuned for ECP applications.

Several ECP applications have already tested and are engaged in adopting lossy compression for their data reduction needs. Other are starting exploring the use of lossy compressor to accelerate computation and/or reduce resource footprint. These ECP applications cover Cosmology, Quantum Chemistry, Molecular Dynamic, Local Weather Prediction, Precision Medicine for Cancer and Data Analytics for Instruments.

Topic
The breakout session will cover 3 main topics of lossy compression for scientific data:

(i) the compression software developed in ECP, including SZ, ZFP, MGARD and how to use them directly or through I/O libraries (HDF5, ADIOS, PnetCDF),
(ii) the ECP use-cases of lossy compression for scientific simulations and experiments, covering six different usages, including streaming intensity reduction, storage footprint reduction, I/O acceleration, accelerating checkpoint/restart, memory footprint reduction and re-computation avoidance and
(iii) methods to evaluate, compare and assess the impacts of lossy compression.

Goal
The first goal of this breakout session is to inform ECP applications looking for data reduction software about the compression technologies developed in ECP and their currently identified use-cases. The second goal is to present examples of performance improvements and reductions of resource occupation provided by ECP lossy compressors to ECP applications. The third goal is to expose challenges that users need to address when they add lossy compression in their pipeline (selection of the error bound type and level, validation of the decompressed data accuracy level, integration in the application code).

Target audience
All ECP applications in needs of effective data reduction techniques beyond lossless compression (that only reduce data size by at most 50%) and decimation in space or time. ECP I/O library developers who are interested in optimizing the integration of lossy compressor in their I/O library. ECP numerical library developers who are interested in understanding how lossy compression affect computation.

Speakers
The speakers will be the designers/developers of the lossy compression technologies developed in ECP. We expect some users to also discuss their usage of lossy compression, how they integrate it in their pipeline and what quality criteria and assessment methods they use.

TITLE **Management of Complex, Heterogeneous, and Hierarchical Memory Resources**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Kamil Iskra (Argonne National Laboratory) iskra@mcs.anl.gov
David Beckingsale (Lawrence Livermore National Laboratory) beckingsale1@llnl.gov
Mike Lang (Los Alamos National Laboratory) mlang@lanl.gov
Stephen Olivier (Sandia National Laboratories) slolivi@sandia.gov
Jeff Vetter (Oak Ridge National Laboratory) vetter@ornl.gov
Maya Gokhale (Lawrence Livermore National Laboratory) maya@llnl.gov
Swann Perarnau (Argonne National Laboratory) swann@anl.gov

ABSTRACT
Using and managing complex, heterogeneous, and hierarchical memory resources in operating systems, libraries, and applications is a challenge facing numerous ECP projects. Current systems already provide a range of memories: high-bandwidth memory (HBM), non-volatile memory (NVM), standard DDR, and GPU accelerator memory. The complexity of this space continues to increase as does the need for a performance-portable solution. In both the A21 exascale system and the CORAL2 exascale systems heterogeneous memory hierarchies are expected to be present. This interactive session will focus on the challenges these memory architectures introduce and the possible solutions provided by various ECP ST projects to the applications that need to harness these resources.

DESCRIPTION
Using and managing complex memory hierarchies in exascale systems presents a challenge for all ECP projects. In this session, we will present the approaches taken by various ECP ST projects to address this challenge and take questions from the audience. We will also present early experiences/technology updates.

The following projects will be represented:

2.3.1.16 - SICM: Simplified Interface to Complex Memories (Mike Lang, Maya Gokhale)
2.3.1.18 - RAJA/Kokkos (Umpire - David Beckingsale)
2.3.1.19 - Argo: Low-level resource management for the OS and runtime (AML - Swann Perarnau; UMap - Maya Gokhale)
2.3.2.10 - PROTEAS-TUNE (Jeff Vetter)
2.3.6.03 - Sandia National Laboratories ATDM (OpenMP - Stephen Olivier)

The planned format for the session is a series of brief project overviews/technology updates/etc. from project representatives, followed by an interactive Q&A, likely in the form of a panel. We will strive for a 50-50 split between these two parts.

This will be a user-focused session, with an emphasis on gathering requirements and sharing possible approaches to tackling these challenges at the application and library level. The target audience for this session are application, library, and runtime developers across the ECP.

TITLE **MPI BoF**

FOCUS AREA **ST**
TYPE BoF
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Pavan Balaji (Argonne National Laboratory) balaji@anl.gov
David Bernholdt (Oak Ridge National Laboratory) bernholdtde@ornl.gov
Alicia Grundhoffer (Oak Ridge National Laboratory) grundhoffeaa@ornl.gov

ABSTRACT
This session provides an opportunity for MPI users to interact with MPI developers and standards committee members, as represented by the two ECP-funded projects. The session will include an update on the activities of the MPI Forum and the status of various proposals for the standard, and brief overviews of the key activities of the two ECP MPI projects. At least half of the time will be devoted to discussions with the audience. We would like to hear from ECP code teams about their questions, concerns, and "unusual" use cases for MPI. As an overarching topic, we also want to discuss how ECP code teams and the MPI community can create a sustainable ecosystem in which both can thrive.

DESCRIPTION
This proposal comes jointly from the two MPI-related ECP projects: 2.3.1.17 OMPI-X: Open MPI for Exascale, and 2.3.1.07 Exascale MPI.

The goal of this session is to provide attendees with updates on the activities of the MPI Forum and the two projects, and to provide an opportunity for discussions with those interested in or concerned about MPI. Accordingly, the plan is to spend no more than 30 minutes in presentations, reserving the remainder for group discussion.

We have not had time to identify specific speakers yet. Our goal is to be able to add them before the agenda is published.

TITLE **Multiprecision Focus Effort in the ECP Math Library Ecosystem**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS

Hartwig Anzt (Karlsruhe Institute of Technology) hanzt@icl.utk.edu
Ulrike Meier Yang (Lawrence Livermore National Laboratory) yang11@llnl.gov
Jakub Kurzak (Advanced Micro Devices, Incorporated) Jakub.Kurzak@amd.com
Piotr Luszczek (The University of Tennessee, Knoxville) luszczek@icl.utk.edu
Siva Rajamanickam (Sandia National Laboratories) srajama@sandia.gov
Terry Cojean (Karlsruhe Institute of Technology) terry.cojean@kit.edu

ABSTRACT

Recently, hardware manufacturers responded to an increasing request for low precision functionality like fp16 by integrating low precision special function units, like, e.g., the NVIDIA Tensor cores. These, however, remain unused even for compute-intensive applications as long as high precision is employed for all arithmetic operations. At the same time, communication-intensive applications suffer from the memory bandwidth of architectures growing at a much slower pace than the arithmetic performance. In both cases, a promising strategy is to abandon the high precision standard (typically fp64), and employ lower or non-standard precision for arithmetic computations or memory operations whenever possible. While employing formats other than working precision can render attractive performance improvements, it also requires careful consideration of the numeric effects. On the other end of the spectrum, precision formats with higher accuracy than the hardware-supported fp64 can be effective in improving robustness and accuracy of numerical methods. With this breakout session, we want to create a platform where all ECP projects working with multiprecision or interested in using multiprecision come together and share their experiences, and present functionality they already provide, but also discuss software and application requirements, hardware limitations, and research synergies. Proceeding the ECP AHM, we anticipate summarizing the central findings of the breakout session in a whitepaper drafting a multiprecision strategy for ECP.

DESCRIPTION

In this session, we focus on multiprecision aspects in numerical libraries for ECP. Special focus is put on the memory bottleneck for sparse linear algebra, and how to benefit from the adoption of special function units for low precision arithmetic. The breakout session combines a schedule of contributed talks (with a preliminary agenda below) with a community discussion and the formation of topic interest groups focusing on specific aspects of multiprecision. The session is supposed to kick off a multiprecision math library effort where regular virtual meetings will enable the collaborative progress in this area.

Preliminary Program:

- Introduction (10min) Ulrike Meier Yang, LBNL and Hartwig Anzt, UTK/KIT
- Multiprecision Focus Effort in the ECP Math Library Ecosystem (10 min)
- Contributed talks (6 x (5min + 2 min questions))
  - Multiprecision BLAS, Jakub Kurzak, ICL
  - Low precision kernels for ML, Siva Rajamanickam, SNL

  - Mixed Precision Iterative Refinement, Piotr Luszczek, ICL
  - Mixed Precision Multigrid, Ulrike Meier Yang, LLNL
  - Multiprecision Preconditioning, Terry Cojean, KIT
  - Modular Precision Ecosystem, Hartwig Anzt, ICL/KIT
- Discussion & Forming of interest groups focusing on specific aspects (30 min)
  - BLAS
  - Preconditioning
  - Krylov solvers
  - Multigrid
  - Modular Precision Ecosystem
  - Application Integration?
- Closing (5 min) Hartwig Anzt, UTK/KIT: announcement Slack channel, xSDK multiprecision meetings, design document community invitation.

TITLE **OLCF ECP Application Engagement**

FOCUS AREA **HI**
TYPE BoF
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 4:30 PM - 5:30 PM central
DURATION 60 min

AUTHORS
Judith Hill (Oak Ridge National Laboratory) hilljc@ornl.gov
Bronson Messer (Oak Ridge National Laboratory) bronson@ornl.gov
Justin Whitt (Oak Ridge National Laboratory) whitt@ornl.gov

ABSTRACT
The OLCF will discuss its ECP application engagement model and how this engagement relates to the OLCF Center for Accelerated Application Readiness (CAAR). Specifically, we will present publicly available information about Frontier, the current and expected future compute resources that are or will be available to ECP teams, and we will introduce the OLCF Center of Excellence model (COE) and the resources available to ECP through the COE. This session will be an opportunity for ECP application teams to directly engage with the OLCF application integration personnel, the CAAR management team, and Frontier vendor partners who are part of the OLCF COE. We encourage ECP teams to communicate their application's requirements and current challenges as they are thinking about Frontier in this session.

DESCRIPTION
Through the Center for Accelerated Application Readiness (CAAR), the OLCF is embarking on its third instantiation of this program following two successful programs readying applications for Titan and Summit. OLCF ECP Engagement Applications are a part of this program, and this session will provide an opportunity for a face-to-face discussion with ECP-funded OLCF Application Integration personnel, CAAR and ECP management, and the Frontier vendor Center of Excellence (COE) personnel. Though some of the Frontier material regarding hardware, software and schedule is likely to be presented in other venues at the ECP meeting, this BOF provides an opportunity for ECP teams to ask questions in a smaller more-focused forum.

Goals:
The goal of this session is three-fold:

(1) Provide the most current OLCF plan of record for the schedule of expected compute resources available to ECP application teams,
(2) Introduce the ECP teams to the OLCF Center of Excellence and the compute and personnel resources available to them within the COE, and
(3) Provide a forum for ECP teams to engage directly with their OLCF Application Integration personnel.

Target Audience:
ECP Application teams interested in the OLCF model for application readiness. Secondary target audience are the ECP ST teams that are interested in the OLCF COE resources available for ST teams.

TITLE **OpenMP Roadmap for Accelerators across DOE Pre-Exascale/Exascale Machines**

FOCUS AREA **HI**
TYPE BoF
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Yasaman Ghadar (Argonne National Laboratory) ghadar@anl.gov
Colleen Bertoni (Argonne National Laboratory) bertoni@anl.gov
Jae Hyuk Kwack (Argonne National Laboratory) jkwack@anl.gov
Jose Monsalve Diaz (Argonne National Laboratory) jmonsalvediaz@anl.gov
Johannes Doerfert (Argonne National Laboratory) jdoerfert@anl.gov
Brian Homerding (Argonne National Laboratory) bhomerding@anl.gov
Tim Williams (Argonne National Laboratory) tjwilliams@anl.gov
Ray Loy (Argonne National Laboratory) rloy@alcf.anl.gov
Ye Luo (Argonne National Laboratory) yeluo@anl.gov
Kalyan Kumaran (Argonne National Laboratory) kumaran@anl.gov
Hal Finkel (Argonne National Laboratory) hfinkel@anl.gov
Jack Deslippe (Lawrence Berkeley National Laboratory) jrdeslippe@lbl.gov
Christopher Daley (Lawrence Berkeley National Laboratory) csdaley@lbl.gov
Rahul Gayatri (Lawrence Berkeley National Laboratory) rgayatri@lbl.gov
Thorsten Kurth (Lawrence Berkeley National Laboratory) tkurth@lbl.gov
Charlene Yang (Lawrence Berkeley National Laboratory) CJYang@lbl.gov
Brian Friesen (Lawrence Berkeley National Laboratory) bfriesen@lbl.gov
Jay Srinivasan (Lawrence Berkeley National Laboratory) jsrinivasan@lbl.gov
Oscar Hernandez (Oak Ridge National Laboratory) oscar@ornl.gov
Reuben Budiardja (Oak Ridge National Laboratory) reubendb@ornl.gov
Bronson Messer (Oak Ridge National Laboratory) bronson@ornl.gov
Gustav Jansen (Oak Ridge National Laboratory) jansengr@ornl.gov
Matthew Norman (Oak Ridge National Laboratory) normanmr@ornl.gov
Piotr Luszczek (The University of Tennessee, Knoxville) luszczek@icl.utk.edu
Stephen Olivier (Sandia National Laboratories) slolivi@sandia.gov
Bronis R. de Supinski (Lawrence Livermore National Laboratory) bronis@llnl.gov
Doug Miles (NVIDIA) dmiles@nvidia.com
Annemarie Southwell (NVIDIA) asouthwell@nvidia.com
Steve Scalpone (NVIDIA) sscalpone@nvidia.com
Gregory Rodgers (Advanced Micro Devices, Inc.) Gregory.Rodgers@amd.com
Carlos Rosales-Fernandez (Intel Corp.) carlos.rosales-fernandez@intel.com
Luiz DeRose (Cray) ldr@hpe.com
Wang Chen (IBM) wdchen@ca.ibm.com
XinMin Tian (Intel Corp.) xinmin.tian@intel.com
Barbara Chapman (Stony Brook Univeristy) barbara.chapman@stonybrook.edu
Vivek Kale (Brookhaven National Laboratory) vkale@bnl.gov

ABSTRACT
The upcoming pre-exascale/exascale machines (i.e., Aurora, El Capitan, Frontier, Perlmutter) across DoE facilities will support the OpenMP programming model for accelerators. OpenMP must not only be

functional across these machines but also provide the performance that users expect. In this session, facilities representatives and vendors including AMD, Cray, Intel, and NVIDIA/PGI will come together to present their roadmaps for the OpenMP 5.0 programming model across next generation supercomputers. We will discuss which portion of the specification will be supported, details of early test beds that users can access, and the support for Fortran-based applications. We will discuss the performance characteristics for each machine and regression tests available. Each vendor will provide details about their tools for the OpenMP programming model. Interoperability among OpenMP and vendor-specific components such as languages (e.g., CUDA, HIP, SYCL) and math libraries will be reviewed. Finally, each vendor will provide guidance on how to port ECP applications to their target architectures.

DESCRIPTION

In this session, vendors will present a roadmap for OpenMP 5.0 features and functionalities to be available on each target machine. In this interactive session with the audience, facilities representatives and vendors will discuss their expectations and differences for running a code using MPI+OpenMP on the various systems' multi-GPU node(s). This could include examples of affinity settings for the different architectures. Vendors will describe the expected path to port a code to the target machine, changes that the user will need to make, and performance expectations.

Other issues of interest would include:

a) Interoperability between other programming models and libraries (for example calling a vendor math library and also having OpenMP target regions)
b) Feedback from application developers about which features in the OpenMP spec should be prioritized in the vendors, OpenMP implementations.What base language features do application developers need to be supported in OpenMP target regions.
c) How the OpenMP implementations map loop constructs to GPU hardware, allocate objects in GPU shared memory, and how programmers should write code to get the highest performance on the target platform and whether this is likely to perform well with other vendor implementations.
d) How vendors evaluate the success of their OpenMP compiler? What regression and performance testing do they use? How do vendors prioritize what to optimize in the OpenMP compiler? How do vendors want the HPC community to help them?

TITLE **Overview of the Argonne Aurora Exascale System**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Scott Parker (Argonne National Laboratory) sparker@anl.gov
Chris Knight (Argonne National Laboratory) knightc@anl.gov
Kevin Harms (Argonne National Laboratory) harms@alcf.anl.gov
Tim Williams (Argonne National Laboratory) tjwilliams@anl.gov
Katherine Riley (Argonne National Laboratory) riley@alcf.anl.gov
Kalyan Kumaran (Argonne National Laboratory) kumaran@anl.gov
Venkat Vishwanath (Argonne National Laboratory) venkat@anl.gov
James Osborn (Argonne National Laboratory) osborn@alcf.anl.gov
Yasaman Ghadar (Argonne National Laboratory) ghadar@anl.gov

ABSTRACT
In 2021 Argonne National Laboratory will take delivery of Aurora, a system delivering greater than 1 Exaflop of sustained computing performance. This session will provide details and updates on the publicly available information on Aurora, covering the systems hardware and software. Software details including supported programming models, tools, and libraries will be covered. In addition, information and roadmaps for the Aurora development hardware and software environments will be presented. An overview of progress by the Aurora Early Science Program and Aurora ECP applications development efforts will also be discussed.

DESCRIPTION
This session is intended to provide details and updates on the Aurora system design to the ECP community. This session will provide a broad overview of the Aurora hardware and software that will be useful for ECP projects intending to use Aurora. To reach the maximum audience, which is anticipated to be applications, libraries, and tools developers only public information will be disclosed. Public information on Aurora useful to developers currently exists and it is anticipated that more information will be made public prior to, or for, the ECP annual meeting. Full details on the hardware are not expected to be public at the time of the ECP annual meeting but all public hardware information will be presented. Nearly all of the Aurora software environment is expected to be public and will be discussed. Details and roadmaps for the development hardware and software environments will presented along with an overview of the development work for Aurora that has been performed by ECP projects and by Argonne.

TITLE **Overview of the National Security Applications**

FOCUS AREA **AD**
TYPE Panel
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Marianne Francois (Los Alamos National Laboratory) mmfran@lanl.gov
David Daniel (Los Alamos National Laboratory) ddd@lanl.gov
Aimee Hungerford (Los Alamos National Laboratory) aimee@lanl.gov
Robert Rieben (Lawrence Livermore National Laboratory) rieben1@llnl.gov
William Rider (Sandia National Laboratories) wjrider@sandia.gov

ABSTRACT
Following from the successful 2019 panel session "Lessons learned from National Security", this year panel session will provide an overview of the accomplishments of the National Security Applications projects. During this session, each project will highlight their success at integrating software technologies for solving multiphysics applications on various computer architectures. The discussion will focus on portability, productivity and code performance.

DESCRIPTION
The topics to be covered in this session include: multiphysics applications, high energy density physics, materials models, plasma physics, multiscale methods, high-order methods, particle-in-cell methods, compressible flows, verification and validation, parallel programming models, software integration (Kokkos, RAJA, Legion, Caliper, Alpine, etc), portability and productivity.

TITLE **PathForward Update – AMD (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Beginner

DAY Monday, February 3, 2020
TIME 4:45 PM - 6:00 PM central
DURATION 75 min

AUTHORS
Gabriel Loh (Advanced Micro Devices, Inc.) gabriel.loh@amd.com
Mike Schulte (Advanced Micro Devices, Inc.) michael.schulte@amd.com
Bill Brantley (Advanced Micro Devices, Inc.) bill.brantley@amd.com
Chip Frietag (Advanced Micro Devices, Inc.) chip.freitag@amd.com

ABSTRACT
This NDA session provides an overview and update of AMD's advanced research activities under the ECP PathForward program. The agenda will include updates on roadmaps, software/tools/applications, research outcomes and impacts, and an overview of systems with AMD technologies.

TITLE **PathForward Update – Cray (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Monday, February 3, 2020
TIME 3:15 PM - 4:30 PM central
DURATION 75 min

AUTHORS
Dan Ernst (Cray) dje@cray.com
Michael Langer (Cray) mlanger@cray.com

ABSTRACT
This session will provide an update of the Cray PathForward project. The session will include NDA material and attendees will be limited to those covered by an appropriate PathForward-specific NDA.

Pre-registration for the session is required in order to provide an opportunity to verify that attendees are covered by an appropriate PathForward-specific NDA.

TITLE **PathForward Update – HPE (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Monday, February 3, 2020
TIME 8:30 AM - 9:45 AM central
DURATION 75 min

AUTHORS
Dennis Floyd (Hewlett Packard Enterprise) dennis.floyd@hpe.com
Paolo Faraboschi (Hewlett Packard Enterprise) paolo.faraboschi@hpe.com
Nicolas Dube (Hewlett Packard Enterprise) nicolas.dube@hpe.com
Michael Lough (Hewlett Packard Enterprise) michael.lough@hpe.com

ABSTRACT
This session will provide an overview of the HPE PathForward project. The session will include NDA material and attendees will be limited to those covered by an appropriate NDA. Pre-registration for the session is required in order to provide an opportunity to verify that attendees are covered by an appropriate NDA.

TITLE **PathForward Update – IBM (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Monday, February 3, 2020
TIME 10:00 AM - 11:15 AM central
DURATION 75 min

AUTHORS
Ruud Haring (IBM Research) ruud@us.ibm.com
Jim Sexton (IBM Research) sextonjc@us.ibm.com
Jim Kahle (IBM Research) jakahle@us.ibm.com

ABSTRACT
This session will provide an overview of the IBM PathForward project. IBM's PathForward research aims develop, assess and optimize Exascale Systems Solution concepts combining IBM POWER processors, GPUs, in-memory processing, high-performance networks, and high-performance storage for possible US Department of Energy (DOE) Exascale and Post Exascale systems in the 2021 to 2028 time frame. IBM is exploring co-designed hardware, software and applications addressing a broad set of DOE workloads, to optimize a systems architecture to meet DOE Exascale and Post-Exascale requirements. The solution will support modularity, scalability and configurability for specific workloads and DOE opportunities. To this end, IBM's PathForward research focuses on architectural innovations, system component innovations and technology innovations.

DESCRIPTION
Overview of IBM's PathForward project -- now nearing completion.

TITLE **PathForward Update – Intel (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Monday, February 3, 2020
TIME 1:45 PM - 3:00 PM central
DURATION 75 min

AUTHORS
Irina Petrova (Intel Corp.) irina.v.petrova@intel.com
Samantika Sury (Intel Corp.) samantika.s.sury@intel.com
Simon Steely (Intel Corp.) simon.c.steely.jr@intel.com
Kent Glossop (Intel Corp.) kent.glossop@intel.com
Carl Beckmann (Intel Corp.) carl.j.beckmann@intel.com
Venkata Krishnan (Intel Corp.) Venkata.Krishnan@intel.com

ABSTRACT
The challenges in reaching exascale computing are well documented; the most common of these are energy-efficiency, resiliency, cost, and usability. Intel's approach to achieve the Department of Energy's (DOE's) PathForward goals leverage the Intel state-of-the-art technology as well as strengthen the entire Intel Architecture (IA) ecosystem for HPC. Intel will present an update on the work done under the PathForward program. The session will include NDA material and attendees will be limited to those covered by an appropriate NDA. Pre-registration for the session is required in order to provide an opportunity to verify that attendees are covered by an appropriate NDA.

TITLE **PathForward Update – NVIDIA (NDA Session)**

FOCUS AREA **HI**
TYPE Breakout
LEVEL N/A

DAY Monday, February 3, 2020
TIME 11:30 AM - 12:45 PM central
DURATION 75 min

AUTHORS
Matthew Bolitho (NVIDIA) mbolitho@nvidia.com
John Danskin (NVIDIA) jdanskin@nvidia.com
Sachin Idgunji (NVIDIA) sidgunji@nvidia.com
Alan Menezes (NVIDIA) amenezes@nvidia.com
Nirmal Saxena (NVIDIA) nsaxena@nvidia.com
Vicki Little (NVIDIA) vlittle@nvidia.com

ABSTRACT
This session will provide an update of the NVIDIA PathForward project and an overview onprogress in the last year. The session will include NDA material and attendees will be limited tothose covered by an appropriate NDA. Pre-registration for the session is required in order toprovide an opportunity to verify that attendees are covered by an appropriate NDA.

TITLE **Performance Portability at the Intersection of the Exascale Computing Project and the DOE Centers of Excellence**

FOCUS AREA **HI**
TYPE BoF
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 1:30 PM - 3:00 PM central
DURATION 90 min

AUTHORS
Douglas Doerfler (Lawrence Berkeley National Laboratory) dwdoerf@lbl.gov
Hai Ah Nam (Los Alamos National Laboratory) hnam@lanl.gov
Rob Neely (Lawrence Livermore National Laboratory) neely4@llnl.gov
Philip Roth (Oak Ridge National Laboratory) rothpc@ornl.gov
Thorsten Kurth (Lawrence Berkeley National Laboratory) tkurth@lbl.gov

ABSTRACT
To satisfy the ever-growing computational demands of scientific and engineering applications, the architectures and software stacks of high performance computing systems have become increasingly complex. The Department of Energy (DOE) computing centers have established Centers of Excellence (COEs) to help users understand and make best use of these complex systems. Over the past several years, COE members have met annually to discuss cross-cutting topics like application performance portability and productivity (PPP). In this birds of a feather session, we update the Exascale Computing Project (ECP) community on recent COE PPP activities and solicit input on future plans. After a brief overview of COE PP activities, a small number of invited speakers give a deeper dive into key COE PP methods and use cases. We conclude the session with an extended, interactive discussion about the PPP challenges that face the ECP community and the potential approaches for addressing them for the near-exascale and planned exascale systems.

DESCRIPTION
To help users deal with the increasing complexity of high performance computing system hardware and software, the Department of Energy (DOE) has established several Centers of Excellence (COEs) at its computing centers. Each COE focuses on specific technologies and is staffed by vendor personnel working in close proximity with DOE computing center personnel.

Over the past several years, DOE computing center personnel have organized a series of meetings focused on the cross-cutting issue of application performance portability. These meetings are an opportunity for the COEs to update each other on recent successes, to discuss challenges that they still face, and to share ideas about how to collaborate to overcome those challenges.

Because performance portability is a cross-cutting topic, it impacts a large segment of the ECP community. Application Development (AD) projects need performance portability techniques that allow their software to perform well on a variety of architectures, and at a range of scales. And Software Technology (ST) projects may provide the libraries and tools that support those AD project needs.

To provide a venue for COE members and ECP members to discuss their common performance portability successes, challenges, and techniques, we propose a 90-minute birds-of-a-feather (BOF) session. We will begin the BOF with an overview of recent and planned COE performance portability activities, followed by a small number of brief presentations covering performance portability methods,

technologies, and application use cases. Potential topics for these presentations include ways of evaluating the effectiveness of various performance portability: approaches, programming models, metrics, and application use cases. After establishing this context, we will transition to a highly interactive discussion about performance portability activities and needs within the intersection of the ECP and the DOE COE communities.

We have two primary goals for this BOF: to solicit feedback from the ECP community on ways that the COE performance productivity activities can be more effective and relevant to the ECP project; and to establish productive relationships between the two communities that will be mutually beneficial as the DOE prepares for and deploys its future exascale system(s).

TITLE **Perlmutter - a Waypoint for ECP Teams**

FOCUS AREA HI
TYPE Breakout
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Jay Srinivasan (Lawrence Berkeley National Laboratory) jsrinivasan@lbl.gov
Jack Deslippe (Lawrence Berkeley National Laboratory) jrdeslippe@lbl.gov
Brandon Cook (Lawrence Berkeley National Laboratory) bgcook@lbl.gov
Brian Friesen (Lawrence Berkeley National Laboratory) bfriesen@lbl.gov

ABSTRACT
The Perlmutter system at NERSC, expected to land in late 2020, features future generation AMD Epyc CPUs, future generation NVIDIA GPUs and upon delivery will be the largest system with Cray's next-generation Slingshot interconnect. The systems represents the final DOE SC pre-exascale system before the delivery of the Exascale Aurora and Frontier systems. In this session, we will discuss the details of the upcoming Perlmutter system. We will also focus on how ECP teams can optimize their codes for Perlmutter as a waypoint for Exascale with an emphasis on performance-portable programming approaches including OpenMP 4.5/5.0, Kokkos, HIP etc. We will finish the session with a series of case-studies from ECP applications participating in NERSC's NESAP application readiness program (though the ECP HI Apps Integration area) - highlighting progress being made toward ECP FOM improvement goals (KPPs) working with NERSC staff and Perlmutter Centre Of Excellence (COE) resources.

DESCRIPTION
The Perlmutter system at NERSC, which will be described and discussed in this session, represents the final DOE SC pre-exascale system before the delivery of the Exascale Aurora and Frontier systems, and will be, upon delivery, the largest system with Cray's next-generation Slingshot interconnect. Of particular interest to ECP developers will be a discussion on how ECP teams can get early access to Perlmutter to test their applications on next-generation GPUs, an all-flash filesystem and the slingshot interconnect. We will also focus on how ECP teams can optimize their codes for Perlmutter as a waypoint for Exascale with an emphasis on performance-portable programming approaches including OpenMP 4.5/5.0, Kokkos, HIP etc. We will discuss NERSC's OpenMP effort with PGI. The bulk of the session will cover a series of case-studies from ECP applications participating in NERSC's NESAP application readiness program (though the ECP HI Apps Integration area) - highlighting progress being made toward ECP FOM improvement goals (KPPs) working with NERSC staff and Perlmutter Centre Of Excellence (COE) resources. These include updates from the EXAALT, ExaBiome, AMReX (WarpX), ExaFEL and WDMAPP. This session is targeted at ECP AD teams, ECP AD Management and key ECP ST projects. Presenters will be NERSC Staff and ECP AD team members.

TITLE **Persistent Memory in the Exascale Memory/Storage Continuum**

FOCUS AREA **HI**
TYPE BoF
LEVEL Intermediate

DAY Thursday, February 06, 2020
TIME 3:30 PM - 4:15 PM central
DURATION 45 min

AUTHORS
Maya Gokhale (Lawrence Livermore National Laboratory) maya@llnl.gov

ABSTRACT
Low latency byte- and block-access persistent memories are now commercially available and offer a new capability in the memory/storage space to support data rich exascale applications including Machine Learning, in-situ data analysis, and novel in-system caching of global data sets. As ECP applications broaden to encompass large scale data processing and analysis, there is a shift in access patterns from write-dominated to increasingly read-heavy. Persistent Memory such as the Intel technology can be used as a very large capacity extended memory, as a light-weight byte-addressable file system, or as a high performance tier in the storage system. The purpose of this BoF is to share early experiences within the ECP community for persistent memory assessments from multiple perspectives, e.g. performance, reliability, power/energy. We hope ECP participants can discuss existing, proposed, and potential use cases for persistent memory.

DESCRIPTION
As ECP applications transition to include ML and data analytics as integral components of workflows, persistent memories offer an attractive newly emerging technology alternative to hold portions of very large global data sets within the fabric of the computing. Hardware technologies to be employed include large battery-backed DRAM, NVDIMMs containing e.g. Intel Optane modules, and commodity NVMe SSDs. The latencies, bandwidths, power/energy/ and durability of these devices are varied. Existing use of persistent memory is to use SSDs either node local or at I/O nodes as output devices for checkpoints that get flushed to global storage. As external data becomes increasingly important for applications such as ML training and data analysis, a shift towards read-dominated access patterns for these applications will occur. New and existing forms of persistent memory are particularly efficient and show highest performance on read-dominated workloads. The purpose of the BoF is to bring together those interested in hardware evaluation of persistent memories, system architectures planners, system software developers including file system and in-transit storage system, and application developers to discuss the many facets of persistent memory and how they can benefit exascale architectures.

TITLE **Preparing Applications for Aurora**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 1:30 PM - 3:00 PM central
DURATION 90 min

AUTHORS
Brian Homerding (Argonne National Laboratory) bhomerding@anl.gov
Scott Parker (Argonne National Laboratory) sparker@anl.gov
Tim Williams (Argonne National Laboratory) tjwilliams@anl.gov
Venkat Vishwanath (Argonne National Laboratory) venkat@anl.gov
Ramesh Balakrishnan (Argonne National Laboratory) bramesh@anl.gov
Jeff Garelick (Intel Corp.) jeff.garelick@intel.com
Jose Nouddhouenou (Intel Corp.) jose.noudohouenou@intel.com
Roy Moore (Intel Corp.) roy.moore@intel.com
Carlos Rosales-Fernandez (Intel Corp.) carlos.rosales-fernandez@intel.com

ABSTRACT
The application preparation effort within ECP for Argonne's Exascale system Aurora is well underway. In this session we will present insights for Aurora application developers gained from this preparation. An overview of the Aurora software stack and public hardware information will be presented from an application developer's point of view. Additionally, results from the Aurora Early Science Program (ECP) and ECP applications development efforts will be described to enable application developers to learn from one another's experiences. These applications will cover several of the programming models being targeted by Aurora and will be of interest to anyone looking to run on Aurora. The ESP effort includes many ECP simulation applications and additionally hosts 10 projects explicitly integrating data-intensive computing, machine learning, and simulation to be run on Aurora. We will discuss aspects of developing and scaling these complex workflows for Aurora, in which optimizing simulation codes written in conventional compiled languages is only part of the story.

DESCRIPTION
In preparation for Aurora there has been significant effort in getting ECP applications ready. This breakout will showcase stories from various ECP applications highlighting the obstacles overcome and successes achieved. These stories will cover many of the programming models supported on Aurora including OpenMP, SYCL/DPC++ and OpenCL. Application developers will find this session useful to learn from other developers experience porting ECP codes to Aurora.

Presented will be a brief overview of the steps we suggest for porting codes to Aurora. Additionally, we will cover the hardware and software that are in use today for the development, porting, and optimizing of code along with a roadmap for future hardware and software availability. This will include a brief overview for transitioning to programming models supported on Aurora along with some best practices for development with an evolving software stack.

From the Aurora ESP effort there will be presentations on ESP simulation application development along with the incorporation of Data and Learning with Simulation. The standard development environment for Data and Learning components is typically high-level frameworks and libraries which encapsulate machine-optimized implementations, and productivity-oriented languages such as Python and R. We will

illustrate some of the project workflows, and how they are being evolved toward early Exascale execution on Aurora through the use of present-day tools and early access to vendor-optimized frameworks and libraries.

Following the overview of hardware and software we will dive into presentations from various ECP/ESP application developers and their work. These short presentations will take the majority of the session.

TITLE **Slingshot BoF (NDA Session)**

FOCUS AREA **HI**
TYPE BoF
LEVEL Expert

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Taylor Groves (Lawrence Berkeley National Laboratory) tgroves@lbl.gov
Sudheer Chunduri (Argonne National Laboratory) sudheer@anl.gov
Ian Karlin (Lawrence Livermore National Laboratory) karlin1@llnl.gov
Scott Atchley (Oak Ridge National Laboratory) atchleyes@ornl.gov

ABSTRACT
The Slingshot network is a crucial component of US Exascale systems, being a part of four of the next DOE systems (NERSC, ALCF, OLCF and LLNL). Developing a deeper understanding of advanced topics and best practices across the labs and Cray is of significant value. Topics of interest include: routing, congestion, traffic classes, hardware offload, optimization for accelerators.

The session will include NDA material and attendees will be limited to those covered by an appropriate NDA. Pre-registration for the session is required in order to provide an opportunity to verify that attendees are covered by an appropriate NDA.

TITLE **Spack Roundtable**

FOCUS AREA **ST**
TYPE BoF
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Todd Gamblin (Lawrence Livermore National Laboratory) tgamblin@llnl.gov
Gregory Becker (Lawrence Livermore National Laboratory) becker33@llnl.gov

ABSTRACT
Spack is a package manager for scientific computing, with a rapidly growing open source community. Spack has over 400 contributors from academia, industry, and laboratories across the world, and it is used to manage software releases for the U.S. Exascale Computing Project. At this roundtable, Spack core developers will give updates on the community, and on new features, container integration, and collaboration efforts with facilities, HI, E4S, and SDK teams. Sites using Spack will describe their experiences, and we will open the floor for discussion. All are invited to provide feedback, request features, and discuss future directions. Help us make installing HPC software simple!

DESCRIPTION
HPC software is becoming increasingly complex. The largest applications require over 100 dependency libraries, and they combine interpreted languages like Python with C, C++, and Fortran libraries. To achieve good performance, developers tune for multiple compilers, build options, and implementations of dependency libraries like MPI, BLAS, and LAPACK. The space of possible build configurations is combinatorial, and developers waste countless hours rebuilding software instead of producing new scientific results.

Spack (https://spack.io) is a package manager for scientific computing. It aims to reduce the complexity of building and installing HPC software on laptops, clusters, and the most powerful supercomputers in the world. There are over 3,300 package recipes in the Spack mainline repository. Spack has a rapidly growing open source community from across the world, with over 400 contributors from academia, industry, and laboratories. End users install complex HPC applications; developers manage dependencies for themselves and for their team; and the largest supercomputing sites in the world use Spack to deploy software for thousands of users. Spack is also the focus of a US/Japan collaboration on the Fugaku Supercomputer, and it is being used to manage software releases for the U.S. Exascale Computing Project.

The goals of this BOF are 1) to inform users within ECP about recent and upcoming developments, 2) to connect sites and developers using Spack to manage software, 3) to solicit feedback to guide future directions, and 4) to build the Spack community within ECP.

In the past, the Spack BOF has been well attended at the ECP meeting, and it has been a great source of user feedback for us. In particular, we have used feedback from past BOFs to guide our feature development.

This year, we would like to highlight a number of new features and integrations in Spack that have been impactful across multiple areas of ECP. These include integration with the GitLab continuous integration

system, build pipelines for users and facilities, Spack Stacks (a new mechanism for facility deployment), and fine-grained microarchitecture builds, and container integration. We would also like to highlight our work across NNSA and SC labs.

TITLE **Speed Dating for ECP: 1-on-1 Conversations between AD Teams and Math Library Developers**

FOCUS AREA **ST**
TYPE Breakout
LEVEL N/A

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Lois Curfman McInnes (Argonne National Laboratory) curfman@mcs.anl.gov
ECP math library teams (ANL, LBNL, LLNL, ORNL, SNL, U Tennessee) <ecp-st-math-lib-teams@elist.ornl.gov>

ABSTRACT
ECP math libraries encapsulate a diverse range of cutting-edge scalable algorithms to support the needs of ECP applications. Partnerships between AD teams and math library developers potentially offer the biggest payoffs for ECP, where AD teams motivate new functionalities in math libraries and also can customize library capabilities as needed for their projects. As we begin a new phase of ECP work, we have a compelling opportunity to re-examine (1) what math functionalities ECP applications need, (2) what capabilities ECP math libraries provide, and (3) how to pursue application-specific customization, so that together we can more effectively work toward next-generation science goals.

In this session AD teams will have an opportunity for 1-on-1 conversations with math library developers---to explain their needs, learn about library capabilities, and explore how to use/extend library functionalities. A key element of this approach is that conversations will focus on the particular circumstances of applications teams, with goals of advancing understanding on both ends and setting the stage for deeper conversations in the future.

Approach:

In fall 2019, the developers of math libraries will provide concise information about current capabilities and future plans.
We will encourage AD teams to review this information and indicate preferences for talking with developers of particular libraries.
We will then make a master schedule for the breakout, indicating pairings (of app/package conversations) for several rotations (about 20-30 minutes each).

Goals for AD teams: Make connections with new math library teams and deepen understanding in current collaborations. Explore how to use multiple math libraries in combination.

Goals for math libraries teams: Increase collective understanding of AD needs, so we can ensure that math library capabilities fully serve the ECP community.

These interactions are an important element of expanding and deepening AD/ST partnerships, so that library developers can help AD teams make more effective use of math packages, and AD teams can partner with library developers to extend and customize capabilities as needed to support their science goals.

Appendix:

Areas of ECP library functionality, as summarized in slides ECP Final Design Report

2.3.3.01: xSDK (multi-package build infrastructure, xSDK community policies compatibility for better software quality & sustainability, interoperability among packages)
2.3.3.06: PETSc (scalable linear & nonlinear solvers, integrators), TAO (numerical optimization), libEnsemble (ensemble management for exascale platforms)
2.3.3.07: STRUMPACK & SuperLU (scalable sparse direct solvers, preconditioners), FFTX (FFT stack, including symbolic analysis and code generation)
2.3.3.12: SUNDIALS (adaptive time integrators, nonlinear solvers), hypre (scalable linear solvers, with emphasis on algebraic multigrid)
2.3.3.13: CLOVER project: SLATE (exascale-capable dense linear algebra), FFT-ECP (scalable FFTs), PEEKS (latency-tolerant preconditioned iterative solvers, via Trilinos, Ginkgo, MAGMA-sparse), KokkosKernels (portable performance kernels for linear algebra and graph algorithms)
2.3.3.14: ALExa project: DTK (parallel data transfer between grids, search tree capability), Tasmanian (uncertainty quantification, surrogate modeling), ForTrilinos (automatic generation of Fortran interfaces for Trilinos)

TITLE **SUNDIALS Plans and Uses**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 8:30 AM - 10:00 AM central
DURATION 90 min

AUTHORS
Carol S. Woodward (Lawrence Livermore National Laboratory) woodward6@llnl.gov
Daniel R. Reynolds (Southern Methodist University) reynolds@smu.edu
David J. Gardner (Lawrence Livermore National Laboratory) gardner48@llnl.gov
Cody J. Balos (Lawrence Livermore National Laboratory) balos1@llnl.gov

ABSTRACT
The goal of this breakout session is to help connect SUNDIALS software developers with stakeholders from ECP applications and software technology projects. The session will start with an overview of the SUNDIALS suite emphasizing new features and plans for the upcoming 1-2 years. The session will continue with brief presentations from 2-3 SUNDIALS application users. Marc Day of the AMReX team, Aaron Fisher of the MFEM/CEED team, and Mark Taylor of the E3SM climate code team are confirmed speakers. The session will end with informal discussion among participants.

The goals of this discussion will be:

Share successes and failures of SUNDIALS uses among applications
Provide an overview of SUNDIALS for ECP-AD projects interested in using the software
Discuss software technology collaborations and identify features needed for further ST collaborations

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-787927.

TITLE **SYCL Programming Model for Aurora**

FOCUS AREA **HI**
TYPE BoF
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Yasaman Ghadar (Argonne National Laboratory) ghadar@anl.gov
Thomas Applencourt (Argonne National Laboratory) applencourt@anl.gov
Brian Homerding (Argonne National Laboratory) bhomerding@anl.gov
Kevin Harms (Argonne National Laboratory) harms@alcf.anl.gov
Jeff Hammond (Intel) jeff.r.hammond@intel.com

ABSTRACT
This presentation will focus on SYCL/DPC++ as a programming model that will be optimized for use on Aurora, ALCF's exascale machine. In this presentation we will discuss some of the main advantages of SYCL and discuss our early experiences of porting kernels to SYCL. We will begin with a very short overview of Aurora system architecture and then provide an overview of the SYCL runtime. We will walk through various code examples to explain the key features of SYCL and how to use them.

DESCRIPTION
This talk will introduce SYCL as a programming model for Aurora, the upcoming Argonne exascale machine. SYCL is a single source heterogeneous programming model based on standard C++. It uses C++ templates and lambda functions for host and device code. SYCL builds on the underlying concepts of portability and efficiency of OpenCL that enables code for heterogeneous processors, however it is less verbose compare to OpenCL. The single-source programming enables the host and kernel code for an application to be contained in the same source file, in a type-safe way and with the simplicity of a cross-platform asynchronous task graph. We will provide an overview of the SYCL concepts, compilation and runtime. No prior knowledge of OpenCL is required. Once we have reviewed the core concepts of SYCL, we will walk through several code examples to highlight the key features of SYCL. SYCL by design is hardware agnostic and offers the potential to be portable across many of DoE's largest machines.

TITLE **System Software - Evaluation and Future for Pre-Exascale and Exascale Systems**

FOCUS AREA **ST**
TYPE BoF
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 4:30 PM - 5:30 PM central
DURATION 60 min

AUTHORS
Jay Srinivasan (Lawrence Berkeley National Laboratory) jsrinivasan@lbl.gov
Douglas Jacobsen (Lawrence Berkeley National Laboratory) dmjacobsen@lbl.gov

ABSTRACT
NERSC's pre-exascale Perlmutter system and other future exascale systems will feature a new System Software stack. This BoF provides a forum to discuss the requirements of the System Software stack and system management features that will be used at NERSC and other centers for upcoming pre-exascale and exascale systems. Discussion topics will also include best practices and approaches to managing large systems for heterogeneous systems and a diverse set of user workloads and workflows and the effective use of system vendor/customer partnerships.

DESCRIPTION
Future pre-exascale and exascale systems will require a new System Software stack that balances the needs of a diverse set of users and allow for reliable operation and management at scales heretofore unseen. The process of evaluating such a system software stack has already begun. The purpose of this BoF session is to provide a forum to discuss the requirements of the System Software stack and system management features.

NERSC and other centers for that will deploy pre-exascale and exascale systems will participate and since the discussion topics will include best practices and approaches to managing large systems for heterogeneous systems and a diverse set of user workloads and workflows and the effective use of system vendor/customer partnerships, we expect this session to be of interest to those sections of the ECP community also involved in system deployment and management of those systems at scale, including the HI area community.

TITLE **The Compiler Will Help You! The LLVM Compiler Infrastructure in ECP**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 1:30 PM - 3:00 PM central
DURATION 90 min

AUTHORS
Hal Finkel (Argonne National Laboratory) hfinkel@anl.gov
Johannes Doerfert (Argonne National Laboratory) jdoerfert@anl.gov
Michael Kruse (Argonne National Laboratory) mkruse@anl.gov
Patrick McCormick (Los Alamos National Laboratory) pat@lanl.gov
George Stelle (Los Alamos National Laboratory) stelleg@lanl.gov
Jeffrey Vetter (Oak Ridge National Laboratory) vetter@ornl.gov
Joel Denny (Oak Ridge National Laboratory) dennyje@ornl.gov
David Poliakoff (Sandia National Laboratories) dzpolia@sandia.gov
Stephen Scalpone (NVIDIA Corporation) sscalpone@nvidia.com
Alexis Perry-Holby (Los Alamos National Laboratory) aperry@lanl.gov

ABSTRACT
The open-source LLVM compiler infrastructure will play a foundational role in the programming environments for DOE's exascale systems. Not only will open-source compilers and tools be available, but LLVM components are increasingly underlying vendor-provided toolchains. ECP is enhancing LLVM for better OpenMP/GPU performance, better support for loop transformations and autotuning, better support for Fortran (i.e., using Flang), and enhancing a wide spectrum of other capabilities relevant to exascale computing. This breakout will start with an overview of LLVM and how ECP is preparing LLVM for exascale. Representatives from the relevant software-technology projects will explain how to access interesting functionality and we'll collect feedback from breakout participants.

DESCRIPTION
There are several ECP subprojects working to enhance LLVM for our upcoming exascale systems. These include SOLLVE (2.3.2.11) targeting OpenMP, PROTEAS and Y-TUNE working on autotuning, loop optimizations, parallelism enhancements, interprocedural optimizations, and more [all part of the PROTEAS-TUNE (2.3.2.10) project as of the start of FY20], Flang (2.3.2.12) working on Fortran support, LANL's ATDM tools project (2.3.6.01), with contributions from other projects as well. These tools will be available as open-source projects on exascale machines, and also, many of these components are being incorporated into vendor toolchains for delivery on our exascale systems. Nearly all users will use LLVM-based technologies to run programs on exascale machines whether they know it or not. ECP's investments in this area are taking advantage of a key opportunity to inject enhancements important to our users into the entire exascale ecosystem in order to benefit our users. This breakout will educate users about the enhancements being developed in LLVM and collect feedback from users on their priorities and places where further investment might be needed. Those working on compilers and tools can use this opportunity to update each other and identify new collaboration opportunities.

The target audience of this sessions is two-fold: First, developers working on application-development projects and software-technology projects who are interested in learning more about compiler technologies that will be available on exascale systems, and moreover, who might wish to provide feedback on their users, desires, pain points, etc. Second, developers working on software-technology

projects related to compilers and tools in order to learn about the most-recent work of their peers on other projects. In order to address the different needs of these two groups, the breakout will start with information and material targeted at the first group. Toward the end of the session, we'll transition to conversations of interest to compiler and tooling experts.

The goals of the session are primarily to educate users of compiler technology and tools relevant to LLVM (which is nearly all developers across ECP), and secondarily to serve as a synchronization point for collaboration between ECP/ST projects in the space of compilers and tools.

TITLE **The Scope and Role of DevOps in HPC**

FOCUS AREA **ST**
TYPE Breakout
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 3:30 PM central
DURATION 60 min

AUTHORS
Michael Heroux (Sandia National Laboratories) maherou@sandia.gov
Manoj Bhardwaj (Sandia National Laboratories) mkbhard@sandia.gov
Todd Munson (Argonne National Laboratory) tmunson@mcs.anl.gov
Ryan Adamson (tentative) (Oak Ridge National Laboratory) adamsonrm@ornl.gov

ABSTRACT
The transition to accelerated architectures includes a highly disruptive change in programming and execution models as we migrate and create software for host-device, CPU-GPU platforms. In this transition, almost every line of code will be refactored in some way. As a result, the quality of the processes and tools we use for development, integration and operations heavily impact our success.

In this session, we discuss the scope and role of DevOps in our efforts to create software capabilities for next-generation platforms. The purpose of the session is to start a conversation between people across the DOE labs, universities and industry partners who have a stake in DevOps efforts. We will begin the session with introductions, then discuss the definition, scope and role of DevOps (which seems to vary across institutions) and end with a set of next steps for those in the community who want to pursue sharing and collaboration in DevOps activities.

TITLE **Tools and Techniques for Application Performance Characterization**

FOCUS AREA **HI**
TYPE Breakout
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 8:30 AM - 9:30 AM central
DURATION 60 min

AUTHORS
Scott Pakin (Los Alamos National Laboratory) pakin@lanl.gov
Shirley Moore (Oak Ridge National Laboratory) mooresv@ornl.gov
Samuel Williams (Lawrence Berkeley National Laboratory) swwilliams@lbl.gov
Jeanine Cook (Sandia National Laboratories) jeacook@sandia.gov
Kevin Barker (Pacific Northwest National Laboratory) Kevin.Barker@pnnl.gov
Andrew Chien (Argonne National Laboratory) achien@anl.gov

ABSTRACT
The ECP Analytical Modeling working group houses expertise in a variety of tools and techniques for gaining insight into application performance. In this talk, we showcase a number of noteworthy approaches for identifying application performance limiters on both current and future computer architectures and present the results of some specific case studies.

DESCRIPTION
There is far more to application performance analysis than executing an application on existing hardware and measuring how long each subroutine takes to run or what computational resources it consumes. The goal of this session is to educate the audience on methods that lie off the beaten path for answering performance questions that computational scientists may have about their codes and that can prove useful as applications are developed and optimized. The target audience is application developers who are interested in preparing their codes now for future supercomputers.

TITLE **Understanding Performance with Exa-PAPI**

FOCUS AREA **ST**
TYPE Breakout
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 4:00 PM - 5:30 PM central
DURATION 90 min

AUTHORS
Heike Jagode (The University of Tennessee, Knoxville) jagode@icl.utk.edu
Anthony Danalis (The University of Tennessee, Knoxville) adanalis@icl.utk.edu
Damien Genet (The University of Tennessee, Knoxville) dgenet@icl.utk.edu
Tony Castaldo (The University of Tennessee, Knoxville) TonyCastaldo@icl.utk.edu

ABSTRACT
The goals of this session are to inform attendees about the most recent developments and activities of the Exascale Performance Application Programming Interface (Exa-PAPI) project.

1. The development of the new SDE API enables ECP software layers to expose software-defined events that performance analysts can use to form a complete picture of the entire application performance.
2. Performance counter monitoring capabilities for new ECP hardware: AMD and NVIDIA GPUs
3. Power monitoring capabilities for new ECP hardware: IBM Power 9
4. Counter Inspection Toolkit (CIT): A tool to assist with native performance counter disambiguation through micro-benchmarks that are used to probe different important aspects of modern CPUs, which, ultimately, aids the classification of native performance events.

Each of these efforts will be presented in four 15 minutes talks. After the presentations, we will have a 30 minutes panel with the speakers as panelists to further explore the topics presented as well as PAPI's path forward with C++.

DESCRIPTION
The goals of this session are to inform attendees about the most recent developments and activities of the Exascale Performance Application Programming Interface (Exa-PAPI) project.

The PAPI package provides tool designers and application engineers with a consistent interface and methodology for the use of low-level performance counter hardware found across the entire compute system (i.e. CPUs, GPUs, on/off-chip memory, interconnects, I/O system, energy/power, etc.).

In this breakout, we will focus on four of PAPI's latest efforts:

1. The development of the new SDE API enables ECP software layers to expose software-defined events that performance analysts can use to form a complete picture of the entire application performance. Because software complexity is one of the fundamental issues the community is facing on the way to exascale, one of Exa-PAPI's central goals is to close the gap between SDE monitoring and hardware performance counter monitoring. The design and development of the new SDE API is nearly complete and, thanks to the strong feedback from various ECP teams, the Exa-PAPI team has successfully integrated SDE's into ECP applications (e.g., NWChemEx), math libraries (e.g., MAGMA-Sparse), and

runtimes (e.g., PaRSEC) to export important internal behavior. Exa-PAPI's new SDE functionality enables scientists to monitor the behavior of low-level linear algebra routines without the need for expert knowledge of the simulation code's full software stack.

2. Performance counter monitoring capabilities for AMD and NVIDIA GPUs:

- The development of two new components, called "rocm" and "rocm_smi", enables PAPI users to monitor hardware performance counters and monitor power usage on AMD GPUs.
- PAPI's latest accelerator support enables monitoring of NVIDIA GPU and NVLINK performance counters, power consumption, fan speed, temperature, and power capping support, to name but a few.

3. Power monitoring capabilities for IBM Power9:

- The development of the "powercap_ppc" PAPI component supports measuring and capping power usage on recent IBM PowerPC architectures (Power9 and later) using the powercap interface exposed through the Linux kernel.
- The newly developed PAPI "sensors_ppc" component supports reading system metrics on recent IBM PowerPC architectures (Power9 and later) using the opal/exports sysfs interface.

4. Counter Inspection Toolkit (CIT): A tool to assist with native performance counter disambiguation. This involves the development of various micro-benchmarks that are used to probe different important aspects of the architecture of modern CPUs, which, ultimately, aids the automatic classification of native performance events.

Each of these efforts will be presented in four 15 minutes talks. After the presentations, we will have a 30 minutes panel with the speakers as panelists to further explore the topics presented as well as PAPI's path forward, which implies the development of a new performance application programming interface, named PAPI++. PAPI++ builds upon classic-PAPI functionality and strengthens its path to exascale with a more efficient and flexible software design, one that takes advantage of C++'s object-oriented nature but preserves the low-overhead monitoring of performance counters and adds a vast testing suite.
The target audience is expected to be a mixture of ST tool developers, ECP application and library developer teams who are interested in performance analysis.

APPENDIX C. DETAILS OF TUTORIALS

The ECP agenda included an exciting series of 60, 90, 150, 180, and 330-minute tutorials offered Tuesday through Thursday of the annual meeting. When planning their schedule for the ECP annual meeting, attendees were encouraged to allow time to attend some tutorials.

See the list of tutorials below.

- A Tutorial Introduction to RAJA Performance Portability Abstractions
- Accelerating Your Application I/O with UnifyFS
- Achieving High Performance I/O with HDF5
- Analyzing ECP Proxy Apps with the Profiling Tool Score-P
- Application-Driven Fault-Tolerance for High Performance Distributed Computing
- Better Scientific Software Tutorial
- CANDLE Hands-on Tutorial
- Co-Design Center for Particle Applications (CoPA) Summary
- CODAR Hands-on Tutorial
- Container Computing for HPC and Scientific Workloads
- Customizing Data Services for Fun and Profit
- Developing Engaged and Productive Virtual Teams
- ECP Continuous Integration Startup Tutorial
- ExaGraph: Graph and Combinatorial Methods for Enabling Exascale Applications
- ExaLearn Tutorial: Co-Design Center for Exascale Machine Learning Technologies
- Exascale I/O with ADIOS: Online Processing, in Situ Visualization and High-Performance Storage I/O
- Flux: Using Next-Generation Resource Management and Scheduling Infrastructure for Exascale Workflows
- Frontier Tutorial
- Getting It Right with Open MPI - Best Practices for Deployment and Tuning of Open MPI
- Hands-on with Progress Tracking Cards: A Lightweight Method for Improving Your Software Practices
- In Situ Visualization and Analysis with Ascent
- Integrating PaRSEC-Enabled Libraries in Scientific Applications
- Introduction to the Capabilities and Use of the SUNDIALS Suite of Nonlinear and Differential/Algebraic Equation Solvers
- Managing HPC Software Complexity with Spack
- Managing Power Efficiency of HPC Applications with Variorum and GEOPM
- Memory Movement Orchestration and Topology-Aware Placement with AML
- Modern CMake Tutorial
- OpenMP 4.5 and 5.0 Tutorial (Offload)
- Performance Evaluation Using the TAU Performance System
- Performance Tuning with the Roofline Model on GPUs and CPUs
- SLATE and MAGMA Linear Algebra & FFT Libraries Tutorial

- STRUMPACK / SuperLU: Fast Parallel Direct Linear Solvers and Preconditioners
- Umpire: Managing Heterogenous Memory Resources
- UPC++: A PGAS/RPC Library for Asynchronous Exascale Communication in C++
- Using HPCToolkit to Measure and Analyze the Performance of GPU-Accelerated Applications

TITLE **A Tutorial Introduction to RAJA Performance Portability Abstractions**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner to Intermediate

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Richard Hornung (Lawrence Livermore National Laboratory) hornung1@llnl.gov
Arturo Vargas (Lawrence Livermore National Laboratory) vargas45@llnl.gov

ABSTRACT
Due to the diversity of computing architectures and programming models, the ability to develop single-source performance portable applications has become greatly important. This is particularly true for large production codes where developing and maintaining hardware-specific versions is untenable.

In this tutorial, we present RAJA, a library of C++ abstractions that enables developers to target multiple hardware and programming model back-ends with a portable single-source application code. Using only C++11 features, RAJA provides mechanisms to write computational loop kernels that can be recompiled to run on different hardware using different programming model back-ends without modifying application source code. RAJA provides abstractions for a wide range of loop structures found in numerical algorithms, and includes support for reductions, data layouts and views, iteration spaces, atomic operations, scans, etc. In this talk, we motivate RAJA features and design and provide working example codes for attendees to experiment with. The goal of the tutorial is to give participants sufficient understanding of how to use RAJA and how it works so they can begin to use it in their own applications.

DESCRIPTION
RAJA is an open-source software project (https://github.com/LLNL/RAJA) that has been in development for about 7 years. It provides performance portable C++ abstractions (requiring only standard C++11) that enable applications to target different hardware architectures and programming model back-ends (OpenMP, CUDA, HIP, TBB, etc.) with single-source application code. It supports a wide range of numerical algorithm loop structures found in HPC simulation codes as well as reductions, scans, atomic operations, multi-dimensional data views and layouts.

Currently, RAJA is in use in most of the ASC production application codes in the Weapons Simulation and Computing Program at Lawrence Livermore National Laboratory. It is also used in a variety of other applications at LLNL and the ECP: including the LLNL ECP/ATDM application, SW4 subsurface wave propagation code (ADSE19-EQSIM), GEOSX subsurface geomechanics code (ADSE05-Subsurface), SUNDIALS math library (WBS 2.3.3.12 SUNDIALS/Hypre), and the MFEM library (part of the ADCD04-CEED co-design project). It is funded primarily through the LLNL ASC and ATDM programs; it has also been funded as an ECP ST project (2.3.1.06 STPM08-RAJA) since the inception of the ECP.

The tutorial will introduce participants to RAJA concepts and interfaces by way of a variety of algorithm examples. The goal is to have participants learn enough about how RAJA works and how to use it to start using it in their own application development work. The target audience is any ECP software developer working on a C++ application or library. Familiarity with C++ templates and lambda expressions will help participants understand the material more quickly, as will familiarity with the basics of parallel programming for CPUs and GPUs. However, knowledge of these topics is not required.

TITLE **Accelerating Your Application I/O with UnifyFS**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Kathryn Mohror (Lawrence Livermore National Laboratory) kathryn@llnl.gov
Sarp Oral (Oak Ridge National Laboratory) oralhs@ornl.gov
Suren Byna (Lawrence Berkeley National Laboratory) sbyna@lbl.gov

ABSTRACT
UnifyFS is a user-level file system that highly-specialized for fast shared file access on HPC systems with distributed burst buffers. In this tutorial, we will present users with an introductory overview of how to use the lightweight UnifyFS file system to improve the I/O performance of their applications. We will begin at a high level describing how UnifyFS works with burst buffers and how users can incorporate it into their jobs. Following this, we will dive into more detail on what kinds of I/O UnifyFS currently supports and what we expect it to support in the future. We will also discuss the interoperation of UnifyFS with HDF5 and DataElevator. We will work with interested potential users in the tutorial and afterwards to learn to use UnifyFS.

DESCRIPTION
In this tutorial, we will present users with an introductory overview of how to use the lightweight UnifyFS file system to improve the I/O performance of their applications. We will begin at a high level describing how UnifyFS works with burst buffers and how users can incorporate it into their jobs. Following this, we will dive into more detail on what kinds of I/O UnifyFS currently supports and what we expect it to support in the future. We will also discuss the interoperation of UnifyFS with HDF5 and DataElevator. We will work with interested potential users in the tutorial and afterwards to learn to use UnifyFS.

Outline:

Part 1: Introduction
- What is UnifyFS and how can it help applications?
- Basic structure of UnifyFS and background on I/O in HPC

Part 2: How to use UnifyFS
- How to build and dependencies
- How to integrate UnifyFS with your application
- How to set up your job with UnifyFS
- How to run with UnifyFS

Part 3: Status of UnifyFS
- Current status of UnifyFS software
- Current status of interoperability with Data Elevator
- Contact information

Part 4: Audience questions and help for interested users

TITLE **Achieving High Performance I/O with HDF5**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 8:30 AM - 12:00 PM central
DURATION 180 min

AUTHORS
Quincey Koziol (Lawrence Berkeley National Laboratory) koziol@lbl.gov
Scot Breitenfeld (The HDF Group) brtnfld@hdfgroup.org
Elena Pourmal (The HDF Group) epourmal@hdfgroup.org
Suren Byna (Lawrence Berkeley National Laboratory) sbyna@lbl.gov
Kathryn Mohror (Lawrence Livermore National Laboratory) mohror1@llnl.gov

ABSTRACT
HDF5 is a data model, file format, and I/O library that has become a de facto standard for HPC applications to achieve scalable I/O and for storing and managing big data from computer modeling, large physics experiments and observations. Several ECP applications are currently using or planning to use HDF5 for I/O. This 3-hour tutorial offers an overview of HDF5 to anyone who works with I/O in an HPC environment. The Tutorial consists of four parts. Part I introduces the HDF5 data model, and APIs for organizing data and performing I/O. Part II gives an overview of parallel file systems and shows how to use HDF5 in a parallel environment to perform I/O on a shareable file or multiple HDF5 files. Part III uses examples from well-known codes and use cases from experimental sciences to demonstrate the tuning techniques such as collective metadata I/O, data aggregation, parallel compression and new HDF5 features that help to utilize HPC storage beyond current files systems (Data Elevator and UnifyFS).

DESCRIPTION
HDF5 (https://www.hdfgroup.org/solutions/hdf5/) is a data model, file format, and I/O library that has become a de facto standard for HPC applications to achieve scalable I/O and for storing and managing big data from computer modeling, large physics experiments and observations. HDF5 is among the top 10 software libraries used at supercomputing facilities, such as the National Energy Research Scientific Computing Center (NERSC), Argonne Leadership Computing Facility (ALCF), and Oak Ridge Leadership Computing Facility (OLCF). Achieving high performance I/O, effective usage of storage, managing data archives, and working with a diversity of parallel file systems presents a challenge to the majority of scientists and researchers who use HPC systems as a tool for their research or data acquisition and analysis. HDF5 hides this complexity and allows the scientists to focus on science instead of I/O and storage.

The tutorial provides a comprehensive introduction to HDF5 and HDF5 hands-on experience to students, researchers, applications developers who need to work with high volumes of simulated or experimental data stored in HDF5 in HPC environment. It features lessons learned as well as key techniques to work with huge volumes of data in HDF5 and to tune I/O performance on HPC systems. The outcome of this tutorial is for students to have a better understanding of how to use HDF5 to achieve faster I/O and efficient data management on the current and the upcoming Exascale systems. The tutorial will empower users with the fundamental knowledge of HDF5 and the best tuning techniques from the storage and file system layer all the way to the HDF5 API layer. The tutorial will also focus on how to acquire, store, and analyze HDF5 data on the HPC systems in almost real time.

This tutorial is designed to bring all participants to the level at which they should be able to use HDF5 with maximum efficiency. Part I (33%) targets beginners, Part II (33%) target beginners, intermediate and advanced users of HDF5 since it will cover the best practices and tuning techniques for both sequential and parallel I/O and efficient storage of big data. Part III (33%) targets beginners, intermediate, and advanced users by building on the first two tutorial parts, training attendees how to gain high levels of performance on HPC systems on examples of real use cases.

Outline

* **Part I** (Elena Pourmal, Scot Breitenfeld) - 1 hr

We will give an overview of the HDF5 data model, software, and format, to ensure that all attendees (new users and users that have some experience with HDF5) have the same basic understanding and a common vocabulary, and prepare attendees for the advanced parts of the tutorial.

- Introduction to HDF5 data model, file format and architecture (20 min)
- Introduction to HDF5 basic APIs to store and organize data in HDF5 (20 min)
- Introduction to HDF5 I/O including partial I/O (20 min)

* **Part II** (Quincey Koziol, Scot Breitenfeld) - 1 hr

This part is designed for users who have had exposure to MPI I/O and would like to learn about the parallel HDF5 library. It will cover parallel HDF5 design, programming models and APIs. C and Fortran examples for different I/O access patterns will be used to demonstrate the capabilities of the HDF5 parallel library. Basic tuning techniques and tools for performance debugging will be demonstrated.

- Overview of parallel file systems, MPI I/O, and parallel HDF5 architecture (20 min)
- Using HDF5 to write to shared files with different access patterns; employing internal compression for parallel I/O (20 min)
- Understanding I/O performance and debugging and achieving performance (20 min)

* **Part III** (Quincey Koziol, Suren Byna, Scot Breitenfeld, Kathryn Mohror) - 1 hr

In this part of the tutorial, we will discuss in depth HDF5 performance on parallel file systems and demonstrate optimization techniques.

- Parallel I/O tuning techniques (20 min)
- Collective metadata I/O
- Collective vs. independent raw data I/O
- Load balancing, data aggregation, open/close/flush optimizations
- Techniques in action: Use cases based on VPIC, Chombo, CGNS (20 min)
- Using new storage layers for accelerating I/O (20 min)
- VOL (Data Elevator)
- UnifyFS

TITLE **Analyzing ECP Proxy Apps with the Profiling Tool Score-P**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 1:30 PM - 4:30 PM central
DURATION 150 min

AUTHORS
George Markomanolis (Oak Ridge National Laboratory) markomanolig@ornl.gov
Michael Brim (Oak Ridge National Laboratory) brimmj@ornl.gov

ABSTRACT
Score-P is a community-developed source code instrumentation infrastructure that supports the profiling and tracing capabilities of several HPC performance analysis tools. Score-P allows for instrumentation of many parallel programming models, including MPI, OpenMP, CUDA, OpenACC, and several others. In profiling mode, it provides summary information for application performance characterization at relatively low overhead. In-depth tracing is also supported, which produces OTF2 traces that can be visualized by the scalable trace viewer Vampir. The purpose of this tutorial is to showcase the functionalities of Score-P and Vampir, including specific instructions on how to instrument various application codes. We will also report any issues encountered and suggested workarounds for known problems. The tutorial will relate our experience in applying a workflow to do performance analysis and identify bottlenecks on the Summit Supercomputer using ECP proxy applications.

DESCRIPTION
Many users of HPC resources often complain that a profiling tool is not working as expected because during the training they attended, the presenters used benchmark with simple cases. We took this opportunity to prepare the current tutorial. We are advanced users of profiling tools and we plan to demonstrate in this tutorial the strengths and weakness of the Score-P profiling tool. We present the full workflow of compilation, instrumentation, and performance analysis of various ECP proxy applications while using the Summit Supercomputer. The ultimate purpose is to showcase the difficulties of using such tools, the methodology to bypass some technical issues when possible, and how Score-P provides the required data to understand better the insights of the applications. Moreover, through this tutorial, a scientist will be possible to answer some questions about applications, such as “why is it slow?” “why it does not scale?” “what part of my code should I optimize?”. This tutorial gave us the opportunity to provide feedback to the developers, to understand what is missing of such tools and how they could be improved. Moreover, we will present lessons learned of this effort and what the users should expect. Finally, there will be a short hands-on session to get more familiar and feel confident that you can use this profiling tool. The targeted audience are scientists who develop applications and would like to know if their application is efficient or more optimization is required to finish their experiments faster and for scalability purposes.

TITLE **Application-Driven Fault-Tolerance for High Performance Distributed Computing**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
George Bosilca (The University of Tennessee, Knoxville) bosilca@icl.utk.edu
Bogdan Nicolae (Argonne National Laboratory) bnicolae@anl.gov

ABSTRACT
As the age of Exascale draws closer and the size of large-scale, distributed applications continues to increase, so does the failure rate and thus the need for advanced resilience techniques to handle them. Over the last years, the resilience topic evolved from an open question to a clear requirement where the failure occurrences are not questioned anymore, but instead the focus is on the frequency of such radical events during the execution of applications at scale. Solutions to transparently manage faults at the system level exist, but their scalability potential and overhead in terms of performance and resource utilization remains high, even for low failure frequency. Therefore, empowering the developers to deal with the failures at application-level instead brings more opportunities to reduce the resilience overhead that needs holistic support from all layers: hardware and software as well as from the parallel programming paradigm. This tutorial highlights application-driven solutions to survive faults and provide a basic understanding of their expected costs at scale. The presented solutions cover two complementary approaches (1) application-defined checkpoint-restart (as demonstrated through the VeloC runtime); and (2) user-level failure mitigation (as demonstrated through ULFM runtime).

DESCRIPTION
This tutorial is a joint effort between 2 ECP projects: OMPI-X and VeloC.

Level: The content is structured such that it covers 30% beginner, 40% intermediate and 30% advanced level. Some previous knowledge with MPI and the C programming language is necessary.

Participants need to have a laptop and be able to download Docker and a docker capsule.

TITLE **Better Scientific Software Tutorial**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL N/A

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
David Bernholdt (Oak Ridge National Laboratory) bernholdtde@ornl.gov
Anshu Dubey (Argonne National Laboratory) adubey@anl.gov
James Willenbring (Sandia National Laboratories) jmwille@sandia.gov
Paul Wolfenbarger (Sandia National Laboratories) prwolfe@sandia.gov

ABSTRACT
Producing scientific software is a challenge. The Exascale Computing Project, in particular, encompasses the confluence of disruptive changes in computing architectures and new opportunities (and demands) for greatly improved simulation capabilities, especially through coupling physics and scales. At the same time, computational science and engineering (CSE), as well as other areas of science, are experiencing increasing focus on scientific reproducibility and software quality.

Computer architecture changes require new software design and implementation strategies, including significant refactoring of existing code. Reproducibility demands require more rigor across the entire software endeavor. Code coupling requires aggregate team interactions including integration of software processes and practices. These challenges demand large investments in scientific software development and improved practices. Focusing on improved developer productivity and software sustainability is both urgent and essential.

This half-day tutorial distills multi-project and multi-years experience from members of the IDEAS Productivity project and the creators of the BSSw.io community website. The tutorial will provide information about software practices, processes, and tools explicitly tailored for CSE. Topics to be covered include: Agile methodologies and tools, software design and refactoring, testing and test-driven development (TDD), and Git workflows for teams. Material will be mostly at the beginner and intermediate levels. There will also be opportunities to discuss topics raised by the audience.

DESCRIPTION
We are proposing a half-day variant of our Better Scientific Software tutorial, which we have been presenting and refining since 2016 at venues including ATPESC (2016-2019), SC (2016-2019), SIAM CSE (2017), and prior ECP AMs (2018-2019). The proposed set of topics was developed in discussion with Ashley Barkey, but we have other modules we could present instead (or in addition), if the committee recommends. Options include (in no particular order): reproducibility, licensing, documentation strategies, code coverage, requirements gathering.

A basic goal of this tutorial is to give participants tools and strategies to start thinking systematically about their software development practices and realize that they don,Äôt have to be overwhelmed by the challenges they face. The material will be at the beginner and intermediate levels, making it accessible to a wide range of ECP AM attendees, but particularly relatively junior participants who may be less experienced with all aspects of computational science and engineering software development. There are

no particular prerequisites for attendees. Because of the 3-hour time frame, we do not plan significant hands-on content.

This tutorial will be complementary (not overlapping) with other activities proposed by the IDEAS team, as well as the Continuous Integration team.

TITLE **CANDLE Hands-on Tutorial**

FOCUS AREA **AD**
TYPE Tutorial
LEVEL Intermediate

DAY Tuesday, February 04, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Emily Dietrich (Argonne National Laboratory) edietrich@anl.gov
John Gounley (Oak Ridge National Laboratory) gounleyjp@ornl.gov
Hong-Jun Yoon (Oak Ridge National Laboratory) Yoonh@ornl.gov
Justin Wozniak (Argonne National Laboratory) woz@anl.gov
Jamaludin Mohd Yusof (Los Alamos National Laboratory) jamal@lanl.gov
Tom Brettin (Argonne National Laboratory) brettin@anl.gov

ABSTRACT
The Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer project focuses on building a scalable deep neural network application called the CANDLE (CANcer Distributed Learning Environment) through the Exascale Computing Project (ECP) in a collaboration of four DOE national laboratories with the NCI and the NCI-supported Frederick National Laboratory for Cancer Research. The challenge problem for CANDLE is to enable the most challenging deep learning problems in cancer research to run on the most capable supercomputers in the DOE and the National Institute of Health (NIH). Initially developed to address three top challenges facing the cancer community, CANDLE increasingly can be used to tackle deep learning problems of various other application areas.

Implementations of CANDLE have been tested on current and past DOE systems Titan, Cori, Theta, and Summit, and on the NIH Biowulf system. The CANDLE software builds on open source deep learning frameworks including Keras, TesnsorFlow and PyTorch. Through collaborations with DOE computing centers, HPC vendors and ECP co-design and software technology projects, CANDLE is being prepared for the coming DOE exascale platforms. Features currently supported in CANDLE include feature selection, hyperparameter optimization, model training, inferencing and UQ. Future release plans call for supporting experimental design, model acceleration, uncertainty guided inference, network architecture search, synthetic data generation and data modality conversion. These features have been used to evaluate over 20,000 models in a single run on a DOE system.

The CANDLE hands-on tutorial will dive into the CANDLE environment, including large-scale model search, inferencing, natural language processing, and discussion of ongoing work with neural architecture search. As a hands-on tutorial, topics will include preparing a deep neural network to run in CANDLE, hyperparameter optimization, model ensembles for uncertainty quantification (UQ), and data parallelism. Discussions will be aimed at understanding of current deep learning research, specific projects, and key requirements of workshop participants. This will allow the projects to pursue development and advancement in a complementary and accelerated manner with CANDLE. There will be plenty of time for questions and answers.

TITLE **Co-Design Center for Particle Applications (CoPA) Summary**

FOCUS AREA **AD**
TYPE Tutorial
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 8:30 AM - 12:00 PM central
DURATION 180 min

AUTHORS
Susan Mniszewski (Los Alamos National Laboratory) smm@lanl.gov
Jim Belak (Lawrence Livermore National Laboratory) belak1@llnl.gov
Stuart Slattery (Oak Ridge National Laboaratory) slatterysr@ornl.gov
Christoph Junghans (Los Alamos National Laboratory) junghans@lanl.gov
Bob Bird (Los Alamos National Laboratory) bird@lanl.gov

ABSTRACT
Particle-based algorithms are becoming ubiquitous in the Exascale Computing Project (ECP). From molecular dynamics to continuum mechanics and plasma physics, these algorithms are growing in popularity due to their numerical advantages for certain classes of problems as well as the advent of more powerful supercomputers that are able to mitigate their expense relative to more traditional computational methods. In this tutorial we will overview the ECP Co-Design Center for Particle Applications (CoPA). We will present the structure of CoPA as well as the current particle landscape in ECP including CoPA core application partners as well as other particle-based work in ECP. We will then introduce the two main software components being developed by CoPA, Cabana and Progress-BML, and discuss their design and development strategy. Finally, we will present results from several CoPA core application partners outlining the development and deployment accomplishments achieved in the last several years.

DESCRIPTION
In this session, we would like to provide a summary of the ECP Co-design Center for Particle Applications (CoPA) describing the project at a high level, details on our library development per the Cabana Particle library and PROGRESS/BML libraries, followed by examples of applications that are using our capabilities. The goal is to provide an understanding of the CoPA project to the ECP community and provide the contacts for questions and potential collaborations.
CoPA is addressing challenges for particle-based applications to run on upcoming exascale computing architectures. This scope is partitioned into four “sub-motifs”: short-range particle-particle interactions (e.g., those which often dominate molecular dynamics and smooth particle hydrodynamics methods), long-range particle-particle (e.g., electrostatic and gravitational) interactions, particle-in-cell (PIC) methods, and linear-scaling electronic structure and quantum molecular dynamics (QMD) algorithms. CoPA provides two types of crosscutting co-designed technologies, proxy apps and libraries.

Libraries are modular instantiations that multiple applications can utilize or be built upon. Cabana is a performance portable library for particle-based simulations (github.com/ECP-copa/Cabana). Applications include but are not limited to Molecular Dynamics (MD) with either short- and/or long-range interactions and various flavors of Particle-in-Cell (PIC) methods, including applications to fluid and solid mechanics and plasma physics. Cabana provides particle data structures, algorithms, and utilities to enable simulations on a variety of platforms including many-core architectures and GPUs. The PROGRESS library provides quantum chemistry solvers for QMD (github.com/lanl/qmd-progress). It uses the basic matrix library package (BML) which provides a common application programming interface (API) for

linear algebra and matrix functions on CPU-based and CPU-GPU architectures and is matrix format independent (github.com/lanl/bml).

The CoPA Cabana Particle Simulation Toolkit has been demonstrated for WDMapp's XGC science study of full-scale Summit for a production case tokamak plasma and geometry run. New capabilities include a Fortran interface for Cabana, rewrite of the GPU kernels in XGC, portability across other architectures, and full-scale execution of XGC on Summit with improved performance.

Additionally, the SWFFT and fftMPI parallel fast Fourier transform libraries have been contributed as part of CoPA. Their performance benefits are demonstrated in ExaSky/HACC.

The CoPA PROGRESS/BML Libraries have been demonstrated for the EXAALT/LAMMPS/LATTE challenge problem on Summit, the DFTB+ quantum chemistry package, and the MGmol finite difference pseudopotential DFT code. Performance portability and improved performance taking advance of the GPU acceleration on full-scale Summit are shown.

The agenda is as follows:

- Overview and CoPA Tools
- Co-design Center for Particle Applications (CoPA) Overview – 30 min
- Cabana Particle Simulation Toolkit – 30 min
- PROGRESS/BML Libraries – 30 min

Applications and Examples:

- WDMapp/XGC – 30 min
- ExaSky/HACC – 30 min
- EXAALT/LAMMPS/LATTE – 30 min

TITLE **CODAR Hands-on Tutorial**

FOCUS AREA **AD**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Kshitij Mehta (Oak Ridge National Laboratory) mehtakv@ornl.gov
Igor Yakushin (Argonne National Laboratory) iyakushin@anl.gov
Matthew Wolf (Oak Ridge National Laboratory) wolfmd@ornl.gov
Jong Yuol Choi (Oak Ridge National Laboratory) choij@ornl.gov
Eric Suchyta (Oak Ridge National Laboratory) suchytaed@ornl.gov
Todd Munson (Argonne National Laboratory) tmunson@mcs.anl.gov
Ian Foster (Argonne National Laboratory) foster@anl.gov
Scott Klasky (Oak Ridge National Laboratory) klasky@ornl.gov

ABSTRACT
The ever-increasing data sizes generated by scientific applications have prompted a shift in data analysis paradigms, from post-hoc ("offline") to online reduction and/or analysis. This shift raises important questions, such as what compression and analysis methods to use, and how to orchestrate online reduction and analysis workflows effectively. The ECP Co-design center for Online Data Analysis and Reduction (CODAR) project aims to assist ECP application projects to understand the characteristics of different data compression and reduction algorithms, the tradeoffs associated with in situ and online analysis, and methods for efficient mapping of online workflows to exascale computers.

In this tutorial, attendees will first learn about Savanna, an in situ runtime engine for executing coupled application workflows. Presenters will explain how Savanna controls mapping application components to resources, mapping processes to cores, and organizing interprocess communications. Attendees will also learn about Cheetah, a tool for defining and running parametric codesign studies (a campaign) of data compression and reduction methods for coupled application workflows using different orchestration strategies. Attendees will learn how to define abstract, Python-based specifications of large campaigns of experiments in a simple way, and how to launch and monitor their campaigns. Users will learn how to perform complex process mapping on a heterogenous system such as the Summit supercomputer at OLCF. Online data will be reduced and analyzed using multiple options, including SZ, ZFP, and MGARD. Prior experience with these tools is not required.

DESCRIPTION
As science applications continue to generate vast amounts of data, I/O and data management have become significant challenges as we progress towards exascale. This has led researchers to increasingly look towards analyzing data online as opposed to the traditional process of analyzing it post-hoc. However, online data analysis poses several questions in terms of the right compression and reduction methods suited for different application domains, the different tradeoffs associated with accuracy and application performance, and optimal ways of orchestrating an online analysis workflow.

The CODAR project seeks to answer these questions by running codesign campaigns of coupled application workflows that explore various options in each of the three categories listed above. Applications are coupled using the ADIOS middleware library. It provides a way to apply different

compression methods such as SZ, ZFP, MGARD etc. on output data, along with fine-tuning their accuracy settings. Lastly, we seek to determine how to best orchestrate a workflow, which involves synchronous versus asynchronous reduction, inline vs online reduction, and on-node vs off-node reduction of data.

To test the many possible workflow composition and orchestration methods, we have developed Cheetah, a tool for running campaigns of coupled application workflows that lets users explore various parameters associated with online data analysis, and Savanna, a runtime engine for executing coupled workflows. In this tutorial, presenters will talk in-depth about Cheetah and its dependent suite of tools. Attendees will learn how to use Cheetah's abstract Python-based specification format and API to set up campaigns of experiments for exploring various data reduction and workflow orchestration mechanisms. With minimal modifications, a campaign can be adapted to different supercomputers. Presenters will show how to use different systems such as Summit, Cori, and local Linux machines. Additionally, attendees will learn how Cheetah can be used for campaign monitoring and generating a performance report for a completed campaign. Examples and mini-apps will be presented to demonstrate the full workflow of setting up codesign campaigns, executing them, and obtaining results from the campaigns using Cheetah.

The target audience are developers, performance engineers, and scientists who want to perform codesign studies for online data analysis. Application users who want to transform their traditional offline workflows to online analysis workflows will benefit from this session. As many factors determine the advantages and adaptability of online data processing, attendees will learn how to utilize Cheetah to perform a systematic study of the impact of different parameters. This will help them optimize their application workflow for the target architecture. In addition to using Cheetah for running codesign campaigns, they will learn about tools for assessing the quality of their reduced data. Finally, we will show how to utilize an optimal configuration determined by Cheetah for production runs.

TITLE **Container Computing for HPC and Scientific Workloads**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS

Shane Canon (Lawrence Berkeley National Laboratory) scanon@lbl.gov
Sameer Shende (University of Oregon) sameer@cs.uoregon.edu
Andrew Younge (Sandia National Laboratories) ajyoung@sandia.gov

ABSTRACT

Containers have gained traction in HPC because they promise significantly greater software flexibility, reliability, and portability for users. Tools like Docker, and Shifter, Charliecloud and Singularity enable users to leverage containers for scientific and technical computing use cases. However, to fully unlock its potential, users and administrators need to understand how to utilize these new approaches. This tutorial includes an overview of linux container technology, options for containers in HPC environments and use cases they are enabling, as well as an interactive period to try out containers on HPC resources. Using the ECP SDK using the E4S containers with these container runtimes will also be covered. A key goal of the tutorial is to provide participants with hands-on experience with linux containers on HPC resources and how this can enable their scientific workflows. The tutorial will also explain how research scientists can utilize container-based computing to accelerate their research and how these tools can boost the impact of their research by enabling better reproducibility and sharing of their scientific process.

DESCRIPTION

Outline

1. Introduction
    1. Google Question Doc
    2. Audience Poll
2. Overview of linux containers (10 minutes)
    1. How containers work
    2. Why containers are useful to scientists
    3. What do containers solve and not solve
3. Building container images (10 minutes)
    1. Docker
    2. Singularity
    3. Buildah
4. Managing images with DockerHub (10 minutes)
5. Hands-on Basics (45 minutes)
    1. Use laptop to build, run and ship an image
6. Break - 15 minutes
7. Container runtimes and HPC (15 minutes)
    1. Common Issues with Containers and HPC
    2. Shifter
    3. Charliecloud
    4. Singularity

8. Hands-on HPC (45 minutes)
    1. Using Shifter on Cori
    2. Using Charliecloud on Cori
    3. Using Singularity
9. E4S Demo
10. Advanced Use Cases (15 minutes)
    1. Building containers that use MPI and GPUs
    2. Considerations for non-x86 architectures (e.g. ARM, PowerPC)
11. End User Examples (10 minutes)
12. Wrap Up

TITLE **Customizing Data Services for Fun and Profit**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 8:30 AM - 4:30 PM central
DURATION 330 min

AUTHORS
Philip Carns (Argonne National Laboratory) carns@mcs.anl.gov
Rob Latham (Argonne National Laboratory) robl@mcs.anl.gov
Matthieu Dorier (Argonne National Laboratory) mdorier@anl.gov
Shane Snyder (Argonne National Laboratory) ssnyder@mcs.anl.gov
Bob Robey (Los Alamos National Laboratory) brobey@lanl.gov
Rob Ross (Argonne National Laboratory) rross@mcs.anl.gov

ABSTRACT
Science teams often find themselves forced to adapt to the quirks of parallel file systems and a limited set of I/O libraries that can be difficult to use and even more difficult to use with high performance. Looking outside the HPC community, we see industry assembling solutions from a variety of technologies, including object stores, column stores, document stores, and other types of services.

The Mochi (https://www.mcs.anl.gov/research/projects/mochi/) project is enabling the rapid development of custom data services for science teams, and a handful of such data services have now been developed. In this tutorial we will present an overview of the Mochi technologies available to date, help attendees run their first Mochi services on a current platform, and work with these teams towards the design of a custom service meeting their specific requirements. This would kick start interactions towards the implementation and integration of this service.

DESCRIPTION
We will spend approximately 2/3rds of the time discussing Mochi and its components and getting our attendees to build and run on a platform at ALCF, OLCF, or NERSC (we are assuming they will have an account at least one of these facilities). We will have already run these demonstrations on all these platforms prior to the Annual Meeting, so the exercise is really just in getting environments right for our attendees.

Next, we will split into small groups and work out a high-level design for a specific service. Based on the components we anticipate using, we would then help teams build and execute additional sample codes using the relevant components and develop a plan for how we might develop this service going forward.

Because of the hands-on / small group aspects of this discussion we cannot handle more than 4 small teams of people in this activity. We would prefer to be in a "cozy" room for this meeting, something that comfortably seats 12-15. We have already identified teams that are interested in participating in this event. Additional teams interested in participating should contact **Rob Ross** <rross@mcs.anl.gov> to discuss.

TITLE **Developing Engaged and Productive Virtual Teams**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 1:30 PM - 4:30 PM central
DURATION 150 min

AUTHORS
David Moulton (Los Alamos National Laboratory) moulton@lanl.gov
Elaine Raybourn (Sandia National Laboratories) emraybo@sandia.gov
Todd Munson (Argonne National Laboratory) tmunson@mcs.anl.gov

ABSTRACT
In this breakout session, we build on our experiences and lessons learned after the "Team of Teams" and "Fostering a Culture of Passion and Productivity in ECP Teams" break out sessions presented at past ECP Annual Meetings to present strategies and interactive activities centered around developing engaged and productive virtual teams. We explore how meeting ECP technical challenges depend critically on meeting social and cultural challenges as well. We will explore these social and cultural development techniques including, informal "tea-time", virtual scrums, human-centered design, empathy building and listening games, and exercises. In addition, we will show how these techniques enable scaling to meet the additional challenges associated with managing a team-of-teams project, with a focus on virtual teams. We will also delve into effective virtual meetings and brainstorming sessions, managing personalities and team dynamics, the role of the liaison and contrarian, and the diffusion of ideas within a team of teams. This session will be broken into two 90 minute segments. Each segment will have roughly 60 minutes of guided small group exercises, followed by roughly 30 minutes of discussion with the panel. In particular, we are developing a panel comprised of successful ECP small team project leads who have experience in fostering a culture that promotes engaged and productive virtual teams.

DESCRIPTION
The ECP core mission is to develop a capable exascale computing ecosystem that accelerates scientific discovery and supports addressing critical challenges in energy and national security. The very nature of this mission has drawn a wide range of talented and successful scientists with diverse backgrounds to work together in new ways to push beyond the status-quo toward this goal. Many of these scientists achieved success through efficient and agile collaborations with a (very) small trusted team that could rapidly innovate, prototype and deliver. However, in this context delivery was often focused on scientific understanding and journal publications, and placed lower expectations on deploying the corresponding capability as a well-documented component or library for other teams to use and integrate. Yet, for the ECP software ecosystem to succeed and truly accelerate scientific discovery, small teams must meet both of these deliverables, while meeting additional challenges as they embrace a new role within their larger ECP project.

In this breakout session, we build on our experiences and lessons learned after the "Team of Teams" and "Fostering a Culture of Passion and Productivity in ECP Teams" breakout sessions presented at past ECP Annual Meetings to present strategies and interactive activities centered around developing engaged and productive virtual teams. We explore how meeting ECP technical challenges depend critically on meeting social and cultural challenges as well. Specifically, many important cultural aspects of very small teams evolve naturally on a positive path because the team was formed by friendly collaborators with common interests and goals. It is relatively straightforward to maintain an aligning narrative that captures the

team's goals, approach and accepted practices because everyone is immersed in it. But as the team grows it faces new challenges, including fragmented time commitments, physically distributed developers, and growing numbers of customers or stakeholders. These challenges make it significantly harder to maintain an aligning narrative and healthy team culture, and although agile methodologies and modern software tools can help developers manage these challenges, these processes alone are insufficient. Thus, a growing number of teams are combining these development methodologies with techniques that support an engaged, prodictive, and collaborative culture.

We will explore these social and cultural development techniques including, informal "tea-time", virtual scrums, human-centered design, empathy building and listening games, and exercises. In addition, we will show how these techniques enable scaling to meet the additional challenges associated with managing a team-of-teams project, with a focus on virtual teams. We will also delve into effective virtual meetings and brainstorming sessions, managing personalities and team dynamics, the role of the liaison and contrarian, and the diffusion of ideas within a team of teams. Pointers to the research behind these activities will be provided.

This session will be broken into two 90-minute segments. Each segment will have roughly 60 minutes of guided small group exercises, followed by roughly 30 minutes of discussion with the panel. In particular, we are developing a panel comprised of successful ECP small team project leads who have experience in fostering a culture that promotes engaged and productive virtual teams. The panel members will serve as facilitators for the small group exercises, which include building trust, effective communication and brainstorming, and exploring and managing personality impacts on team dynamics. The discussion with the panel will recap lessons learned from the exercises, and explore how these critical aspects of culture affect the performance of the small teams, as well as its effectiveness in the increasingly common interdisciplinary team-of-teams setting. Participants will take away an appreciation for the vital role that culture plays in overall productivity, and strategies they can use to foster an engaged and productive team.

TITLE **ECP Continuous Integration Startup Tutorial**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Paul Bryant (New Mexico Consortium) pbryant@newmexicoconsortium.org
Tiffany Connors (Lawrence Berkeley National Laboratory) tconnors@lbl.gov
Mario Melara (Lawrence Berkeley National Laboratory) mamelara@lbl.gov
Mark Miller (Lawrence Livermore National Laboratory) miller86@llnl.gov
David Nicholaeff (New Mexico Consortium) dnicholaeff@newmexicoconsortium.org
Kyle Shaver (Argonne National Laboratory) kshaver@anl.gov

ABSTRACT
Continuous Integration (CI) has become an industry standard, utilized by teams to better facilitate testing and integration of projects by using a robust, repeatable, and scalable process. The ECP CI project has been developing functionality, leveraging GitLab CI, and working with the DOE HPC Centers to bring capability to software teams and cross site testing.

For the past year the ECP CI project has been focusing on developing key functionality into GitLab in order to best support the HPC workflow model. This includes adding a feature allowing the runner to interface with different schedulers to submit jobs and ensuring jobs are executed under the appropriate user/service account. In addition, security features for cross-site user authentication and control have been integrated to support job submissions potentially targeting multiple remote facilities. We believe that with these enhancements HPC facilities and application developers can use CI thus improving software engineering practices, and ensuring a robust code base.

Join the members of the CI team as we provide an overview of the ongoing efforts, tutorials ranging from the basics of GitLab's CI framework to the creation of complex pipelines targeting multiple heterogeneous machines and architectures, and an open Q&A with the goal of providing a clear path forward for ECP teams interested in getting started. Please note that you will need access to an internet browser and SSH terminal in order to participate in the tutorials.

DESCRIPTION
The ECP CI effort has been working for over a year in adding needed security features and added capability to the GitLab CI products to allow ECP ST and AD teams to test their software across E6 Facility HPC machines. This has been in development and not available to software teams until very recently. This session is to provide an update, hands on tutorial, and getting started guidelines to ECP users that are interested in getting started with ECP CI.

Topics

- Brief introduction regarding the purpose of GitLab CI and the benefits it can provide.
- Update from the ECP CI project team to outline the current status of enhancements as well as site collaboration to better provide context to the tutorials.
- Provide an overview of on-going development efforts.

- Walk through of the ECP CI ecosystem workflow - The role of the Office of Scientific and Technical Information (OSTI), cross-site authentication procedures, and gaining access to site resources.
- (Hands-on) Connecting to training resources.
- (Hands-on) Setting up a project, generating a "gitlab-ci.yml" file, and developing an introductory level CI pipeline.
- (Hands-on) A more advanced CI utilization targeting different facility resources, queues, and policies. Focused on highlighting and helping to address the complexities present in cross-site testing.
- Question / answers regarding any of the topics covered, potential concerns, and desired functionality.
- Outline a path forward and introduce points of contact that can assist ECP AD/ST projects get started.

Goals

- Update on project status, ongoing development, and availability of cross-site resources.
- Provide a basic understanding of GitLab CI and how to get started.
- Insure ECP AD/ST attendees have a path forward and contacts provided to help get started after the event has concluded.

Target Audience

- ECP AD/ST project teams interested in leveraging the ECP CI infrastructure through OSTI's GitLab instance.
- Anyone with an interest in leveraging GitLab CI at their facility, optionally with the ECP enhancements.
- Those interested in using or participating in the ECP CI infrastructure.

TITLE **ExaGraph: Graph and Combinatorial Methods for Enabling Exascale Applications**

FOCUS AREA **AD**
TYPE Tutorial
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Mahantesh Halappanavar (Pacific Northwest National Laboratory) hala@pnnl.gov
Aydin Buluc (Lawrence Berkeley National Laboratory) abuluc@lbl.gov
Erik Boman (Sandia National Laboratories) egboman@sandia.gov
Alex Pothen (Purdue University) apothen@purdue.edu

ABSTRACT
Combinatorial algorithms in general and graph algorithms in particular play a critical enabling role in numerous scientific applications. The irregular memory access nature of these algorithms makes them one of the hardest algorithmic kernels to implement on parallel systems. In this tutorial, we will cover several aspects related to the design of parallel algorithms, introduction to software tools and their execution, and application of graph algorithms to solve exascale applications. We aim to stimulate the use of graph algorithms to solve novel problems in scientific computing and will target beginner and intermediate levels of expertise.

DESCRIPTION
We aim to cover algorithms for graph matching, coloring, clustering, partitioning and graph traversals. We will also include short overview of the tools that are available for each problem and provide basic information on their execution. We will cover relevant applications that exploit the graph algorithms to solve different aspects. For example, we will present the use of graph algorithms in the context of computational biology and sparse direct solvers.

The principal goal of this tutorial is to stimulate algorithmic thinking for novel users and to update intermediate users of the recent developments. We target audience at the beginner level to stimulate application of graph algorithms to their specific domain (or problem). We target intermediate level audience for further use of software tools developed by the ExaGraph project. We will also discuss the recent developments targeting GPU-based architectures (single and multi-gpu implementations).

TITLE **ExaLearn Tutorial: Co-Design Center for Exascale Machine Learning Technologies**

FOCUS AREA **AD**
TYPE Tutorial
LEVEL Beginner

DAY Wednesday, February 05, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Francis Alexander (Brookhaven National Laboratory) falexander@bnl.gov
Brian Van Essen (Lawrence Livermore National Laboratory) vanessen1@llnl.gov
Peter Nugent (Lawrence Berkeley National Laboratory) penugent@lbl.gov
Christine Ahrens (Los Alamos National Laboratory) cahrens@lanl.gov
Sudip Seal (Oak Ridge National Laboratory) sealsk@ornl.gov
Ian Foster (Argonne National Laboratory) foster@anl.gov
Siva Rajamanickam (Sandia National Laboratories) srajama@sandia.gov
Sutanay Choudhury (Pacific Northwest National Laboratory) Sutanay.Choudhury@pnnl.gov

ABSTRACT
We will provide an overview of the ExaLearn Codesign Project as well as a detailed description of the software and algorithms for applications in the development of surrogate models and systems control, inverse solvers and design. The goal of this tutorial is to provide potential adoptors of the ExaLearn software with a path to getting started using it for their ML problems. Issues of scalability and performance will be emphasized.

TITLE **Exascale I/O with ADIOS: Online Processing, in Situ Visualization and High-Performance Storage I/O**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Scott Klasky (Oak Ridge National Laboratory) klasky@ornl.gov
Norbert Podhorszki (Oak Ridge National Laboratory) pnorbert@ornl.gov
Chuck Atkins (Kitware Inc) chuck.atkins@kitware.com
Philip Davis (Rutgers University) philip.e.davis@rutgers.edu
Greg Eisenhauer (Georgia Institute of Technology) eisen@cc.gatech.edu
William Godoy (Oak Ridge National Laboratory) godoywf@ornl.gov
Junmin Gu (Lawrence Berkeley National Laboratory) jgu@lbl.gov
Manish Parashar (Rutgers University) parashar@rutgers.edu
Matthew Wolf (Oak Ridge National Laboratory) wolfmd@ornl.gov
K. John Wu (Lawrence Berkeley National Laboratory) kwu@lbl.gov

ABSTRACT
As concurrency and complexity continue to increase on high-end machines, I/O performance is rapidly becoming a fundamental challenge to achieving exascale computing. To address this challenge, wider adoption of higher-level I/O abstractions will be critically important. Modern I/O libraries provide data models, portable APIs, storage abstractions, in situ infrastructure, and self-describing data containers. These capabilities enable reproducible science, allow for more effective data management throughout the data life cycle, and facilitate parallel I/O with high performance and scalability. The ADIOS framework provides the abstractions needed to address many exascale I/O challenges. In this tutorial, participants will learn about the ADIOS concept and APIs and participants can run a pipeline of simulation, analysis and visualization, using both storage I/O and in situ data staging. Participants will also learn how to use state-of-the art compression techniques with ADIOS. Finally, we will discuss how to use ADIOS on the LCFs for the best storage performance, in situ data analysis and code coupling.

DESCRIPTION

- Part 1. Presentation: Overview of ADIOS (45 min)
  ADIOS framework, I/O abstraction, streaming data format, and APIs
- Part 2. Hands-on (90 min)
  ADIOS code examples for publishing and subscribing to ECP data
- Part 3. Tuning on Summit, Theta and Cori (45 min)
  ADIOS performance tuning / ECP goals, status

This tutorial aims to teach the ADIOS concept of data publisher/subscriber interface, and the way we envision writing scientific pipelines. We provide a VirtualBox VM for participants to run a simulation-analysis-plotting-visualization pipeline. Participants will learn how to use an I/O abstraction to decouple the application data movement/storage needs from the underlying infrastructure, which will enable them to write simulations that consist of multiple applications coupled together in a flexible way, to perform in situ analysis and data reduction pipelines to deal with overwhelming amount of data that cannot be stored permanently, to use intermittent and interactive in situ visualization, and to write/read files with high I/O

performance. We will also discuss how to run these things on the LCFs and how to tune ADIOS for the best performance by showing the settings for WDMApp, E3SM-MMF and EXAALT applications on the LCFs.

TITLE **Flux: Using Next-Generation Resource Management and Scheduling Infrastructure for Exascale Workflows**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 8:30 AM - 4:30 PM central
DURATION 330 min

AUTHORS
Dong Ahn (Lawrence Livermore National Laboratory) ahn1@llnl.gov
Stephen Herbein (Lawrence Livermore National Laboratory) herbein1@llnl.gov
Daniel Milroy (Lawrence Livermore National Laboratory) milroy1@llnl.gov
Tapasya Patki (Lawrence Livermore National Laboratory) patki1@llnl.gov
Francesco Di Natale (Lawrence Livermore National Laboratory) dinatale3@llnl.gov
Anthony Agelastos (Sandia National Laboratories) amagela@sandia.gov

ABSTRACT
Present day scientific workflows require far more complex interplays among physics simulations, analysis and visualization tools, data stores, etc. This has outpaced the capabilities of today's HPC resource management and scheduling (RMS) systems. Various scalability and programming effort challenges have already begun to appear, and several ad-hoc approaches to workflow management have come into existence. Exascale platforms will present significantly greater challenges and such ad-hoc solutions for management of scientific applications will not be sufficient. In this tutorial, we will introduce Flux, a next-generation resource management framework that supports diverse upcoming scientific workflows and complex resources such as power and IO. With the help of practical demonstrations, attendees will learn about easy-to-use, portable solutions for scientific workflows with the Flux API as well as about hierarchical, user-level scheduling techniques that can be used for high-throughput computing and ensemble management.

DESCRIPTION
Relevance of Topic:
The next-generation of supercomputing requires support for diverse resource types and complex workloads. Workload types are changing from traditional long-running jobs, to a mix containing jobs with in-situ analysis and visualization, co-scheduled components, high-throughput computing and data-driven codes. Despite this changing landscape, HPC resource managers and job schedulers continue to use restrictive and old approaches to solve emerging problems, creating a technology gap.

Flux is an open-source, next-generation resource management framework that addresses this technology gap and expands the scheduler's view beyond the single dimension of nodes. It provides support for management of complex workflows as well as diverse resources such as power, network, IO with ease. The Flux infrastructure is designed to be fully hierarchical in order to expose high levels of scheduler parallelism while allowing for workload customization and enforcement of multi-level constraints through its flexible graph-based resource model. Most current HPC scheduling frameworks, such as SLURM or PBSPro, are not designed to support complex workflows and rely on a centralized daemon, resulting in scalability bottlenecks and limited throughput. This is an important thrust area for exascale, making this topic highly relevant for ECP.

Content Level:
This tutorial introduces users to Flux and will equip the audience with mechanisms to easily manage advanced scientific workflows using the Flux API as opposed to the current ad-hoc techniques. The tutorial will target general HPC system users/practitioners, scientific workflow developers, and scheduling researchers. We will show use cases based on ECP Pilot2 and ensemble (UQ/V&V) workflows, where we will illustrate how Flux can scale up to over a million simultaneous job requests with ease. While basic knowledge in the HPC area is useful, there are no prerequisites to understand the tutorial. Our previous Flux tutorial at the ECP 2019 meeting was highly successful and well-received.

With the help of compelling use cases, practical demonstrations, and hands-on exercises, attendees will be able to:

- Understand the limitations of current job schedulers for scientific workflows, and learn how Flux can significantly reduce scripting effort and scheduling complexity,
- Understand and apply the Flux API to manage complex scientific workflows in a portable manner with a hands-on demo (Docker and/or AWS)
- Understand user-level hierarchical scheduling and apply such user-level schedulers to manage their ensembles and science workflows
- Understand how Flux can be installed and set up with the help of detailed examples
- Understand how Flux can be integrated to manage their specific workflow (support for debugging and testing will be provided),
- Understand the graph-based data model and how to write customized scheduling plugins
- Develop awareness for future directions and roadmap for Flux

Hands-on exercises:
We will make the content available through Docker images and/or AWS instances as well as detailed documentation and presentations, which attendees can utilize based on their convenience, both during and after the session. We will also reserve an hour to help users integrate Flux in their applications and workflows and provide debugging support.

Previous material: https://github.com/flux-framework/Tutorials

TITLE **Frontier Tutorial**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 8:30 AM - 12:00 PM central
DURATION 180 min

AUTHORS

Justin Whitt (Oak Ridge National Laboratory) whitt@ornl.gov
Bronson Messer (Oak Ridge National Laboratory) bronson@ornl.gov
Judith Hill (Oak Ridge National Laboratory) hilljc@ornl.gov
Noah Reddell (Cray) nreddell@cray.com
Nicholas Malaya (Advanced Micro Devices, Inc.) Nicholas.Malaya@amd.com
Jakub Kurzak (Advanced Micro Devices, Inc.) Jakub.Kurzak@amd.com
Rene van Oostrum (Advanced Micro Devices, Inc.) rene.vanoostrum@amd.com
Damon McDougall (Advanced Micro Devices, Inc.) damon.mcdougall@amd.com

ABSTRACT

The Oak Ridge Leadership Computing Facility (OLCF) will host a Frontier Application Readiness Tutorial where the OLCF and Cray's Center of Excellence (CoE), including AMD, will deliver information about the Frontier system to the ECP application and software teams. This tutorial will largely be a condensed repeat of material presented during the OLCF CAAR kick-off workshop held in October in Knoxville, TN.

Topics include:

- Frontier architecture: Node architecture, Slingshot interconnect, and Lustre/flash filesystem
- AMD GPU basics: Hardware, execution and memory hierarchies
- Frontier Programming Environment: HIP, OpenMP Offload, and abstraction models

Some updates since the October workshop, particularly on the Frontier programming environment, may be presented. The intent of this tutorial is to make material available to those who were unable to attend the October workshop due to either scheduling constraints or in-person space limitations.

DESCRIPTION

At the OLCF hosted an Application readiness kick-off workshop in October, due to space constraints at the workshop venue the OLCF was limited to only allowing each ECP AD team to send 3 people. The entirety of the ECP ST focus area was limited to only 10 members in total. The goal of this workshop is to repeat much of the material presented at that workshop to the broader ECP community to ensure that the information is distributed as widely as is possible.

Topics are expected to be a subset of what was presented at the October workshop (posted under the Agenda tab at https://www.olcf.ornl.gov/frontier-application-readiness-kick-off-workshop/). Due to time limitations, there will not be an opportunity for hands-on session in this session.

Goals:
The goal of this session is to communicate information about the Frontier architecture and programming environment to the broader ECP community.

Target Audience:
ECP Application and Software Technology teams that were unable to attend the OLCF Frontier Application Readiness workshop in October.

TITLE **Getting It Right with Open MPI - Best Practices for Deployment and Tuning of Open MPI**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Howard Pritchard (Los Alamos National Laboratory) howardp@lanl.gov
George Bosilca (The University of Tennessee, Knoxville) bosilca@icl.utk.edu
Thomas Naughton (Oak Ridge National Laboratory) naughtont@ornl.gov
David Bernholdt (Oak Ridge National Laboratory) bernholdtde@ornl.gov

ABSTRACT
Open MPI is an open source implementation of the MPI-3 standard that is widely available at ECP facilities. Open MPI is based on a Modular Component Architecture (MCA) design which allows for a great deal of flexibility in its deployment and usage. The choice of Open MPI installation and usage options can have a significant impact on the available MPI functionality, the MPI performance realized by applications, and its ability to be used in containerized environments and more complex workflows. In addition, Open MPI is becoming increasingly dependent on networking middleware packages like OpenUCX and OFI libfabric. The options used to build these packages can also have a significant impact on the performance of Open MPI. This tutorial will provide guidance on how best to deploy Open MPI and its dependencies, runtime tuning options, use of Open MPI with various containers technologies, and special considerations for applications using GPUs.

DESCRIPTION
It can be challenging for system administrators and application developers to deploy optimally configured version of Open MPI for their site or specific workflows due to the wide range of options available for building and deploying Open MPI as well as several important dependencies including OpenUCX, OFI libabric, PMIx. This has led to the notion that tuning Open MPI for a given system is somewhat of a black art. In this tutorial, we intend to dispel this notion by providing guidance on the deployment of Open MPI, special considerations when using Spack-based deployments, best practices for Open MPI use with containers, and special considerations for GPU accelerated applications. Additional topics that will be covered include interoperability of Open MPI with various debuggers and performance analysis tools and interoperability with third-party job launchers.

The target audience is primarily those who build and install Open MPI in different operational environments, although power users of Open MPI may find the content useful as well.

TITLE **Hands-on with Progress Tracking Cards: A Lightweight Method for Improving Your Software Practices**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL Beginner

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 4:00 PM central
DURATION 90 min

AUTHORS
Elsa Gonsiorowski (Lawrence Livermore National Laboratory) gonsiorowski1@llnl.gov
Reed Milewicz (Sandia National Laboratories) rmilewi@sandia.gov
Elaine Raybourn (Sandia National Laboratories) emraybo@sandia.gov

ABSTRACT
This hands-on session is an opportunity for project contributors to level up their team's software engineering practices. Led by the IDEAS-ECP team, this session will provide a brief overview of the Productivity and Sustainability Improvement Planning (PSIP) process. Attendees will learn about the value of user stories (as a flexible means for communicating about software requirements) and how to organize a PSIP cycle within their own team. For the hands-on portion, attendees may select a Progress Tracking Card (PTC) from the BSSw catalog and customize it to fit their current team and project, or define their own PTC.

PSIP is a software process improvement methodology that can help teams increase software quality while decreasing the effort, time, and cost to develop, deploy, maintain, and extend the software over its intended lifetime. The PSIP workflow is lightweight, intended to fit in with a project,Äôs standard planning and development process. A core tool of the PSIP process is the PTC – brief document containing the goal of the planning activity and a step-by-step list of qualitative descriptions and values that help to track progress. The purpose of the PTC is to help a team set and achieve improvement goals. The PTC is not a tool for external assessment or comparison with other projects.

DESCRIPTION
This session will provide an overview of the Productivity and Sustainability Improvement Planning (PSIP) process: a lightweight, iterative workflow where teams identify their most urgent software development and sustainability bottlenecks and track progress on work to overcome them. It will include a hands-on session where participants are given the opportunity to reflect on existing practices in their current projects and create a plan for improvement. These plans will come in the form of Progress Tracking Cards (PTCs) which will be adapted from the existing BSSw catalog.

The objectives of the PSIP process are to capture and convey the practices, processes, policies, and tools of a given software project. The PSIP workflow is intended to be lightweight and fit within a project's planning and development process. It is not meant to be an assessment or evaluation tool. Instead PSIP captures the tacit, more subjective aspects of team collaboration, workflow planning, and progress tracking. Additionally, in the potential absence of planning and development processes, and as scientific software teams scale to larger, more diverse, aggregate teams, unforeseen disruptions or inefficiencies can often impede productivity and innovation. PSIPs are designed to bootstrap aggregate team capabilities into best practices, introduce the application of appropriate resources, and encourage teams to adopt a culture of process improvement.

At the core of a PSIP cycle is a PTC. This brief document contains the target, or goal of the planning activity, title of the topic of improvement, and a step-by-step list of activities or outcomes that incrementally lead to improvements in team effectiveness and efficiency. Teams may select PTCs from the PTC catalog, define their own PTC, or modify PTCs found in the catalog. The purpose of the PTC is to help teams set and achieve improvement goals. The PTC is not a tool for external assessment or comparison with other projects. In fact, since PTCs are custom-designed for each project, comparisons are typically not possible.

Goals: The outcome of this session will concrete steps that teams can take to improve their development practices, in the form of a PTC. Participants will learn some of the skills needed to approach their teams and precipitate process improvement.

Target audience: Process improvement is valuable for all teams, but the lightweight nature of PSIP and PTCs is ideal for scientific software developers. Scientific software teams are typically focused on obtaining scientific results from the software they write. Funding is for generating results, not software. In this session, we will target individuals and teams that have little or modest formal software engineering training.

TITLE **In Situ Visualization and Analysis with Ascent**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Hank Childs (University of Oregon) hank@uoregon.edu
Matt Larsen (Lawrence Livermore National Laboratory) larsen30@llnl.gov
Cyrus Harrison (Lawrence Livermore National Laboratory) cyrush@llnl.gov
Ken Moreland (Sandia National Laboratories) kmorel@sandia.gov
David Rogers (Los Alamos National Laboratory) dhr@lanl.gov
Steve Petruzza (University of Utah) spetruzza@sci.utah.edu
Ayan Biswas (Los Alamos National Laboratory) ayan@lanl.gov
Soumya Dutta (Los Alamos National Laboratory) sdutta@lanl.gov

ABSTRACT
In situ visualization and analysis is an important capability for addressing slow I/O on modern supercomputers. With this 3-hour tutorial, we will spend the majority of our time (two hours) going into detail on Ascent, an in situ visualization and analysis library being developed by ECP ALPINE. Ascent is from the makers of ParaView Catalyst and VisIt LibSim, and it will soon be able to directly connect with both of those products. The tutorial will be practical in nature: how to integrate Ascent into a simulation code, Ascent's data model, build and linking issues, and capabilities. The remaining hour will be spent highlighting other visualization efforts in ECP, such as in situ-specific visualization algorithms, VTK-m, and CINEMA.

DESCRIPTION
This tutorial is on in situ visualization and analysis. We believe the tutorial will primarily be relevant to ECP Application Teams. We also believe the tutorial will be relevant for ST collaborators, as the software we describe is a delivery vehicle for other ST projects, and this tutorial will allow to familiarize themselves with the delivery vehicle.

The tutorial contains two thrusts:

- The Ascent in situ library for visualization and analysis (110 minutes)
- Capabilities within Ascent that come from other ECP Visualization and Analysis groups (60 minutes)

== Why Ascent? ==

The ECP ALPINE project plans for Ascent to be its primary delivery vehicle for ECP Applications. The infrastructure is designed for leading-edge supercomputers and has support for both distributed-memory and shared-memory parallelism. It can take advantage of the computing power on both conventional CPU architectures and on many-core architectures such as NVIDIA GPUs or the Intel Xeon Phi. Further, it has a flexible design that supports for integration of new visualization and analysis routines and libraries.

Ascent is being developed by the makers of ParaView Catalyst and VisIt LibSim. It has advantages over those libraries via its flyweight nature. Further, in the next year, Ascent will be extended to directly connect to Catalyst and LibSim, which will enable increases in capability.

== What Other Visualization and Analysis Capabilities Will Be Described? ==

We will discuss three areas:

1) CINEMA, which is a system for post hoc exploration of images that were extracted in situ. David Rogers of Los Alamos will present material on this.
2) VTK-m, which is a key library for Ascent for supporting many-core architectures. Ken Moreland will present this material.
3) In situ appropriate algorithms. ALPINE is developing algorithms with in situ processing in mind, specifically involving topology, moments, Lagrangian flow, and sampling. We will give several minute overviews of each.

All three of these areas are available in Ascent, as Ascent can produce CINEMA databases, Ascent exposes VTK-m, and Ascent contains the in situ appropriate algorithms.

== What Is the Plan for the Tutorial? ==

10 minutes: brief introduction
15 minutes: install software
20 minutes: data model description
20 minutes: how to integrate Ascent with simulations codes
55 minutes: capabilities in Ascent, including 4 in depth examples
20 minutes: CINEMA overview 20 minutes: in situ algorithms from ALPINE (4 x 5 minutes)
20 minutes: VTK-m

== What Will the Outcome Be for Attendees? ==

The tutorial will be hands on for the Ascent portion. As such, we believe attendees will develop enough experience with Ascent that they can try integrations on their own after attendance. Further, the tutorial's final hour will expose them to other activities within the ECP visualization and analysis portfolio, which they could then use via Ascent.

TITLE **Integrating PaRSEC-Enabled Libraries in Scientific Applications**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
George Bosilca (Innovative Computing Laboratory, The University of Tennessee, Knoxville)
bosilca@icl.utk.edu
Thomas Herault (Innovative Computing Laboratory, The University of Tennessee, Knoxville)
herault@icl.utk.edu
Aurélien Bouteiller (Innovative Computing Laboratory, The University of Tennessee, Knoxville)
bouteill@icl.utk.edu

ABSTRACT
PaRSEC is a task runtime system for distributed, hybrid, large scale supercomputers that supports the execution of task systems represented by a variety of APIs and DSLs. This tutorial will present PaRSEC by focusing on the libraries and applications that use it. The tutorial includes two parts: first, we will present examples of libraries and applications that use PaRSEC, how to use these libraries, and how to take advantage of the task system to leverage GPU acceleration on ECP machines, in particular summit; then, we will give an overview of the the PaRSEC development environment: how to get detailed information on the efficiency of the task systems execution, how to visualize the task systems and debug it.

DESCRIPTION

Topic
Using the PaRSEC Tasks Runtime System to leverage accelerators in parallel environments.

Goal
Present the PaRSEC Tasks runtime system, and the mathematical libraries that are built on top of it; demonstrate its use and how to integrate within applications.

The tutorial will include 2 parts:

First, we will introduce mathematical libraries and applications that use PaRSEC: how they were programmed, the performance they reach on ECP platforms (in particular summit), and their interface. This part will present both results and working tools in traditional dense linear algebra operations (the dplasma library), sparse and low-rank representation of data for factorizations, and irregular dense linear algebra operations that are used in ECP applications.

Then, in the second part, we will provide some technical details on how to configure the runtime system in order to extract the best performance on ECP machines. This part will in particular present the internal instrumentation available to any PaRSEC-enabled application, to get quantitative and visual feedback on the performance of the task systems; we will also present the debugging tools and visualization tools available to the developers to understand the behavior of their programs at scale.

TITLE **Introduction to the Capabilities and Use of the SUNDIALS Suite of Nonlinear and Differential/Algebraic Equation Solvers**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Carol S. Woodward (Lawrence Livermore National Laboratory) woodward6@llnl.gov
Daniel R. Reynolds (Southern Methodist University) reynolds@smu.edu
David J. Gardner (Lawrence Livermore National Laboratory) gardner48@llnl.gov
Cody J. Balos (Lawrence Livermore National Laboratory) balos1@llnl.gov

ABSTRACT
SUNDIALS is a SUite of Nonlinear and DIfferential/ALgebraic equation Solvers consisting of six packages which provide robust time integrators and nonlinear solvers. The suite is designed to be easily incorporated into existing simulation codes and to require minimal information from the user. The modular implementation allows users to easily supply their own data structures underneath SUNDIALS packages and to easily incorporate user-supplied linear solvers and preconditioners. SUNDIALS consists of the following six packages: CVODE, solves initial value problems for ordinary differential equation (ODE) systems with linear multi-step methods; CVODES, solves ODE systems and includes sensitivity analysis capabilities (forward and adjoint); ARKODE, solves initial value ODE problems with explicit, implicit, IMEX, and multirate additive Runge-Kutta methods; IDA, solves initial value problems for differential-algebraic equation (DAE) systems with BDF methods; IDAS, solves DAE systems and includes sensitivity analysis capabilities (forward and adjoint); KINSOL, solves nonlinear algebraic systems.

This tutorial will include an overview of SUNDIALS capabilities and target uses followed by a discussion of user interfaces for the packages as well as for the solvers and data structures that can be applied underneath the SUNDIALS packages. Lastly, we will discuss build and installation procedures for SUNDIALS and how to configure on various architectures.

TITLE **Managing HPC Software Complexity with Spack**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 06, 2020
TIME 8:30 AM - 4:30 PM central
DURATION 330 min

AUTHORS
Todd Gamblin (Lawrence Livermore National Laboratory) tgamblin@llnl.gov
Gregory Becker (Lawrence Livermore National Laboratory) becker33@llnl.gov

ABSTRACT
HPC software is becoming increasingly complex. The largest applications require over 100 dependency libraries, and they combine interpreted languages like Python with C, C++, and Fortran libraries. To achieve good performance, developers tune for multiple compilers, build options, and implementations of dependency libraries like MPI, BLAS, and LAPACK. The space of possible build configurations is combinatorial, and developers waste countless hours rebuilding software instead of producing new scientific results.

This tutorial focuses on Spack, an open-source tool for HPC package management. Spack uses concise package recipes written in Python to automate builds of applications with hundreds of dependencies. Spack can build with arbitrary combinations of compilers, MPI versions, and dependency libraries. Spack allows users to build and install over 3,300 community-maintained packages, developers can quickly build on others' work, and HPC center staff can deploy many versions of software for thousands of users. This tutorial provides a thorough introduction to Spack's capabilities: basic software installation, creating new packages, using Spack for HPC software development, advanced multi-user deployment for facilities, environments, continuous integration, and container use cases. We provide detailed use cases from our experiences in production HPC environments. Most sessions involve hands-on demonstrations, so attendees should bring a laptop computer.

DESCRIPTION
This tutorial will introduce attendees to Spack, an open-source HPC package management tool. Attendees will learn how to automatically build and install over 3,300 packages, develop their own packages, and contribute packages to Spack for others to use. We will have a session specifically for developers of large HPC codes, showing how Spack can be used to test and manage releases with many dependencies. Finally, attendees will learn to deploy software for users at HPC facilities, including best practices from LLNL, EPFL, ANL, and NERSC. We focus on Spack, but we also compare it with other tools (like EasyBuild), so that users can pick the tool that best fits their workflow.

Spack has been chosen as the software deployment tool for the US Exascale Computing Project, and this tutorial will be highly relevant for ECP attendees, many of whom will need to use Spack on their projects. Software complexity issues have broad appeal for ECP attendees. HPC applications can require over 100 dependency libraries. They must run on cutting edge machines with different compilers, languages, and BLAS, LAPACK, or MPI implementations. HPC developers struggle with obscure errors when building by hand on bleeding-edge machines. Manual builds are not consistent or reproducible, and time spent building and rebuilding software takes time away from scientific goals.

Spack is rapidly gaining traction in the broader HPC community. It has a very active open source community, with over 250 contributors. 30-60 pull requests are merged to the Spack GitHub project each week. Spack is currently used to manage clusters at OLCF, ANL, NERSC, LLNL, EPFL, AWE, LRZ, Genentech, and many other production HPC centers.
Spack simplifies combinatorial software builds. It enables users to create reproducible, reusable software recipes, and its powerful specification semantics allow users to customize and optimize builds for their own use. Spack saves users, developers, and administrators many hours of time, and tools like it are becoming increasingly necessary to build software HPC systems. Moreover, Spack is an open source project. It has over 400 contributors from over 60 institutions, including many of the largest HPC centers in the world. This tutorial will bring attendees into the broader community, allowing them to contribute to the Spack project and share their work with other HPC users across the globe.
This tutorial targets anyone who builds software on HPC machines: Users build and run packages written by others. Developers manage complex dependencies for their own software releases and testing across a wide array of machines. Facility staff and administrators work at HPC centers and support many users by building, installing, and maintaining system software packages.

General Description: Background: We begin with background material (II) to establish key concepts for building software on supercomputers. We cover fundamentals of libraries, linking, and loading, and how dynamic loaders find dependencies (LD LIBRARY PATH and RPATH). We briefly cover ABI compatibility, and we describe major build systems for compiled languages (make, CMake, autotools) and for interpreted languages (Python, R) at a high level. This helps attendees understand what Spack can automate.

Basic Usage: The first hands-on session (III), introduces how to query available packages, how to install and uninstall, and how to query what has been installed. Attendees will understand how to use Spack to install already-packaged software.

Packaging: Sessions IV-V focus on adding custom packages to Spack. We give an overview of Spack's installation model, and we teach attendees to write their own build recipes in Spack's embedded DSL. We start with automating the builds of simple packages with no dependencies, which Spack can typically generate automatically, then add dependencies, patching, and variants. Participants should be able to create packages that can be built with arbitrary compiler/MPI combinations.
Deploying Software at HPC Facilities: Session VI will show users workflows and best practices for Spack usage at LLNL, from experiences at EPFL, NERSC, and ANL. We provide a hands-on tutorial on customizing the module files that Spack can generate. Session VII will show users how to cache and share builds using Spack's binary packaging features.
Spack for developers: Spack provides flexible configuration in order to support the needs of HPC developers. Sections VIII-IX focus on configuration options and abstractions that allow developers to use external libraries and handle the more complex aspects of HPC builds (linking MPI, BLAS, LAPACK, etc.) Session X focuses on combinatorial test matrices, which Spack simplifies, and Section XI describes Spack's new support for creating environments that can be replicated across platforms.

New features: We will add material to our past tutorials covering new features such as environments, continuous integration, and container use cases.

Contributing: Our last hands-on session (XII) will allow attendees to package their own software using Spack, with help from experts. Attendees build a new Spack package or improve an existing one. Spack developers will provide 1x1 help and suggestions to attendees. We will then show users how they can submit pull requests on GitHub to include their packages in the Spack mainline.

TITLE **Managing Power Efficiency of HPC Applications with Variorum and GEOPM**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Tapasya Patki (Lawrence Livermore National Laboratory) patki1@llnl.gov
Aniruddha Marathe (Lawrence Livermore National Laboratory) marathe1@llnl.gov
Stephanie Brink (Lawrence Livermore National Laboratory) brink2@llnl.gov
Barry Rountree (Lawrence Livermore National Laboratory) rountree4@llnl.gov

ABSTRACT
As we push the limits of supercomputing toward exascale, resources such as power and energy are becoming expensive and limited. HPC PowerStack is an end-to-end holistic approach to power management that includes node-level, job-level and cluster-level policy integration. The focus of this tutorial will introduce attendees to two key components of the PowerStack. We will discuss the challenges in power management and the supporting hardware features that can be leveraged by the HPC software stack. We will introduce Variorum, a user-space, vendor-neutral library for power management and control that supports Intel and IBM platforms through a common API. The tutorial will then discuss GEOPM as a cross-community, cross-platform solution for job-level power management.

DESCRIPTION

Overview

As we push the limits of supercomputing toward exascale, resources such as power and energy are becoming expensive and limited. HPC PowerStack is an end-to-end holistic approach to power management that includes node-level, job-level and cluster-level policy integration. The focus of this tutorial will introduce attendees to two key components of the PowerStack. We will discuss the challenges in power management and the supporting hardware features that can be leveraged by the HPC software stack. We will introduce Variorum, a user-space, vendor-neutral library for power management and control that supports Intel and IBM platforms. The tutorial will then discuss GEOPM as a cross-community, cross-platform solution. Our GEOPM tutorial at the ECP 2019 meeting was well received.

Goals and Content

The tutorial will target general HPC system users, application programmers, and power researchers. No prior experience is necessary. It aims to provide theoretical and practical demonstration on how Variorum can be used for power measurement and control, and how GEOPM can be used for safe execution and performance optimization of applications running in a power-constrained environment. The following challenges will be addressed during the session:

1. Vendor-neutral APIs for Power Management: Exascale architectures are diverse and include vendors such as Intel, IBM, AMD, ARM, and NVIDIA. Each architecture introduces a different power measurement and control mechanism, making it challenging for application developers and users to manage power. Variorum introduces a simple and accessible vendor-neutral API that masks the complicated details of the underlying architecture from the users. We will discuss two to three architectures in detail (Intel, IBM, NVIDIA, for instance) and present a common API for managing power with Variorum.

2. Systems with varying efficiency demands: The GEOPM framework is capable of running asynchronously on compute nodes of an HPC system, and leveraging hardware capabilities to dynamically tune applications for multiple efficiency metrics. It uses learning and control system techniques to discover runtime patterns in the application and tune hardware knobs in the underlying hardware platform to exploit those patterns.

3. Non-deterministic factors contributing to imbalance during application execution: Factors like the degree of heterogeneity across different system architectures, manufacturing variation in hardware components, unpredictability of the physical layout of the nodes selected during job allocation, contribute to variation in performance on HPC systems. The automated and decentralized operation of GEOPM runtime helps mitigate this by enabling asynchronous steering of power among compute nodes.

4. Ease of extension: GEOPM supports extensible tuning strategies through its scalable plugin-based infrastructure. This enables system programmers to design multiple energy/power management algorithms in the form of plugins. The framework also provides a streamlined path for deploying new ideas and algorithms on different platforms.

At the end of this tutorial, attendees will:

1. be familiarized with the state-of-the-art solutions for power management provided by current hardware platforms and understand how to use Variorum API
2. clearly understand the multiple factors that contribute to performance variability in HPC applications
3. be able to use Variorum and GEOPM to monitor and control their application efficiency

We will have demos of each component and discuss integration and support for ECP applications.

TITLE **Memory Movement Orchestration and Topology-Aware Placement with AML**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 8:30 AM - 9:30 AM central
DURATION 60 min

AUTHORS
Swann Perarnau (Argonne National Laboratory) swann@anl.gov
Nicolas Denoyelle (Argonne National Laboratory) ndenoyelle@anl.gov
Florence Monna (Argonne National Laboratory) fmonna@anl.gov
Brice Videau (Argonne National Laboratory) bvideau@anl.gov
Kamil Iskra (Argonne National Laboratory) iskra@mcs.anl.gov
Pete Beckman (Argonne National Laboratory) beckman@anl.gov

ABSTRACT
Exascale systems are expected to exhibit complex memory hierarchies and topologies, including multiple types of memory (HBM, DDR, NVM, Accelerator-side memory) and multiple accelerator devices per node. ECP Projects might be facing a challenge taking full advantage of these platforms, in particular for workloads requiring the use of multiple devices per rank or needing to orchestrate data movement across multiple memory types.

The Argo project is developing AML, a memory management library, to help application developers easily and portably integrate data movement orchestration facilities in their applications. This tutorial session will present the library and how it can be used to improve application performance on future platforms.

DESCRIPTION
The Argo project is developing AML, a memory management library, to help application developers easily and portably integrate data movement orchestration facilities in their applications. Intra-node, in-memory data movement will be of critical importance for ECP application workloads that require the use of multiple accelerator devices per rank or that need to move data between multiple types of memory during runtime.

This tutorial session aims to provide ECP application developers with a full picture of the challenges of memory management on future exascale platforms.

The tutorial will cover:

- a presentation of the expected memory topologies on future platforms,
- an overview of the performance trade-offs between the different memory types as well as challenges when using multiple devices from a single rank,
- key application characteristics to consider before designing complex memory management policies
- common abstractions provided by AML to reason about memory management and how to combine them for efficient data orchestration.

The tutorial will in particular include a discussion of:

- data layout in memory and its impact on performance
- allocating and moving data between memories
- transforming and tiling data across a topology

TITLE **Modern CMake Tutorial**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL N/A

DAY Thursday, February 6, 2020
TIME 1:30 PM - 4:30 PM central
DURATION 150 min

AUTHORS
Zack Galbreath (Kitware Inc.) zack.galbreath@kitware.com

ABSTRACT
Best practices for using CMake to build, test, and share your software.

This tutorial will serve as an introduction to the CMake cross-platform build system. We will cover the modern, target-centric approach to designing builds with CMake. Other topics will include automated testing, continuous integration, and packaging.

TITLE **OpenMP 4.5 and 5.0 Tutorial (Offload)**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Expert

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Oscar Hernandez (Oak Ridge National Laboratory) oscar@ornl.gov
Tom Scogland (Lawrence Livermore National Laboratory) scogland1@llnl.gov
Johannes Doerfert (Argonne National Laboratory) jdoerfert@anl.gov
Colleen Bertoni (Argonne National Laboratory) bertoni@anl.gov
Jae Hyuk Kwack (Argonne National Laboratory) jkwack@anl.gov

ABSTRACT
OpenMP is the defacto directive-based parallel programming API to program applications on shared memory and accelerator architectures. In this tutorial we will focus on how we can use directives to program accelerators using OpenMP 4.5 and the latest features in OpenMP 5.0.

We propose a tutorial that describes recent and new features of OpenMP, with a focus on those features that have proven important for ECP. We will cover emerging programming styles and best practices for early adopters that they can use to efficiently program accelerators, orchestrate work across CPUs, GPUs and the network, and take advantage of the memory hierarchy. We will also discuss how OpenMP 4.5 offload can be used as a backend to other frameworks like Raja and Kokkos that target accelerators which target C++ applications.

DESCRIPTION
While most HPC developers have MPI experience, many are not familiar with the latest OpenMP features and how to use them in their codes. Modern OpenMP can be used in creative ways to map parallelism to current and emerging parallel architectures. Yet it is much broader and more complex than OpenMP was even just three years ago, and a significant amount of training is needed if programmers are to be able to exploit it fully. We propose a tutorial that describes recent and new features of OpenMP, with a focus on those features that have proven important for ECP. We will cover emerging programming styles and best practices for early adopters that they can use to efficiently program accelerators, orchestrate work across CPUs, GPUs and the network, and take advantage of the memory hierarchy.

The OpenMP 5.0 specification was released at SC'18 with exciting new functionality. Thus it is important not only for developers to be aware of the current standard, and what is available today, but also what is coming next and what will be available in the exascale time frame. Finally, we also want to use this opportunity to discuss with the application teams areas where they can influence the design choices and the open challenges that we have not solved yet but that we will address in the future.

The tutorial is expected to cover the following topics:

- Overview of what is available today in OpenMP 4.5 and will be available in OpenMP 5.0
- An overview of the accelerator programming model

- Examples on how to map data structures
- How to exploit the parallelism in the target regions
- Latest features in OpenMP tasking
- How to use OpenMP target and tasking constructs to manage and orchestrate work and communication between the CPUs and accelerators and inter-node communication.
- Discuss some examples of success stories on how applications have used OpenMP to scale on leadership class systems.
- Other uses: Using OpenMP with other frameworks like Raja/Kokkos
- Provide a deeper explanation of OpenMP 5.0 and discuss the latest proposed features in the OpenMP Technical report for OpenMP 5.1

TITLE **Performance Evaluation Using the TAU Performance System**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS

Sameer Shende (University of Oregon) sameer@cs.uoregon.edu

ABSTRACT

The TAU Performance System provides HPC researchers a powerful profiling and tracing toolkit. This tutorial will focus on performance data collection, analysis, and performance optimization of parallel applications. The tutorial will introduce profiling and debugging support in TAU, cover performance evaluation of parallel programs written in Fortran, Python, C++, and C. The tutorial will also cover parallel performance analysis of applications using MPI, OpenMP, OpenACC, OpenCL, CUDA, and AMD ROCm. The tutorial will demonstrate different techniques for program instrumentation including compiler-based instrumentation for LLVM and highlight TAU's support for memory debugging and I/O evaluation. The tutorial will guide the developers through the instrumentation, measurement, and analysis process steps in TAU. Performance data will include MPI timings, I/O and memory, and hardware performance counters from PAPI.

DESCRIPTION

Outline:

1. Introduction to performance evaluation tools.
2. Instrumentation for OpenMP, MPI, AMD ROCm, CUDA, OpenACC, OpenCL, LLVM compilers.
3. Measurement options in TAU.
4. Performance analysis using ParaProf, demo.
5. Performance evaluation in the E4S HPC container environments (Docker, Shifter, Singularity, Charliecloud) with TAU.
6. Demonstration of TAU with AWS and E4S

TITLE **Performance Tuning with the Roofline Model on GPUs and CPUs**

FOCUS AREA **HI**
TYPE Tutorial
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 2:30 PM - 6:00 PM central
DURATION 180 min

AUTHORS
Samuel Williams (Lawrence Berkeley National Laboratory) swwilliams@lbl.gov
Jack Deslippe (Lawrence Berkeley National Laboratory) jrdeslippe@lbl.gov
Charlene Yang (Lawrence Berkeley National Laboratory) CJYang@lbl.gov

ABSTRACT
The HPC community is on a never ending quest for better performance and scalability. Performance models and tools are an integral component in the optimization process as they quantify performance relative to machine capabilities, track progress towards optimality, and identify performance bottlenecks. The Roofline performance model offers an insightful and intuitive method for extracting the key computational characteristics of HPC applications and comparing them against the performance bounds of CPUs and GPUs. Its capability to abstract the complexity of memory hierarchies and identify the most profitable optimization techniques has made Roofline-based analysis increasingly popular in the HPC community. This 180-minute tutorial is centered around five components. First, we will introduce the Roofline model and discuss how changes in data locality and arithmetic intensity visually manifest in the context of the Roofline model. Next, we will show how one can use ERT and NVProf to perform basic Roofline analysis on NVIDIA GPUs. After a short break, we will explore advanced Roofline concepts on GPUs including instruction rooflines and Roofline for AI and reduced precision codes. Moreover, we will demonstrate how to use LIKWID and Intel Advisor to use Roofline to analyze applications running on x86 CPUs. Finally, we will conclude with a series of vignettes of how we have used Roofline to analyze various CPU- and GPU-accelerated applications.

DESCRIPTION
All of DOE's pre-exascale machines (Summit, Sierra, and Perlmutter) as well as its exascale machines (Aurora, Frontier, and El Capitan) will leverage GPUs for performance. Although porting to a GPU will be covered by other tutorials, this tutorial is focused on ensuring programmers can effectively analyze the performance of GPU-accelerated applications and take actions that ensure exascale performance potential is not squandered. Unfortunately, many domain scientists lack the low-level computer architectural knowledge required to effectively understand many vendor performance tools.

The Roofline model provides a visually-intuitive means of analyzing performance across GPUs and CPUs, tracking progress towards optimality, and motivating algorithmic changes required to exploit emerging architectures. Its ability to abstract the complexity of memory hierarchies and identify the most profitable optimization techniques has made Roofline-based analysis increasingly popular among DOE's computer scientists, computational scientists, and applied mathematicians.

This 180-minute tutorial is focused on ensuring computer scientists, computational scientists, and applied mathematicians tasked with porting applications and frameworks to emerging GPU architectures, can effectively analyze performance, quantify the ultimate performance potential, and understand what properties an optimization must embody to deliver better performance. Every attendee should come out of

the tutorial full versed in the Roofline model and methodologies to analyze GPU-accelerated and CPU-based codes using the Roofline model.

The tutorial is comprised of 5 components. The first is an introduction to the Roofline model. This will all attendees are familiar with the theory and can understand how it can be applied to specific architectures. The second component will be a how-to on using NVProf and NSight compute to collect the basic performance metrics for Roofline analysis on NVIDIA GPUs (note, we will use NVIDIA GPUs in this tutorial as they represent all of DOE's current pre-exascale machines, but acknowledge the methodology is easily applicable to those from AMD or Intel). After a short break, we will return and proceed with advanced topics for Roofline on GPUs. This includes instruction throughput Rooflines (for integer-heavy codes), effects of strided global memory access and shared memory bank conflicts, and Rooflines for Tensor Cores and mixed precision.

As all accelerators require a host, who's performance can be a bottleneck, we will spend 30 minutes demonstrating how to use Roofline on CPUs. For simplicity we will demonstrate two methodologies for x86 CPUs (AMD and Intel). The former is a performance-counter based solution using LIKWID, while the latter uses Intel Advisor with Integrated Roofline. Once again, the performance counter methodology can be easily extended to POWER or ARM architectures. Finally, the tutorial will conclude with a series of case studies of how Roofline has been used to aide in the performance optimization process.

TITLE **SLATE and MAGMA Linear Algebra & FFT Libraries Tutorial**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Mark Gates (The University of Tennessee, Knoxville) mgates3@icl.utk.edu
Stan Tomov (The University of Tennessee, Knoxville) tomov@icl.utk.edu

ABSTRACT
This tutorial focuses on the SLATE and MAGMA linear algebra libraries and the FFT-ECP library. These libraries are designed to harness today's supercomputers with multicore CPUs and multiple GPU accelerators. The tutorial covers practical aspects in setting up matrices and calling the libraries in your application. No prior knowledge is required.

SLATE is a modern C++ library developed as part of ECP to replace ScaLAPACK for solving dense linear algebra problems on distributed systems. It supports multicore CPU nodes or hybrid CPU--GPU nodes, using MPI for communication and OpenMP tasks for node-level scheduling. It covers parallel BLAS, solving linear systems (LU, Cholesky, QR, symmetric indefinite), least squares, symmetric eigenvalue and SVD solvers. SLATE includes a ScaLAPACK compatible interface to aide transitioning existing applications.

MAGMA is a C library for accelerating dense and sparse linear algebra on a node using multiple GPUs. It covers a large part of LAPACK's functionality: LU, Cholesky, QR, symmetric indefinite, eigenvalue and singular value solvers. It has a batch component for solving many small problems in parallel, and tensor contractions for high-order finite element methods. The sparse component accelerates many iterative algorithms such as CG, GMRES, and LOBPCG. MAGMA also includes a Fortran interface.

FFT-ECP is an ECP library for multidimensional Fast Fourier Transforms (FFTs) for Exascale platforms. FFTs are in the software stack for almost all ECP applications. The tutorial will cover FFT-ECP,Äôs APIs, approach, performance, and use in applications. FFT-ECP leverages existing FFT capabilities, including 1-D FFTs , FFTMPI, and SWFFT.

DESCRIPTION
This is primarily targeted at application users that need to solve linear algebra and FFT problems. Other libraries that build on linear algebra and FFTs may also benefit.

Goals include:
- compiling & installing libraries
- explaining data format (matrices, etc.)
- covering calling sequence for major operations (linear solves, etc.)
- showing simple but complete examples that users can try

Compared to the 2019 ECP tutorial, this updates the SLATE and MAGMA information, and adds the FFT component.

TITLE **STRUMPACK / SuperLU: Fast Parallel Direct Linear Solvers and Preconditioners**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Thursday, February 6, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
Sherry Li (Lawrence Berkeley National Laboratory) xsli@lbl.gov
Pieter Ghysels (Lawrence Berkeley National Laboratory) pghysels@lbl.gov
Yang Liu (Lawrence Berkeley National Laboratory) liuyangzhuan@lbl.gov

ABSTRACT
Matrix factorizations and the accompanying solution algorithms (e.g., triangular solution associated with the LU factorization) are often the most robust algorithmic choices for linear systems from multi-physics and multi-scale simulations. They are indispensable tools for building various algebraic equation solvers. They can be used as direct solvers, as coarse-grid solvers in multigrid, or as preconditioners for iterative solvers. As we are approaching the exascale computing era, demand for algorithm innovation is increasingly high. It is imperative to develop optimal-complexity scalable algorithms both in flop count and more importantly in data movement, such as, in the form of communication-avoiding formulation, and low-rank and butterfly compressions. On the software and implementation side, it is imperative to exploit multiple levels parallelism presented by the heterogeneous node architectures through well orchestrated use of MPI, OpenMP and GPU programming like CUDA. In this tutorial, we will present our recently developed novel techniques to address the scalability gaps. We will demonstrate their efficacies through three solver libraries: SuperLU, STRUMPACK and ButterflyPACK, with representative use cases from simulations and data analytics. Through hands-on exercises, the participants will learn how to use each solver most effectively for their target problems and parallel machines.

DESCRIPTION
This tutorial provides an overview of distributed-memory direct solvers for large-scale linear systems arising from various scientific and engineering applications. These techniques provide an appealing avenue for increasing capability and efficiency of numerical simulations and data analysis, and the effectiveness of domain scientists for solving complex problems.

Direct solvers are indispensable tools for solving linear systems that are inherently ill-conditioned and/or require multiple right-hand sides. Although traditional direct solvers leveraging LU decompositions require cubic complexities for dense systems, significant progress has been made in recent years to develop optimal-complexity solvers including communication-avoiding and optimization techniques, low-rank and butterfly compressions and their hierarchical constructs, and scalable memory reduction ordering algorithms. Several open-source direct solver packages are being developed at the Scalable Solvers Group at LBNL including:

- SuperLU_DIST -- a supernodal solver for large unsymmetric sparse systems,
- STRUMPACK -- a multifrontal solver enhanced with hierarchical matrix techniques for large sparse and dense systems, and
- ButterflyPACK -- a low-rank and butterfly enhanced hierarchical matrix solver for large dense systems, among a few others. Typical usages of these packages will be illustrated in this tutorial.

Most HPC practitioners rely, directly or indirectly, on the efficient solution of sparse or dense linear systems. Some consider the problem to be solved when it is formulated as a linear system, while others focus on maximizing performance in basic linear algebra building blocks. This tutorial can be useful for both. We mainly anticipate application developers looking for more efficient linear algebra methods. Target users can interface with our solver libraries directly or through other libraries such as hypre, PETSc, Trilinos or SUNDIALS.

Tutorial attendees will gain a basic understanding of state-of-the-art numerical techniques used in fast linear solvers as implemented in the SuperLU, STRUMPACK and ButterflyPACK libraries Familiarity with the available solver options can benefit practitioners from a wide range of application domains. For instance for several classes of dense linear systems, the techniques available in STRUMPACK and ButterflyPACK can achieve linear complexity with respect to the problem size N, while a traditional approach using ScaLAPACK results in O(N^3) complexity. We want to bring recent algorithmic developments and high performance parallel implementations to the attention of the broader ECP application domains.

We will prepare a VirtualBox image with all the libraries and their dependencies preinstalled, also including the demo materials and the tutorial slides.

TITLE **Umpire: Managing Heterogeneous Memory Resources**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Tuesday, February 4, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
David Beckingsale (Lawrence Livermore National Laboratory) beckingsale1@llnl.gov
Marty McFadden (Lawrence Livermore National Laboratory) mcfadden8@llnl.gov

ABSTRACT
Advanced architectures provide a wide range of memory resources that must often be carefully controlled by the user. These resources will have varying capacities, access timing rules, and visibility to different compute resources. Applications must intelligently allocate data in these spaces, and depending on the total amount of memory required, applications may also be forced to move data between different parts of the memory hierarchy. Finally, applications using multiple packages must coordinate effectively to ensure that each package can use the memory resources it needs. Umpire is a library to address these challenges from an application-oriented perspective.

In this tutorial, we will demonstrate how Umpire allows computer scientists and computational scientists to efficiently program the memory hierarchies of current and future high-performance computing architectures, without tying their application to specific hardware or software. We will use examples in C++, C, and Fortran that show how developers can allocate memory using complex algorithms, modify data using abstract operations, and use Umpire's tools to gather insight into their application memory usage patterns.

DESCRIPTION
This tutorial will present a beginner/intermediate look at Umpire, a library developed to allow powerful and portable access to complex memory resources.

The goals of this session are for each attendee to understand the concepts Umpire uses to describe memory, and how to use these concepts in a concrete fashion to allocate, modify, and introspect memory using either the C++, C, or FORTRAN APIs provided by Umpire.

The target audience of this session is a computer scientist or application developer with a basic understanding of how memory is organized in current high-performance computing architectures. A basic understanding of one of the programming languages we support is also expected.

TITLE **UPC++: A PGAS/RPC Library for Asynchronous Exascale Communication in C++**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Beginner

DAY Thursday, February 6, 2020
TIME 1:30 PM - 4:30 PM central
DURATION 150 min

AUTHORS
Amir Kamil (Lawrence Berkeley National Laboratory) akamil@lbl.gov
John Bachan (Lawrence Berkeley National Laboratory) jdbachan@lbl.gov
Scott B. Baden (Lawrence Berkeley National Laboratory (retired)) baden@lbl.gov
Dan Bonachea (Lawrence Berkeley National Laboratory) dobonachea@lbl.gov
Rob Egan (Lawrence Berkeley National Laboratory) regan@lbl.gov
Paul Hargrove (Lawrence Berkeley National Laboratory) PHHargrove@lbl.gov
Steven Hofmeyr (Lawrence Berkeley National Laboratory) shofmeyr@lbl.gov
Mathias Jacquelin (Lawrence Berkeley National Laboratory) MJacquelin@lbl.gov
Katherine Yelick (Lawrence Berkeley National Laboratory) kayelick@lbl.gov

ABSTRACT
UPC++ is a C++ library providing classes and functions that support Partitioned Global Address Space (PGAS) programming. The UPC++ API offers low-overhead one-sided RMA communication and Remote Procedure Calls (RPC), along with futures and promises. These constructs enable the programmer to express dependencies between asynchronous computations and data movement. UPC++ supports the implementation of simple, regular data structures as well as more elaborate distributed data structures where communication is fine-grained, irregular, or both. The library's support for asynchrony enables the application to aggressively overlap and schedule communication and computation to reduce wait times.

UPC++ is highly portable and runs on platforms from laptops to supercomputers, with native implementations for HPC interconnects. As a C++ library, it interoperates smoothly with existing numerical libraries and on-node programming models (e.g., OpenMP, CUDA).

In this tutorial we will introduce basic concepts and advanced optimization techniques of UPC++. We will discuss the UPC++ memory and execution models and walk through basic algorithm implementations. We will also look at irregular applications and show how they can take advantage of UPC++ features to optimize their performance.

DESCRIPTION
The tutorial goals are as follows:

- Introduce the Partitioned Global Address Space (PGAS) model and its importance to exascale applications, as well as the design of the UPC++ library that enables scalable PGAS programming.
- Present example UPC++ code which takes advantage of the library's features, including a distributed hash table and a 1-D Poisson kernel.

- Discuss UPC++ programming constructs that provide both productivity and performance in irregular applications.
- Introduce advanced optimization techniques and illustrate how one may leverage them to improve performance.
- Present case studies from two collaborations which have adopted UPC++, including the HipMer genome assembler and a multi-frontal sparse solver.

The topic is relevant to anyone who implements irregular applications on distributed-memory systems, including both physics-based algorithms such as adaptive mesh refinement, and non-physics-based applications such as metagenomics and graph analytics. Irregular applications are a challenge because they employ fine-grained communication. To support such applications, UPC++ provides both one-sided communication (RMA put, get and atomics) and Remote Procedure Calls (RPC). UPC++ is a C++ template library, providing the programmer with access to C++ productivity features, including the type system, standard library and lambdas. Both RMA and RPC incur low latency, thanks to the use of GASNet-EX in the UPC++ runtime implementation.

The tutorial will introduce basic concepts and advanced optimization techniques of UPC++. We will discuss the UPC++ memory and execution models and walk through how to implement basic algorithms in UPC++. We will also look at irregular applications and how to take advantage of UPC++ features to optimize communication performance. The tutorial concludes with a brief application performance study. Advanced topics will be discussed only briefly, as the session requires no advanced knowledge of PGAS programming.

The tutorial assumes no prior knowledge of UPC++. However, participants should be comfortable with C++11 (or newer) programming, including use of templates. We also expect attendees to have some prior knowledge or experience in parallel programming.

TITLE **Using HPCToolkit to Measure and Analyze the Performance of GPU-accelerated Applications**

FOCUS AREA **ST**
TYPE Tutorial
LEVEL Intermediate

DAY Wednesday, February 5, 2020
TIME 10:30 AM - 12:00 PM central
DURATION 90 min

AUTHORS
John Mellor-Crummey (Rice University) johnmc@rice.edu
Keren Zhou (Rice University) kz21@rice.edu

ABSTRACT
Tailoring applications for GPU-accelerated compute nodes is essential to harness the power of ORNL's Summit exascale testbed as well as the forthcoming GPU-accelerated A21 and Frontier exascale platforms.

This tutorial will (1) introduce new capabilities for performance measurement and analysis of GPU-accelerated codes that are available in the HPCToolkit performance tools and (2) describe how to use them to analyze and tune GPU-accelerated applications.

To support efficient monitoring of accelerated computations, HPCToolkit employs a novel wait-free data structures to coordinate measurement and attribution of performance metrics while a GPU-accelerated computation executes. To help developers understand the performance of accelerated applications as a whole, HPCToolkit tool attributes metrics to rich heterogeneous calling contexts that span both CPUs and GPUs and displays traces that include both CPU and GPU activity time lines. To help developers understand the performance of complex GPU code generated from high-level programming models such as OpenMP or template-based programming abstractions, HPCToolkit constructs sophisticated approximations of call path profiles for GPU computations from flat PC samples collected by NVIDIA GPUs. To support fine-grain analysis and tuning, HPCToolkit uses GPU PC samples to derive and attribute metrics, including measures of GPU latency and throughput at all levels in a heterogeneous calling context.

To make effective use of HPCToolkit for tuning GPU-accelerated applications, one must understand what performance metrics that HPCToolkit can collect, and how to use them to guide analysis and tuning.

We will illustrate the capabilities of HPCToolkit with case studies of various codes and mini-applications.

DESCRIPTION
Today, a prototype version of HPCToolkit uses NVIDIA's CUPTI interface to collect profiles and traces of GPU-accelerated codes written using CUDA or accelerated OpenMP. The addition of capabilities to support AMD GPUs is also underway. Intel has not yet provided information about performance tools APIs for the forthcoming A21 system, but we expect to resolve this issue soon.

The tutorial is intended to target ECP application, library, and framework developers who are interested in trying to use HPCToolkit to analyze the performance of a GPU-accelerated codes within and/or across GPU-accelerated nodes.

Besides describing how HPCToolkit works, we will discuss its current capabilities, its strengths and weaknesses, and how to use it to measure, analyze, and tune GPU-accelerated code. Measurement of GPU-accelerated compute nodes can generate a lot of information. Knowing how to explore this information in a top-down fashion is the key to using HPCToolkit effectively to guide program analysis and tuning.

APPENDIX D. LIST OF SESSION AUTHORS

Approximately 332 highly engaged session authors worked hard to prepare material for the 102 sessions (Keynote, plenaries, BoF/breakout/panels and tutorials) presented at the 2020 ECP annual meeting.

See list of sessions authors below.

First Name	Last Name	Institution
Shrirang	Abhyankar	Pacific Northwest National Laboratory
Anthony	Agelastos	Sandia National Laboratories
Hadia	Ahmed	Lawrence Berkeley National Laboratory
Dong	Ahn	Lawrence Livermore National Laboratory
Christine	Ahrens	Los Alamos National Laboratory
Francis	Alexander	Brookhaven National Laboratory
Ann	Almgren	Lawrence Berkeley National Laboratory
Hartwig	Anzt	Karlsruhe Institute of Technology
Thomas	Applencourt	Argonne National Laboratory
Scott	Atchley	Oak Ridge National Laboratory
Chuck	Atkins	Kitware Inc.
John	Bachan	Lawrence Berkeley National Laboratory
Scott	Baden	Lawrence Berkeley National Laboratory (retired)
Pavan	Balaji	Argonne National Laboratory
Ramesh	Balakrishnan	Argonne National Laboratory
Cody	Balos	Lawrence Livermore National Laboratory
Prith	Banerjee	ANSYS, Incorporated
Deborah	Bard	Lawrence Berkeley National Laboratory
Kevin	Barker	Pacific Northwest National Laboratory
Gregory	Becker	Lawrence Livermore National Laboratory
David	Beckingsale	Lawrence Livermore National Laboratory
Pete	Beckman	Argonne National Laboratory
James	Belak	Lawrence Livermore National Laboratory
Matt	Belhorn	Oak Ridge National Laboratory
John	Bell	Lawrence Berkeley National Laboratory
David	Bernholdt	Oak Ridge National Laboratory
Colleen	Bertoni	Argonne National Laboratory
Manoj	Bhardwaj	Sandia National Laboratories
Amitava	Bhattacharjee	Princeton Plasma Physics Laboratory
Robert	Bird	Los Alamos National Laboratory
Ayan	Biswas	Los Alamos National Laboratory
Matthew	Bolitho	NVIDIA Corporation
Erik	Boman	Sandia National Laboratories
Dan	Bonachea	Lawrence Berkeley National Laboratory

First Name	Last Name	Institution
George	Bosilca	The University of Tennessee, Knoxville
Kathlyn	Boudwin	Oak Ridge National Laboratory
Aurélien	Bouteiller	The University of Tennessee, Knoxville
Jim	Brandt	Sandia National Laboratories
Bill	Brantley	Advanced Micro Devices, Inc.
Scot	Breitenfeld	The HDF Group
Tom	Brettin	Argonne National Laboratory
Michael	Brim	Oak Ridge National Laboratory
Stephanie	Brink	Lawrence Livermore National Laboratory
Jed	Brown	University of Colorado Boulder
Paul	Bryant	New Mexico Consortium
Reuben	Budiardja	Oak Ridge National Laboratory
Aydin	Buluc	Lawrence Berkeley National Laboratory
Suren	Byna	Lawrence Berkeley National Laboratory
Shane	Canon	Lawrence Berkeley National Laboratory
Franck	Cappello	Argonne National Laboratory
Philip	Carns	Argonne National Laboratory
Tony	Castaldo	The University of Tennessee, Knoxville
Mohamad	Chaarawi	Intel Corporation
Barbara	Chapman	Stony Brook University
Wang	Chen	IBM
Jackie	Chen	Sandia National Laboratories
Andrew	Chien	Argonne National Laboratory
Hank	Childs	University of Oregon
Jong	Choi	Oak Ridge National Laboratory
Sutanay	Choudhury	Pacific Northwest National Laboratory
Sudheer	Chunduri	Argonne National Laboratory
Susan	Coghlan	Argonne National Laboratory
Terry	Cojean	Karlsruhe Institute of Technology
Tiffany	Connors	Lawrence Berkeley National Laboratory
Brandon	Cook	Lawrence Berkeley National Laboratory
Jeanine	Cook	Sandia National Laboratories
Chris	Daley	Lawrence Berkeley National Laboratory
Anthony	Danalis	The University of Tennessee, Knoxville
David	Daniel	Los Alamos National Laboratory
John	Danskin	NVIDIA Corporation
Philip	Davis	Rutgers University
Bronis	de Supinski	Lawrence Livermore National Laboratory
Joel	Denny	Oak Ridge National Laboratory
Nicolas	Denoyelle	Argonne National Laboratory

First Name	Last Name	Institution
Luiz	DeRose	Cray Inc.
Jack	Deslippe	Lawrence Berkeley National Laboratory
Sheng	Di	Argonne National Laboratory
Francesco	Di Natale	Lawrence Livermore National Laboratory
Lori	Diachin	Lawrence Livermore National Laboratory
Emily	Dietrich	Argonne National Laboratory
Veselin	Dobrev	Lawrence Livermore National Laboratory
Johannes	Doerfert	Argonne National Laboratory
Douglas	Doerfler	Lawrence Berkeley National Laboratory
Matthieu	Dorier	Argonne National Laboratory
Erik	Draeger	Lawrence Livermore National Laboratory
Nicolas	Dube	Hewlett Packard Enterprise
Anshu	Dubey	Argonne National Laboratory
Soumya	Dutta	Los Alamos National Laboratory
Rob	Egan	Lawrence Berkeley National Laboratory
Greg	Eisenhauer	Georgia Institute of Technology
Dan	Ernst	Cray Inc.
Thomas	Evans	Oak Ridge National Laboratory
Mark	Fahey	Argonne National Laboratory
Paolo	Faraboschi	Hewlett Packard Enterprise
Hal	Finkel	Argonne National Laboratory
Paul	Fischer	University of Illinois at Urbana-Champaign
Dennis	Floyd	Hewlett Packard Enterprise
Ian	Foster	Argonne National Laboratory
Marianne	Francois	Los Alamos National Laboratory
Chip	Freitag	Advanced Micro Devices, Inc.
Brian	Friesen	Lawrence Berkeley National Laboratory
Zack	Galbreath	Kitware Inc.
Todd	Gamblin	Lawrence Livermore National Laboratory
David	Gardner	Lawrence Livermore National Laboratory
Jeff	Garelick	Intel Corporation
Rene	Gassmoeller	University of California, Davis
Mark	Gates	The University of Tennessee, Knoxville
Rahul	Gayatri	Lawrence Berkeley National Laboratory
Damien	Genet	The University of Tennessee, Knoxville
Yasaman	Ghadar	Argonne National Laboratory
Pieter	Ghysels	Lawrence Berkeley National Laboratory
Roscoe	Giles	Boston University
Joe	Glenski	Hewlett Packard Enterprise
Kent	Glossop	Intel Corporation

First Name	Last Name	Institution
William	Godoy	Oak Ridge National Laboratory
Maya	Gokhale	Lawrence Livermore National Laboratory
Elsa	Gonsiorowski	Lawrence Livermore National Laboratory
Kevin	Gott	Lawrence Berkeley National Laboratory
John	Gounley	Oak Ridge National Laboratory
Jen	Green	Los Alamos National Laboratory
Taylor	Groves	Lawrence Berkeley National Laboratory
Alicia	Grundhoffer	Oak Ridge National Laboratory
Junmin	Gu	Lawrence Berkeley National Laboratory
Chin	Guok	Lawrence Berkeley National Laboratory
Rinku	Gupta	Argonne National Laboratory
Salman	Habib	Argonne National Laboratory
Mahantesh	Halappanavar	Pacific Northwest National Laboratory
Frank	Ham	Cascade Technologies, Incorporated
Jeff	Hammond	Intel Corporation
Paul	Hargrove	Lawrence Berkeley National Laboratory
Ruud	Haring	IBM Research
Kevin	Harms	Argonne National Laboratory
Austin	Harris	Oak Ridge National Laboratory
Cyrus	Harrison	Lawrence Livermore National Laboratory
William	Hart	Sandia National Laboratories
Scott	Hemmert	Sandia National Laboratories
Thomas	Herault	The University of Tennessee, Knoxville
Stephen	Herbein	Lawrence Livermore National Laboratory
Oscar	Hernandez	Oak Ridge National Laboratory
Michael	Heroux	Sandia National Laboratories
Judith	Hill	Oak Ridge National Laboratory
Steven	Hofmeyr	Lawrence Berkeley National Laboratory
Brian	Homerding	Argonne National Laboratory
Richard	Hornung	Lawrence Livermore National Laboratory
Aimee	Hungerford	Los Alamos National Laboratory
Sachin	Idgunji	NVIDIA Corporation
Kamil	Iskra	Argonne National Laboratory
Douglas	Jacobsen	Lawrence Berkeley National Laboratory
Mathias	Jacquelin	Lawrence Berkeley National Laboratory
Heike	Jagode	The University of Tennessee, Knoxville
Gustav	Jansen	Oak Ridge National Laboratory
Terry	Jones	Oak Ridge National Laboratory
Christoph	Junghans	Los Alamos National Laboratory
Jim	Kahle	IBM Research

First Name	Last Name	Institution
Vivek	Kale	Brookhaven National Laboratory
Amir	Kamil	Lawrence Berkeley National Laboratory
Ian	Karlin	Lawrence Livermore National Laboratory
Brunon	Kepczynski	General Electric
Jason	Kincl	Oak Ridge National Laboratory
Scott	Klasky	Oak Ridge National Laboratory
Chris	Knight	Argonne National Laboratory
Tzanio	Kolev	Lawrence Livermore National Laboratory
Doug	Kothe	Oak Ridge National Laboratory
Quincey	Koziol	Lawrence Berkeley National Laboratory
Michael	Kruse	Argonne National Laboratory
Kalyan	Kumaran	Argonne National Laboratory
Thorsten	Kurth	Lawrence Berkeley National Laboratory
Jakub	Kurzak	Advanced Micro Devices, Inc.
Jae Hyuk	Kwack	Argonne National Laboratory
Ignacio	Laguna Peralta	Lawrence Livermore National Laboratory
Mike	Lang	Los Alamos National Laboratory
Michael	Langer	Cray Inc.
Matt	Larsen	Lawrence Livermore National Laboratory
Rob	Latham	Argonne National Laboratory
Ti	Leggett	Argonne National Laboratory
Matt	Legrande	Lawrence Livermore National Laboratory
Sherry	Li	Lawrence Berkeley National Laboratory
Peter	Lindstrom	Lawrence Livermore National Laboratory
Vicki	Little	NVIDIA Corporation
Yang	Liu	Lawrence Berkeley National Laboratory
Gabriel	Loh	Advanced Micro Devices, Inc.
Johann	Lombardi	Intel Corporation
Michael	Lough	Hewlett Packard Enterprise
Ray	Loy	Argonne National Laboratory
Ye	Luo	Argonne National Laboratory
Piotr	Luszczek	The University of Tennessee, Knoxville
Brian	Mackie-Mason	Argonne National Laboratory
Nick	Malaya	Advanced Micro Devices, Inc.
Tanu	Malik	DePaul University
Aniruddha	Marathe	Lawrence Livermore National Laboratory
George	Markomanolis	Oak Ridge National Laboratory
David	Martin	Argonne National Laboratory
Mike	Mason	Los Alamos National Laboratory
David	McCallen	Lawrence Berkeley National Laboratory

First Name	Last Name	Institution
Patrick	McCormick	Los Alamos National Laboratory
Damon	McDougall	Advanced Micro Devices, Inc.
Marty	McFadden	Lawrence Livermore National Laboratory
Lois	McInnes	Argonne National Laboratory
Kshitij	Mehta	Oak Ridge National Laboratory
Mario	Melara	Lawrence Berkeley National Laboratory
John	Mellor-Crummey	Rice University
Thomas	Mendoza	Lawrence Livermore National Laboratory
Alan	Menezes	NVIDIA Corporation
Bronson	Messer	Oak Ridge National Laboratory
Mamun	Miah	Lawrence Berkeley National Laboratory
Doug	Miles	NVIDIA Corporation
Reed	Milewicz	Sandia National Laboratories
Mark	Miller	Lawrence Livermore National Laboratory
Daniel	Milroy	Lawrence Livermore National Laboratory
Misun	Min	Argonne National Laboratory
Susan	Mniszewski	Los Alamos National Laboratory
Jamaludin	Mohd Yusof	Los Alamos National Laboratory
Kathryn	Mohror	Lawrence Livermore National Laboratory
Florence	Monna	Argonne National Laboratory
Jose	Monsalve Diaz	Argonne National Laboratory
David	Montoya	Trenza Synergy
Roy	Moore	Intel Corporation
Shirley	Moore	Oak Ridge National Laboratory
Ken	Moreland	Sandia National Laboratories
David	Moulton	Los Alamos National Laboratory
Misbah	Mubarak	Argonne National Laboratory
Todd	Munson	Argonne National Laboratory
Andrew	Myers	Lawrence Berkeley National Laboratory
Hai Ah	Nam	Los Alamos National Laboratory
Thomas	Naughton	Oak Ridge National Laboratory
Jeff	Neel	Argonne National Laboratory
Rob	Neely	Lawrence Livermore National Laboratory
David	Nicholaeff	New Mexico Consortium
Bogdan	Nicolae	Argonne National Laboratory
Kyle	Niemeyer	Oregon State University
Bill	Nitzberg	Altair Engineering, Incorporated
Matthew	Norman	Oak Ridge National Laboratory
Jose	Noudohouenou	Intel Corporation
Peter	Nugent	Lawrence Berkeley National Laboratory

First Name	Last Name	Institution
Lenny	Oliker	Lawrence Berkeley National Laboratory
Stephen	Olivier	Sandia National Laboratories
Sarp	Oral	Oak Ridge National Laboratory
James	Osborn	Argonne National Laboratory
Scott	Pakin	Los Alamos National Laboratory
Michael	Papka	Argonne National Laboratory
Manish	Parashar	Rutgers University
Scott	Parker	Argonne National Laboratory
Tapasya	Patki	Lawrence Livermore National Laboratory
John	Pennycook	Intel Corporation
Swann	Perarnau	Argonne National Laboratory
Alexis	Perry-Holby	Los Alamos National Laboratory
Anders	Petersson	Lawrence Livermore National Laboratory
Irina	Petrova	Intel Corporation
Steve	Petruzza	University of Utah
Arben	Pitarka	Lawrence Livermore National Laboratory
Norbert	Podhorszki	Oak Ridge National Laboratory
David	Poliakoff	Sandia National Laboratories
Alex	Pothen	Purdue University
Elena	Pourmal	The HDF Group
Howard	Pritchard	Los Alamos National Laboratory
Terri	Quinn	Lawrence Livermore National Laboratory
Siva	Rajamanickam	Sandia National Laboratories
Jini	Ramprakash	Argonne National Laboratory
Elaine	Raybourn	Sandia National Laboratories
Noah	Reddell	Cray Inc.
Daniel	Reynolds	Southern Methodist University
Ryan	Richard	Ames Laboratory
David	Richards	Lawrence Livermore National Laboratory
William	Rider	Sandia National Laboratories
Robert	Rieben	Lawrence Livermore National Laboratory
Katherine	Riley	Argonne National Laboratory
Bob	Robey	Los Alamos National Laboratory
Gregory	Rodgers	Advanced Micro Devices, Inc.
Arthur	Rodgers	Lawrence Livermore National Laboratory
David	Rogers	Los Alamos National Laboratory
Carlos	Rosales-Fernandez	Intel Corporation
Rob	Ross	Argonne National Laboratory
Philip	Roth	Oak Ridge National Laboratory
Barry	Rountree	Lawrence Livermore National Laboratory

First Name	Last Name	Institution
Nicholas	Sauter	Lawrence Berkeley National Laboratory
Nirmal	Saxena	NVIDIA Corporation
Stephen	Scalpone	NVIDIA Corporation
Mike	Schulte	Advanced Micro Devices, Inc.
Tom	Scogland	Lawrence Livermore National Laboratory
Sudip	Seal	Oak Ridge National Laboratory
Jim	Sexton	IBM Research
Arjun	Shankar	Oak Ridge National Laboratory
Kyle	Shaver	Argonne National Laboratory
Sameer	Shende	University of Oregon
Mark	Shephard	Rensselaer Polytechnic Institute
Andrew	Siegel	Argonne National Laboratory
Alex	Sim	Lawrence Berkeley National Laboratory
Stuart	Slattery	Oak Ridge National Laboratory
Shane	Snyder	Argonne National Laboratory
Annemarie	Southwell	NVIDIA Corporation
Jay	Srinivasan	Lawrence Berkeley National Laboratory
Simon	Steely	Intel Corporation
George	Stelle	Los Alamos National Laboratory
Rick	Stevens	Argonne National Laboratory
Greg	Stoner	Intel Corporation
Eric	Suchyta	Oak Ridge National Laboratory
Samantika	Sury	Intel Corporation
Madhava	Syamlal	National Energy Technology Laboratory
Houjun	Tang	Lawrence Berkeley National Laboratory
XinMin	Tian	Intel Corporation
Suzy	Tichenor	Oak Ridge National Laboratory
Stan	Tomov	The University of Tennessee, Knoxville
John	Tramm	Argonne National Laboratory
Christian	Trott	Sandia National Laboratories
Tom	Uram	Argonne National Laboratory
Brian	Van Essen	Lawrence Livermore National Laboratory
Rene	Van Oostrum	Advanced Micro Devices, Inc.
Arturo	Vargas	Lawrence Livermore National Laboratory
Jean-Luc	Vay	Lawrence Berkeley National Laboratory
Jeffrey	Vetter	Oak Ridge National Laboratory
Brice	Videau	Argonne National Laboratory
Venkat	Vishwanath	Argonne National Laboratory
Tim	Warburton	Virginia Polytechnic Institute and State University
Justin	Whitt	Oak Ridge National Laboratory

First Name	Last Name	Institution
Jeremiah	Wilke	Sandia National Laboratories
Jim	Willenbring	Sandia National Laboratories
Tim	Williams	Argonne National Laboratory
Samuel	Williams	Lawrence Berkeley National Laboratory
Frank	Willmore	Argonne National Laboratory
Theresa	Windus	Ames Laboratory / Iowa State University
Matthew	Wolf	Oak Ridge National Laboratory
Paul	Wolfenbarger	Sandia National Laboratories
Carol	Woodward	Lawrence Livermore National Laboratory
Justin	Wozniak	Argonne National Laboratory
K. John	Wu	Lawrence Berkeley National Laboratory
Igor	Yakushin	Argonne National Laboratory
Charlene	Yang	Lawrence Berkeley National Laboratory
Ulrike	Yang	Lawrence Livermore National Laboratory
Katherine	Yelick	Lawrence Berkeley National Laboratory
Hong-Jun	Yoon	Oak Ridge National Laboratory
Andrew	Younge	Sandia National Laboratories
Keren	Zhou	Rice University

APPENDIX E. LIST OF POSTER PRESENTERS

To the best of our knowledge, 146 poster presenters participated in the poster sessions at the 2020 ECP Annual Meeting.

See the list of poster presenters below.

First Name	Last Name	Institution
Omar	Aaziz	Sandia National Laboratories
Dong	Ahn	Lawrence Livermore National Laboratory
Francis	Alexander	Brookhaven National Laboratory
Vassil	Alexandrov	Hartree Centre, UK Research and Innovation
Hartwig	Anzt	The University of Tennessee, Knoxville
Muaaz	Awan	Lawrence Berkeley National Laboratory
Alan	Ayala	The University of Tennessee, Knoxville
Prasanna	Balaprakash	Argonne National Laboratory
Deborah	Bard	Lawrence Berkeley National Laboratory
Prithayan	Barua	Georgia Institute of Technology
David	Beckingsale	Lawrence Livermore National Laboratory
Jim	Belak	Lawrence Livermore National Laboratory
John	Bell	Lawrence Berkeley National Laboratory
David	Bernholdt	Oak Ridge National Laboratory
Sridutt	Bhalachandra	Lawrence Berkeley National Laboratory
Michael	Blocksome	Intel
Erik	Boman	Sandia National Laboratories
Aurelien	Bouteiller	The University of Tennessee, Knoxville
Peter	Boyle	Brookhaven National Laboratory
Hugo	Brunie	Lawrence Berkeley National Laboratory
Paul	Bryant	New Mexico Consortium
Suren	Byna	Lawrence Berkeley National Laboratory
Jackie	Chen	Sandia National Laboratories
Hank	Childs	University of Oregon
Terry	Cojean	Karlsruhe Institute of Technology
Tiffany	Connors	Lawrence Berkeley National Laboratory
David	Daniel	Los Alamos National Laboratory
Wibe	de Jong	Lawrence Berkeley National Laboratory
Joel	Denny	Oak Ridge National Laboratory
Jack	Deslippe	Lawrence Berkeley National Laboratory
Sheng	Di	Argonne National Laboratory
Veselin	Dobrev	Lawrence Livermore National Laboratory
Johannes	Doerfert	Argonne National Laboratory
Tom	Ebbott	Goodyear

First Name	Last Name	Institution
Stephane	Ethier	Princeton Plasma Physics Laboratory
Darel	Finkbeiner	Department of Energy
Ian	Foster	Argonne National Laboratory
Todd	Gamblin	Lawrence Livermore National Laboratory
David	Gardner	Lawrence Livermore National Laboratory
Richard	Gerber	Lawrence Berkeley National Laboratory
Yasaman	Ghadar	Argonne National Laboratory
Sayan	Ghosh	Pacific Northwest National Laboratory
Pieter	Ghysels	Lawrence Berkeley National Laboratory
William	Godoy	Oak Ridge National Laboratory
Anastasia	Gunina	Ames Laboratory
Yanfei	Guo	Argonne National Laboratory
Rinku	Gupta	Argonne National Laboratory
Salman	Habib	Argonne National Laboratory
Mahantesh	Halappanavar	Pacific Northwest National Laboratory
Mary	Hall	University of Utah
Steven	Hamilton	Oak Ridge National Laboratory
Paul	Hargrove	Lawrence Berkeley National Laboratory
Thomas	Herault	The University of Tennessee, Knoxville
Richard	Hornung	Lawrence Livermore National Laboratory
Kevin	Huck	University of Oregon
Aimee	Hungerford	Los Alamos National Laboratory
Keita	Iwabuchi	Lawrence Livermore National Laboratory
Shintaro	Iwasaki	Argonne National Laboratory
Heike	Jagode	The University of Tennessee, Knoxville
Ivo	Jimenez	University of California, Santa Cruz
Hans	Johansen	Lawrence Berkeley National Laboratory
Seth	Johnson	Oak Ridge National Laboratory
Balint	Joo	Thomas Jefferson National Accelerator Facility
Vivek	Kale	Brookhaven National Laboratory
Amir	Kamil	Lawrence Berkeley National Laboratory
Daniel	Kasen	Lawrence Berkeley National Laboratory
Alison	Kennedy	Hartree Centre, UK Research and Innovation
Paul	Kent	Oak Ridge National Laboratory
Jungwon	Kim	Oak Ridge National Laboratory
Jason	Kincl	Oak Ridge National Laboratory
Tzanio	Kolev	Lawrence Livermore National Laboratory
Venkata	Krishnan	Intel
Michael	Kruse	Argonne National Laboratory
Michael	Lang	Los Alamos National Laboratory

First Name	Last Name	Institution
Yu	Lechen	Georgia Institute of Technology
Sherry	Li	Lawrence Berkeley National Laboratory
Peter	Lindstrom	Lawrence Livermore National Laboratory
Yang	Liu	Lawrence Berkeley National Laboratory
Ray	Loy	Argonne National Laboratory
Piotr	Luszczek	The University of Tennessee, Knoxville
Jonathan	Maack	National Renewable Energy Laboratory
Allen	Malony	University of Oregon
Osni	Marques	Lawrence Berkeley National Laboratory
Luke	Mason	Hartree Centre, UK Research and Innovation
Patrick	McCormick	Los Alamos National Laboratory
Peter	McCorquodale	Lawrence Berkeley National Laboratory
Lois	McInnes	Argonne National Laboratory
Bronson	Messer	Oak Ridge National Laboratory
Daniel	Milroy	Lawrence Livermore National Laboratory
Sue	Mniszewski	Los Alamos National Laboratory
Kathryn	Mohror	Lawrence Livermore National Laboratory
Todd	Munson	Argonne National Laboratory
Bogdan	Nicolae	Argonne National Laboratory
Leonid	Oliker	Lawrence Berkeley National Laboratory
Stephen	Olivier	Sandia National Laboratories
Sarah	Osborn	Lawrence Livermore National Laboratory
Scott	Pakin	Los Alamos National Laboratory
Slaven	Peles	Pacific Northwest National Laboratory
Ivy	Peng	Lawrence Livermore National Laboratory
Amedeo	Perazzo	SLAC National Accelerator Laboratory
Danny	Perez	Los Alamos National Laboratory
Cosmin	Petra	Lawrence Livermore National Laboratory
Swaroop	Pophale	Oak Ridge National Laboratory
Rebel	Powell	Department of Energy
David	Pugmire	Oak Ridge National Laboratory
Siva	Rajamanickam	Sandia National Laboratories
Jini	Ramprakash	Argonne National Laboratory
Timothy	Randles	Los Alamos National Laboratory
Elaine	Raybourn	Sandia National Laboratories
Valentin	Reis	Argonne National Laboratory
Robert	Rieben	Lawrence Livermore National Laboratory
David	Rogers	Los Alamos National Laboratory
Rob	Ross	Argonne National Laboratory
Michael	Rowan	Lawrence Berkeley National Laboratory

First Name	Last Name	Institution
Olivier	Serres	Intel
Mallikarjun	Shankar	Oak Ridge National Laboratory
Kyle	Shaver	Argonne National Laboratory
Sameer	Shende	University of Oregon
Galen	Shipman	Los Alamos National Laboratory
Bjorn	Sjogreen	Lawrence Livermore National Laboratory
Stuart	Slattery	Oak Ridge National Laboratory
Michael	Sprague	National Renewable Energy Laboratory
George	Stelle	Los Alamos National Laboratory
Miroslav	Stoyanov	Oak Ridge National Laboratory
Madhava	Syamlal	National Energy Technology Laboratory
Houjun	Tang	Lawrence Berkeley National Laboratory
Mark	Taylor	Sandia National Laboratories
Stan	Tomov	The University of Tennessee, Knoxville
Richard	Tran Mills	Argonne National Laboratory
Christian	Trott	Sandia National Laboratories
Aristeidis	Tsaris	Oak Ridge National Laboratory
Zhao	Tuowen	University of Utah
John	Turner	Oak Ridge National Laboratory
Terry	Turton	Los Alamos National Laboratory
Brian	Van Straalen	Lawrence Berkeley National Laboratory
Arturo	Vargas	Lawrence Livermore National Laboratory
Jean-Luc	Vay	Lawrence Berkeley National Laboratory
Jeffrey	Vetter	Oak Ridge National Laboratory
Gregory	Watson	Oak Ridge National Laboratory
Evan	Weinberg	NVIDIA
Sam	Williams	Lawrence Berkeley National Laboratory
Ulrike	Yang	Lawrence Livermore National Laboratory
Asim	YarKhan	The University of Tennessee, Knoxville
Hyunseung	Yoo	Argonne National Laboratory
Chuck	Yoon	SLAC National Accelerator Laboratory
Andrew	Younge	Sandia National Laboratories

APPENDIX F. EXECUTIVE SUMMARY SLIDES

In addition to this report, the committee prepared a presentation to quantitatively capture the numbers behind the meeting. The slide deck contains detailed statistics and graphs with information about the sessions, the registration and attendance, services provided and summary results of the final survey. This Appendix contains the slides for reference.

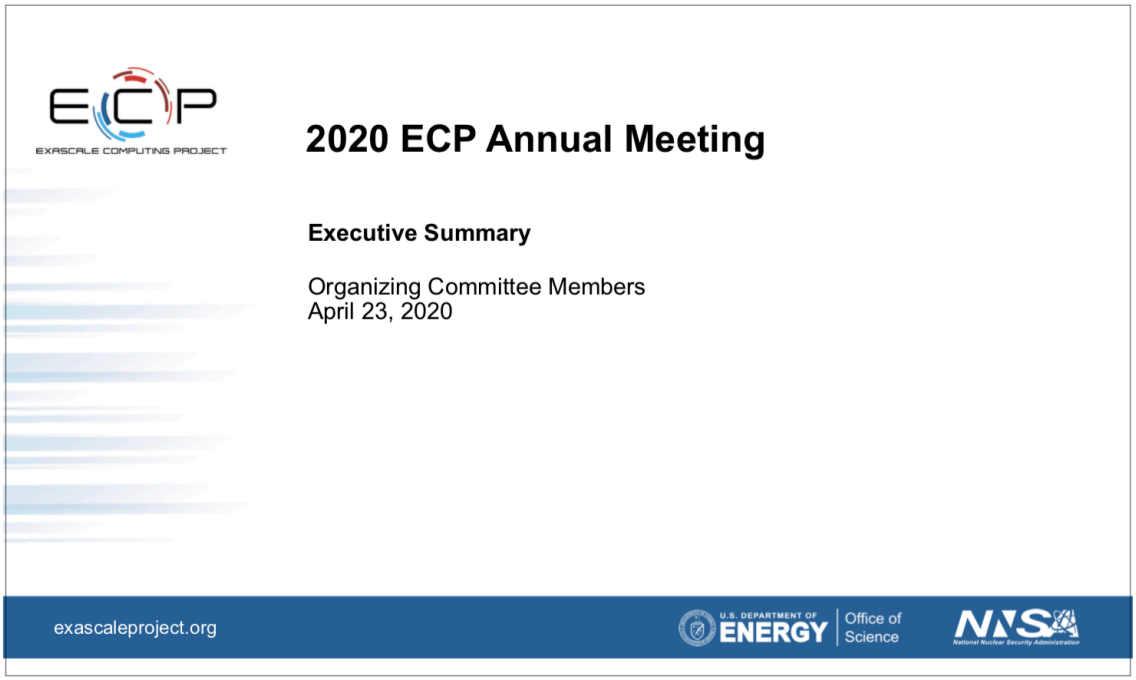

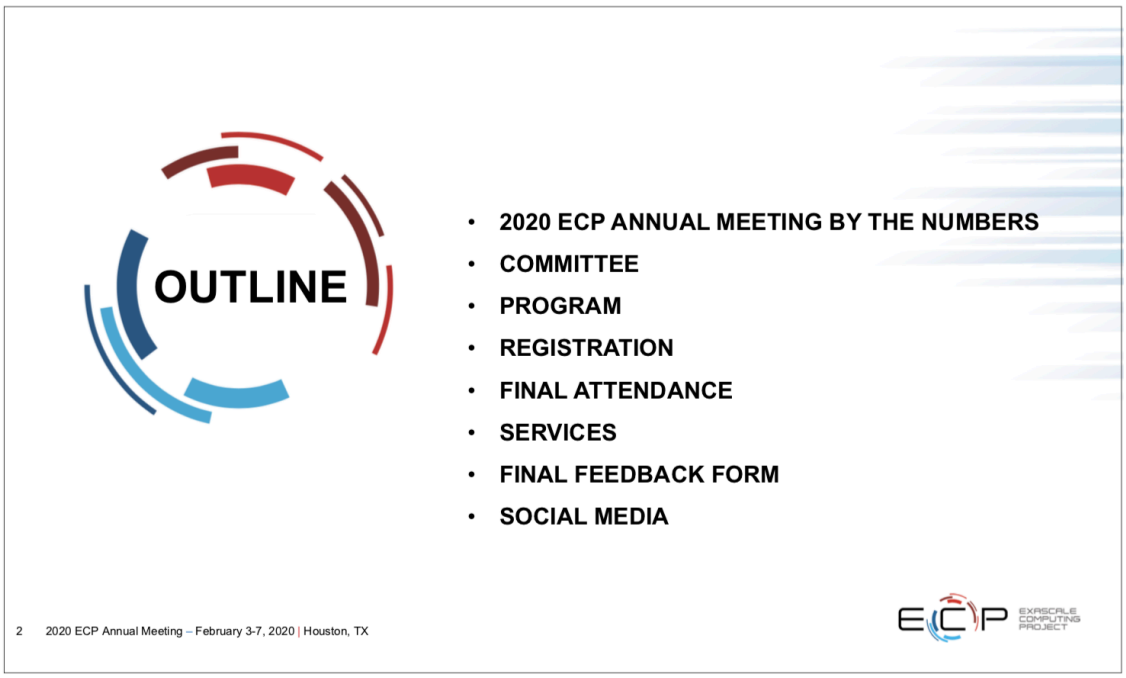

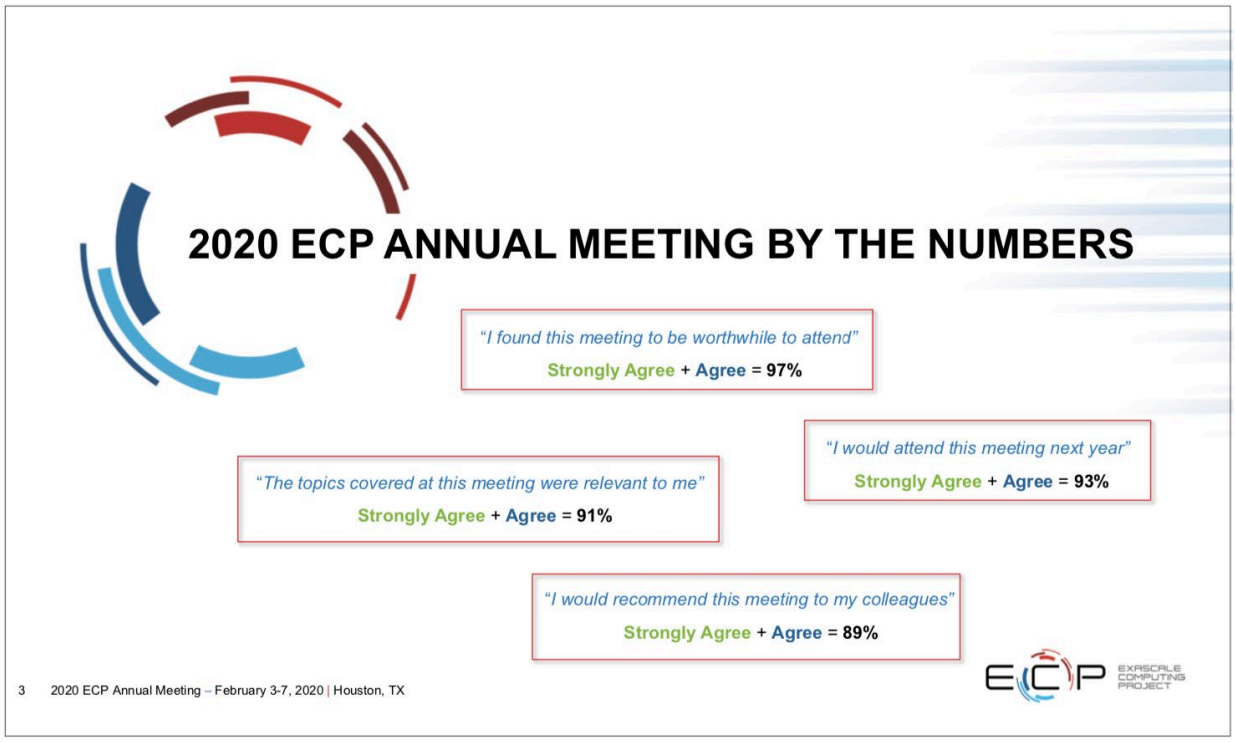

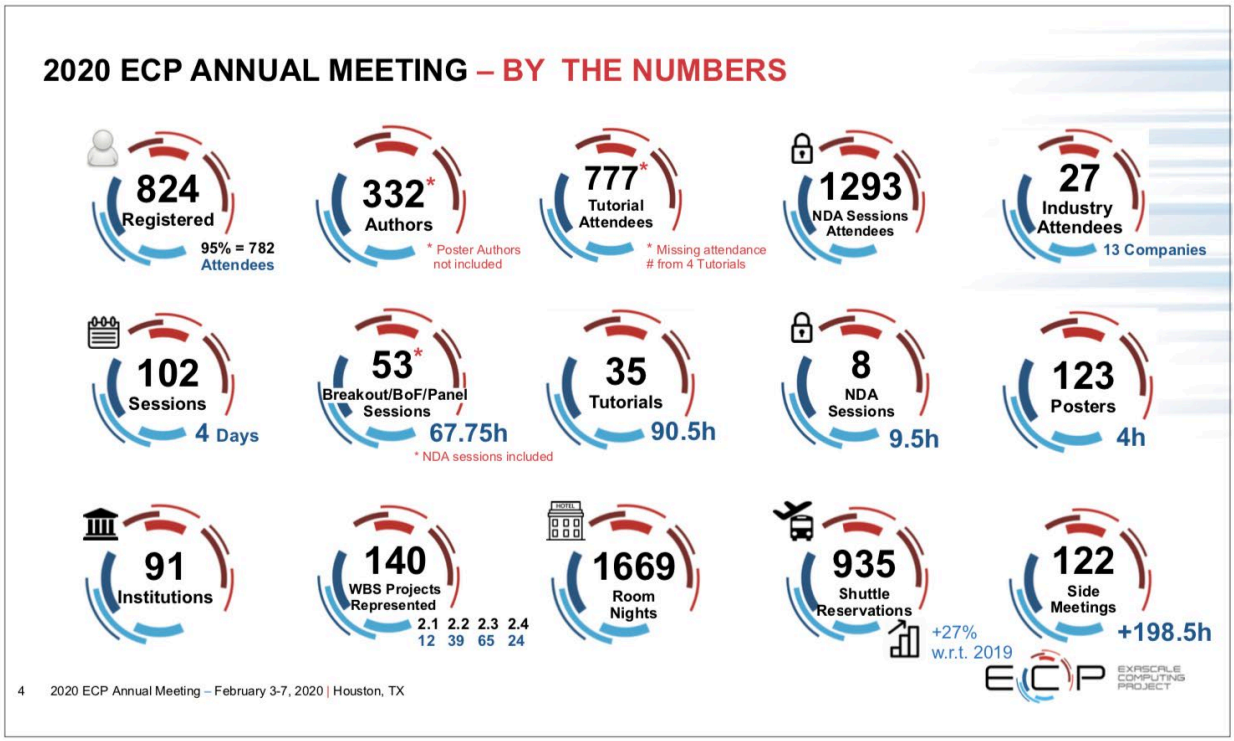

- COMMITTEE MEMBERS
- COMMITTEE WORKING AREA
- COMMITTEE TOOL

COMMITTEE MEMBERS

The organizing committee is composed of skilled members from different national laboratories to assist on the process of planning, managing and executing the numerous tasks that need to be done for the 2020 ECP Annual Meeting. All the representatives of the committee worked together on the organization and planning of the meeting to ensure a successful ECP Annual Meeting.

The committee opened calls and hosted sessions on key topics and areas of interest to the ECP community including discussions of future systems, software stack plans, and interactions with facilities. In addition, the committee prepared the meeting schedule and was responsible for the technical content, including breakouts, tutorial and poster sessions offered through the week.

KEYNOTE
Kathlyn Boudwin – ORNL Email

COMMUNICATIONS
Mike Bernhardt – ORNL Email

PROJECT OFFICE
Doug Collins – ORNL Email

PROGRAM
Marta Garcia – ANL Email
Todd Munson – ANL Email

INDUSTRY COUNCIL
David Martin – ANL Email
Suzy Tichenor – ORNL Email

LOGISTICS
Corliss Thompson – ORNL Email
Ashton Chatman – ORNL Email
Lora Wolfe – ORNL Email
Judy Potok – ORNL Email

POSTERS
Lora Wolfe – ORNL Email
Elaine Raybourn – SNL Email

PATHFORWARD VENDORS
Bronis de Supinski – LLNL Email

INFORMATION TECHNOLOGY
Doug Collins – ORNL Email
Jonathan Wilson – ORNL Email

TRAINING
Ashley Barker – ORNL Email
Osni Marques – LBNL Email

COMMITTEE WORKING AREA

https://confluence.exascaleproject.org/display/2020ECPAMW

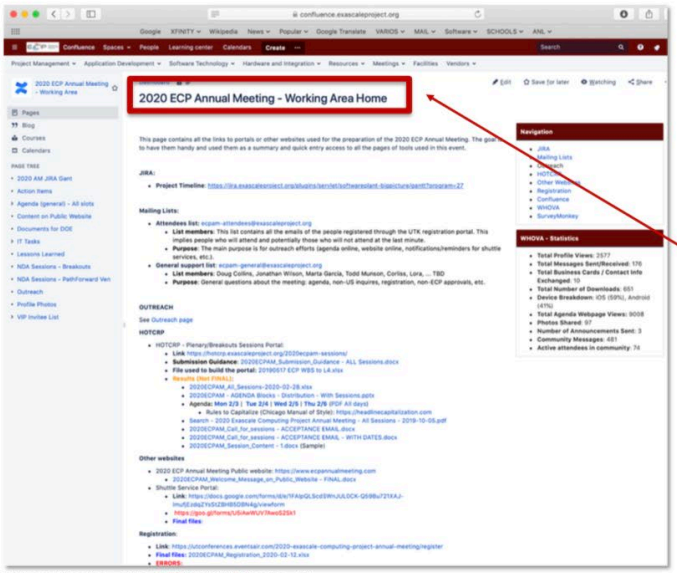

This page contains all links, results and final files created and/or gathered during the preparation, organization and execution of the annual meeting.

Any important file used by committee members is here as a record of their work and for reference for future committee members.

COMMITTEE TOOL

https://jira.exascaleproject.org/plugins/servlet/softwareplant-bigpicture/gantt.html?program=PROG-218

New this year!

Gantt chart on JIRA with task/subtasks to track project progress

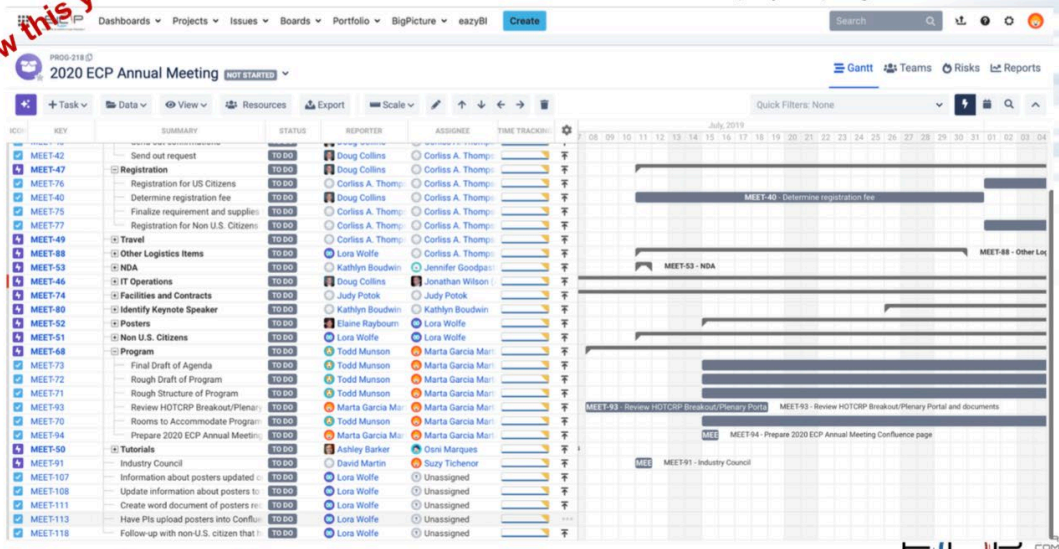

PROGRAM

- AGENDA
- AUTHORS FROM ALL SESSIONS
- KEYNOTE ADDRESS
- PLENARIES
- BOF/BREAKOUT/PANEL SESSIONS
- NDA SESSIONS
- TUTORIALS
- POSTERS
- ECP-O

AGENDA

102 Sessions
170h of content

4* days
* No agenda on Fri, 2/7/20. Day available for side meetings only

Number of total sessions by category

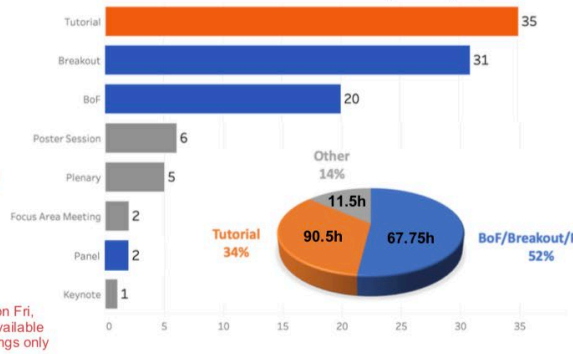

Number of sessions by Focus Area Topic

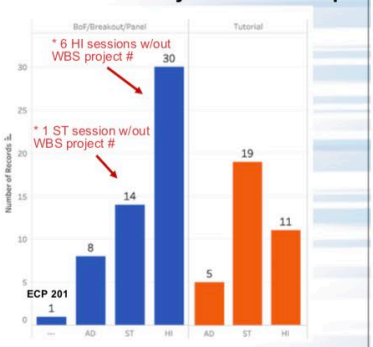

Monday, February 3, 2020

Tuesday, February 4, 2020

Wednesday, February 5, 2020

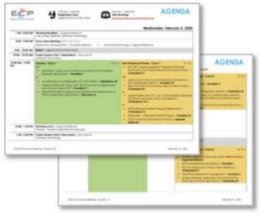

Thursday, February 6, 2020

AGENDA

To optimize the agenda of the 2020 ECP Annual Meeting, the committee used graph techniques and followed methodology from conflict resolution theory

Scientific Achievement

Maximization of the number of sessions, contributing authors, available seats and number of rooms that constituted the technical program of the 2020 Exascale Computing Project Annual Meeting.

Significance and Impact

In-person meetings that bring together teams are essential to foster more complex strategic thinking and create the psychological safety that ensures full engagement, promotes collaboration and facilitates the exchange of ideas.

Strategic planning of thousands of cumulative hours together for experts from National Labs, Universities, U.S. HPC vendors, DOE facilities, Industry Partners, and program sponsors.

Research Details

- Study of layouts and movements to increase opportunities for small talks, specially during poster sessions and breaks.
- Facilitate opportunities to get people out of their normal environment to make them think differently.
- Assign agenda slots while considering 10 conflict resolution strategies.
- Promote and enhance side-meetings experience.

AUTHORS FROM ALL SESSIONS

332* Authors
* Poster Authors not included

44 Institutions

1-5* Submissions per author
* Impact in sessions scheduling

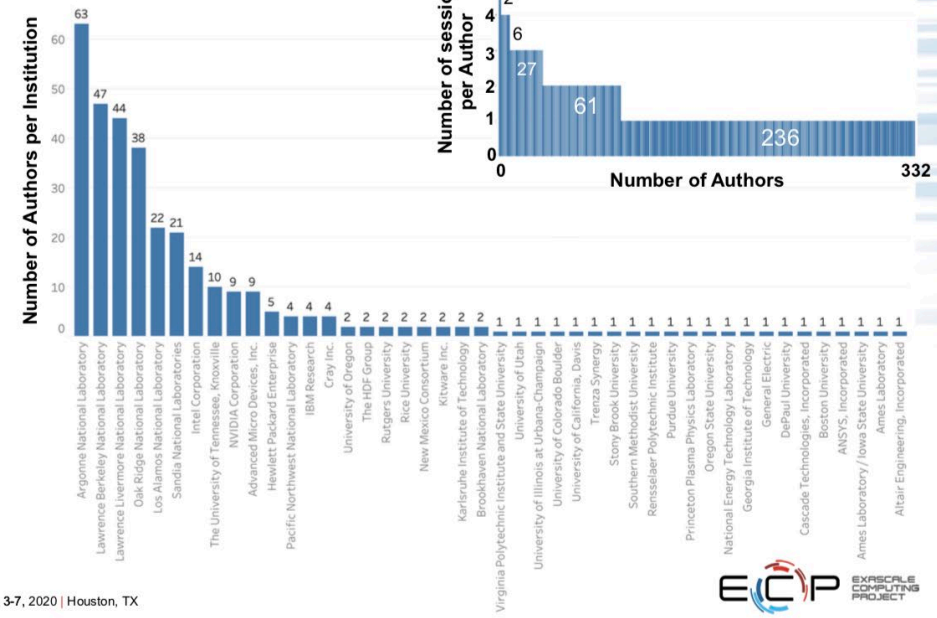

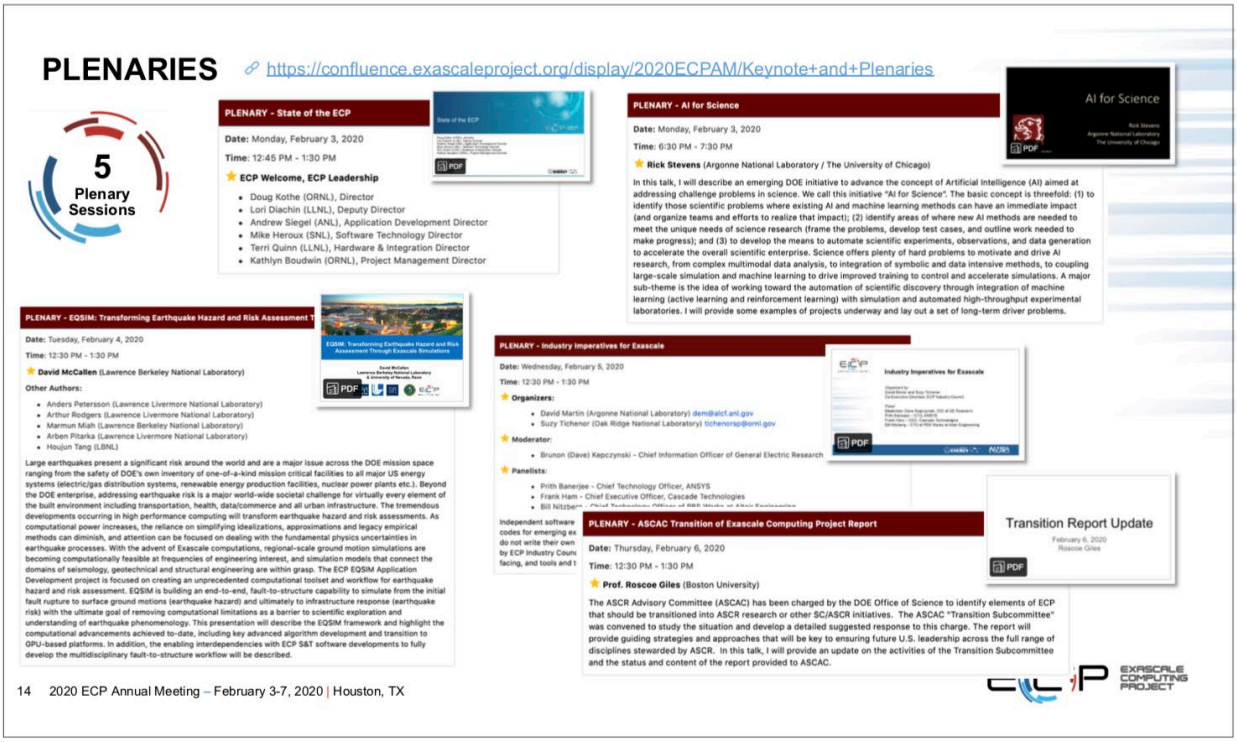

BOF/BREAKOUT/PANEL SESSIONS

https://confluence.exascaleproject.org/display/2020ECPAM/Sessions

53* Breakout Sessions

* NDA Sessions included

67.75 Hours

35 WBS Projects involved

+ 7 submissions: w/out WBS Project #

DURATION [MIN]

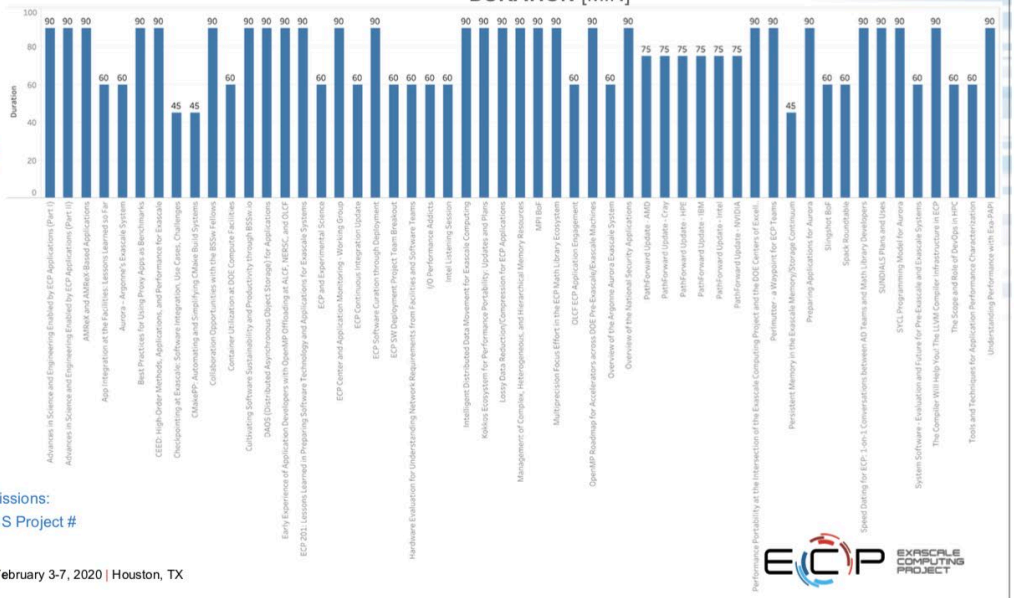

BOF/BREAKOUT/PANEL SESSIONS

2.2 APPLICATION DEVELOPMENT (AD)

8 WBS AD Projects involved

8 WBS AD Sessions presented

NUMBER OF PROJECTS

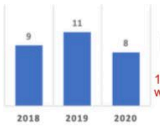

14% decrease w.r.t. 2019

NUMBER OF SESSIONS

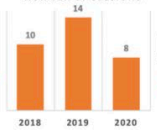

14% decrease w.r.t. 2019

- 2020 and 2019 data come directly from the information authors included in the proposal submission.
- 2018 data is an estimation based on project names and final sessions in the agenda.

Number of sessions with speakers/topics from WBS AD projects

WBS Name	# of Sessions 2018	# of Sessions 2019	# of Sessions 2020
2.2 Application Development			
2.2.1 ---- Chemistry and Materials Applications			1
2.2.1.01 LatticeQCD: Exascale Lattice Gauge Theory Opportunities/Reqmts for Nuclear & High Energy Physics	1		
2.2.1.02 NWChemEx: Tackling Chemical, Materials, & Biomolecular Challenges in Exascale	1		1
2.2.1.03 GAMESS: Enabling GAMESS for Exascale Computing in Chemistry & Materials			
2.2.1.04 EXAALT: Molecular Dynamics at the Exascale - Materials Science		1	
2.2.1.05 ExaAM: Transforming Additive Manufacturing through Exascale Simulation (ExaAM)			
2.2.1.06 QMCPACK: Predictive and Improvable Quantum-mechanics Based Simulations			
2.2.2 ---- Energy Applications			
2.2.2.01 ExaWind: Exascale Predictive Wind Plant Flow Physics Modeling	1		
2.2.2.02 Combustion-Pele: Transforming Combustion Science & Technology with Exascale Simulations			
2.2.2.03 ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors			
2.2.2.04 MFIX-Exa: Perf Prediction of Multiphase Energy Conversion Device			
2.2.2.05 WDMApp: High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas			
2.2.2.06 WarpX: Exascale Modeling of Advanced Particle Accelerators			
2.2.3 ---- Earth and Space Science Applications			
2.2.3.01 ExaStar: Exascale Models of Stellar Explosions: Quintessential Multi-physics Simulation			
2.2.3.02 ExaSky: Computing the Sky at Extreme Scales			
2.2.3.03 EQSIM: High Perf, Multidisciplinary Simulations for Regional Scale Earthquake Hazard/Risk Assmts			
2.2.3.04 Subsurface: Exascale Subsurface Simulator of Coupled Flow, Transport, Reactions & Mechanics			
2.2.3.05 E3SM-MMF: Cloud-resolving Climate Modeling of the Earth's Water Cycle			
2.2.4 ---- Data Analytics and Optimization Applications			1
2.2.4.01 Urban: Multiscale Coupled Urban Systems	1	1	
2.2.4.02 ExaSGD: Optimizing Stochastic Grid Dynamics at Exascale			
2.2.4.03 CANDLE: Exascale Deep Learning Enabled Precision Medicine for Cancer	1		
2.2.4.04 ExaBiome: Exascale Solutions for Microbiome Analysis		1	
2.2.4.05 ExaFEL: Data Analytics at the Exascale for Free Electron Lasers	1		1
2.2.5 ---- National Security Applications		1	1
2.2.5.01 ASC ATDM LANL Application		1	
2.2.5.02 ASC ATDM LLNL Application			
2.2.5.03 ASC ATDM SNL Application		1	
2.2.6 ---- Co-Design			
2.2.6.01 Proxy Applications	1	3	1
2.2.6.02 Application Assessment			
2.2.6.03 CODAR: Co-Design Center for Online Data Analysis and Reduction at the Exascale	2	2	
2.2.6.04 COPA: Co-Design Center for Particle Applications		1	
2.2.6.05 AMREX: Block-Structured AMR Co-Design Center		1	1
2.2.6.06 CEED: Center for Efficient Exascale Discretizations	1	1	1
2.2.6.07 ExaGraph: GraphEx Co-Design Center			
2.2.6.08 ExaLearn: CoDesign Center for Exascale Machine Learning Technologies			

BOF/BREAKOUT/PANEL SESSIONS

2.3 SOFTWARE TECHNOLOGY (ST)

10 WBS ST Projects involved

NUMBER OF PROJECTS

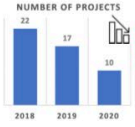

41% decrease w.r.t. 2019

13 WBS ST Sessions presented

NUMBER OF SESSIONS

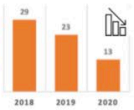

43% decrease w.r.t. 2019

Number of sessions with speakers/topics from WBS ST projects

WBS Name	# of Sessions 2018	# of Sessions 2019	# of Sessions 2020
2.3 Software Technology		3	2
2.3.1 ---- Programming Models & Runtimes			
2.3.1.01 PMR SDKs: Programming Models & Runtimes Software Development Kit			
2.3.1.02 LANL ATDM PMR Projects		1	
2.3.1.03 LANL ATDM PMR Projects			
2.3.1.04 SNL ATDM PMR Projects		3	
2.3.1.05 xGA: Global Arrays on Extreme Scale Architectures	1		
2.3.1.06 RAJA: Integrated Software Components for Managing Computation and		1	
2.3.1.07 Exascale MPI	2	1	
2.3.1.08 Legion: Enhancing and Hardening the Legion Programming System	1		
2.3.1.09 ParSEC: Distributed Tasking for Exascale	2		
2.3.1.10 Kokkos Support: ECP Applications Effective use of Kokkos to Achieve	1		
2.3.1.11 OMPI-X: Open MPI for Exascale	1	1	
2.3.1.12 Runtime System for Application-Level Power Steering on Exascale Systems			
2.3.1.13 SOLLVE: Scaling OpenMP with LLVm for Exascale Performance and	2	2	
2.3.1.14 UPC++ & GASNet: Lightweight Communication and Global Address Space Support for Exascale Applications	2		
2.3.1.15 Enhancing Qthreads for ECP Science and Energy Impact			
2.3.1.16 SICM: Simplified Interface to Complex Memory	1	1	
2.3.1.17 OMPI-X: Open MPI for Exascale			1
2.3.1.18 RAJA/Kokkos			1
2.3.1.19 Argo: Low-level resource management for the OS and runtime			1
2.3.2 ---- Development Tools			
2.3.2.01 Development Tools Software Development Kit			
2.3.2.02 LANL ATDM Development Tools Projects		1	
2.3.2.03 LLNL ATDM Development Tools Projects	2		
2.3.2.04 SNL ATDM Development Tools Projects			
2.3.2.05 Exascale Code Generation Toolkit			
2.3.2.06 EXA-PAPI: The Exascale Performance Application Programming Interface			1
2.3.2.07 Autotuning Compiler Technology for Cross-Architecture Transformation and		1	
2.3.2.08 Extending HPCToolkit to Measure and Analyze Code Performance on			
2.3.2.09 PROTEAS: Programming Toolchain for Emerging Architectures and Systems	1	1	

Number of sessions with speakers/topics from WBS ST projects

WBS Name	# of Sessions 2018	# of Sessions 2019	# of Sessions 2020
2.3.3 ---- Mathematical Libraries			
2.3.3.01 Extreme-scale Scientific xSDK for ECP	1		2
2.3.3.02 LANL ATDM Mathematical Libraries Projects			
2.3.3.03 LLNL ATDM Mathematical Libraries Projects			
2.3.3.04 SNL ATDM Mathematical Libraries Projects		1	
2.3.3.05 Enabling Time Integrators for Exascale Through SUNDIALS			
2.3.3.06 Preparing PETSc/TAO for Exascale			
2.3.3.07 STRUMPACK/SuperLU: Factorization Based Sparse Solvers and Preconditioners			
2.3.3.08 ForTrilinos: Sustainable Production Fortran Interoperability with Trilinos	1		
2.3.3.09 SLATE: Software for Linear Algebra Targeting at Exascale			
2.3.3.10 PEEKS: Production-ready, Exascale-Enabled, Krylov Solvers for Exascale			
2.3.3.11 ALExa: Accelerated Libraries for Exascale			
2.3.3.12 Enabling Exascale Simulations with SUNDIALS and hypre			1
2.3.4 ---- Data and Visualization			
2.3.4.01 Data and Visualization Software Development Kit			
2.3.4.02 LANL ATDM Data & Visualization Projects			
2.3.4.03 LLNL ATDM Data & Visualization Projects		1	
2.3.4.04 SNL ATDM Data & Visualization Projects			
2.3.4.05 VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart	2	1	
2.3.4.06 Ez: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for	1	1	
2.3.4.07 UCNIFYCR: A Checkpoint/Restart File System for Distributed Burst Buffers	2		
2.3.4.08 ExaHDF5: Delivering Efficient Parallel I/O on Exascale Computing Systems			
2.3.4.09 ADIOS Framework for Scientific Data on Exascale Systems	1		
2.3.4.10 DataLib: Data Libraries and Services Enabling Exascale Science			1
2.3.4.11 ZFP: Compressed Floating-Point Arrays	1		
2.3.4.12 ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis	1		
2.3.4.13 ECP/VTK-m	1		
2.3.4.14 VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart/SZ: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for Scientific Data			1
2.3.5 ---- Software Ecosystem and Delivery			
2.3.5.01 Software Ecosystem and Delivery Software Development Kit			
2.3.5.02 LANL ATDM Software Ecosystem & Delivery Projects			
2.3.5.03 LLNL ATDM Software Ecosystem & Delivery Projects		2	
2.3.5.04 SNL ATDM Software Ecosystem & Delivery Projects			
2.3.5.05 Argo: Operating System and Resource Management for Exascale	1		
2.3.5.06 FLANG LLVM Work w/NVIDIA	1	1	
2.3.5.07 ECP Software Stack Releases			
2.3.5.09 Software Packaging Technologies			2

- 2020 and 2019 data come directly from the information authors included in the proposal submission.
- 2018 data is an estimation based on project names and final sessions in the agenda.

BOF/BREAKOUT/PANEL SESSIONS

2.4 HARDWARE & INTEGRATION (HI)

16 WBS HI Projects involved

NUMBER OF PROJECTS

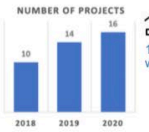

14% increase w.r.t. 2019

24 WBS HI Sessions presented

NUMBER OF SESSIONS

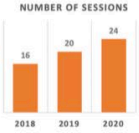

20% increase w.r.t. 2019

- 2020 and 2019 data come directly from the information authors included in the proposal submission.
- 2018 data is an estimation based on project names and final sessions in the agenda.

Number of sessions with speakers/topics from WBS HI projects

WBS Name	# of Sessions 2018	# of Sessions 2019	# of Sessions 2020
2.4 Hardware and Integration			1
2.4.1 ---- PathForward			
2.4.1.01 PathForward Vendor - Cray	1	2	1
2.4.1.02 PathForward Vendor - Intel	1	1	2
2.4.1.03 PathForward Vendor - HPE	1	1	1
2.4.1.04 PathForward Vendor - AMD	1	1	1
2.4.1.05 PathForward Vendor - IBM	1	2	1
2.4.1.06 PathForward Vendor - NVIDIA	1	1	1
2.4.1.07 Contract Management			
2.4.1.08 Testbed Support			
2.4.1.09 ATDM LLNL FastForward 2			
2.4.2 ---- Hardware Evaluation		1	
2.4.2.01 Memory Technologies		2	2
2.4.2.04 Compute Node Simulation			
2.4.2.03 High-Level Architecture and Abstract Machine Models			
2.4.2.04 Compute Node Simulation			
2.4.2.05 Analytical Modeling			
2.4.3 ---- Application Integration at Facilities	3	2	2
2.4.3.01 Aurora Application Integration			2
2.4.3.02 Frontier Application Integration			1
2.4.3.03 Pre-exascale System Application Integration		1	
2.4.3.04 Facility System Milestones			
2.4.4 ---- Software Deployment at Facilities		2	1
2.4.4.01 Continuous Integration Testing and Release	1		3
2.4.4.02 Integration and Deployment of Software			
2.4.5 ---- Facility Resource Utilization	5	1	
2.4.5.01 Facility Resource Utilization			
2.4.6 ---- Training and Productivity		1	
2.4.6.01 Training			2
2.4.6.02 Productivity	1	2	2
2.4.6.90 Management			1

NDA SESSIONS

- 8 NDA Sessions — 9.5h
- 6 PathForward Updates — 7.5h
  - PathForward Update – AMD
  - PathForward Update – Cray
  - PathForward Update – HPE
  - PathForward Update – IBM
  - PathForward Update – Intel
  - PathForward Update – NVIDIA
- 2 Breakout/BoF Sessions — 2h
  - Slingshot BoF
  - Aurora - Argonne's Exascale System

1293 Attendees of NDA Sessions

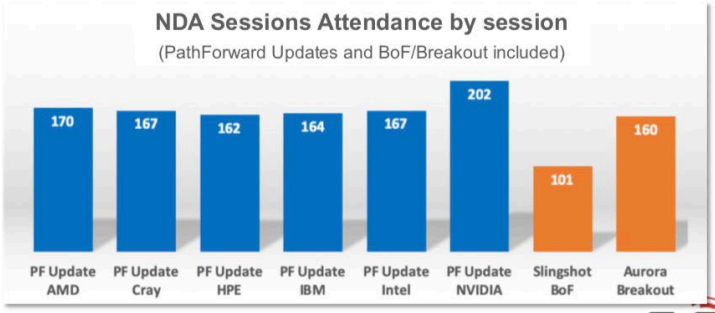

TUTORIALS

https://confluence.exascaleproject.org/display/2020ECPAM/Sessions

- 35 Tutorials
- 90.5 Hours
- 777* Tutorial Attendees

* Missing attendance # from 4 Tutorials

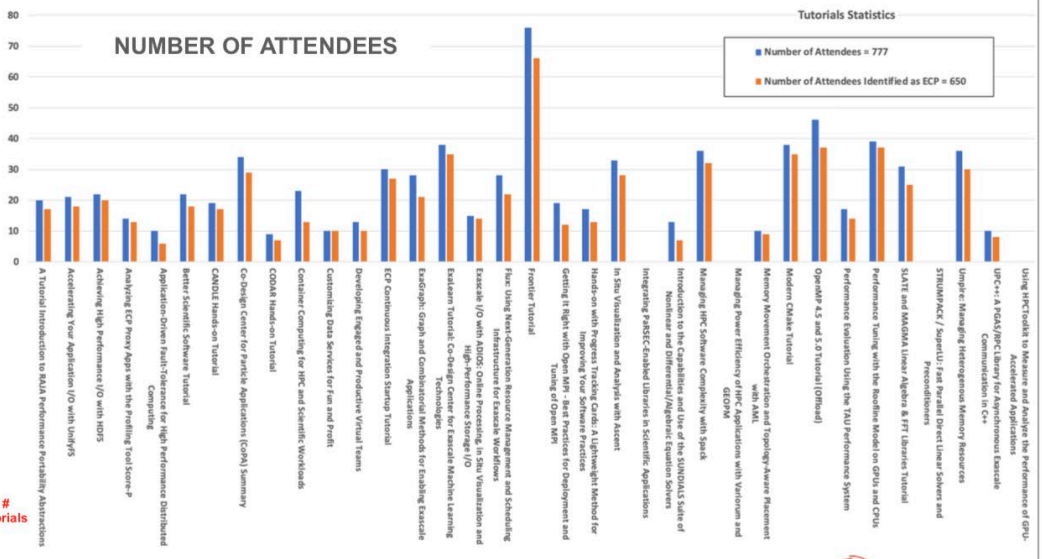

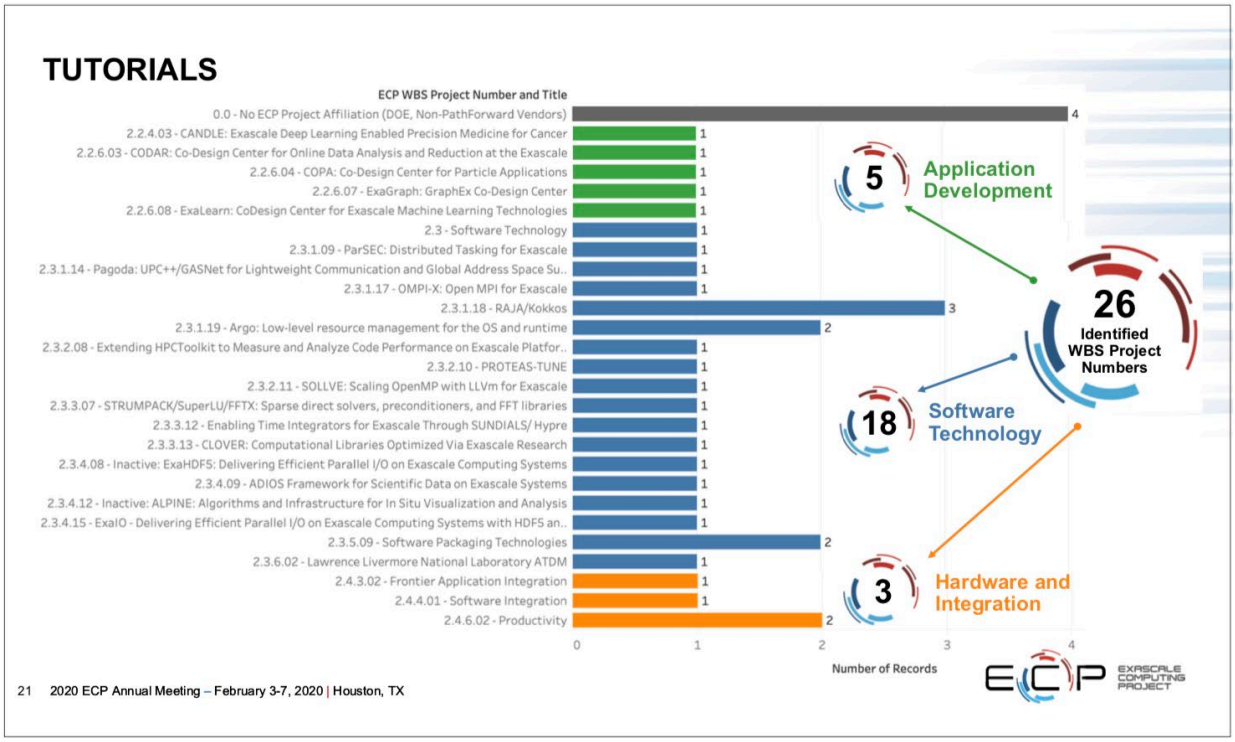

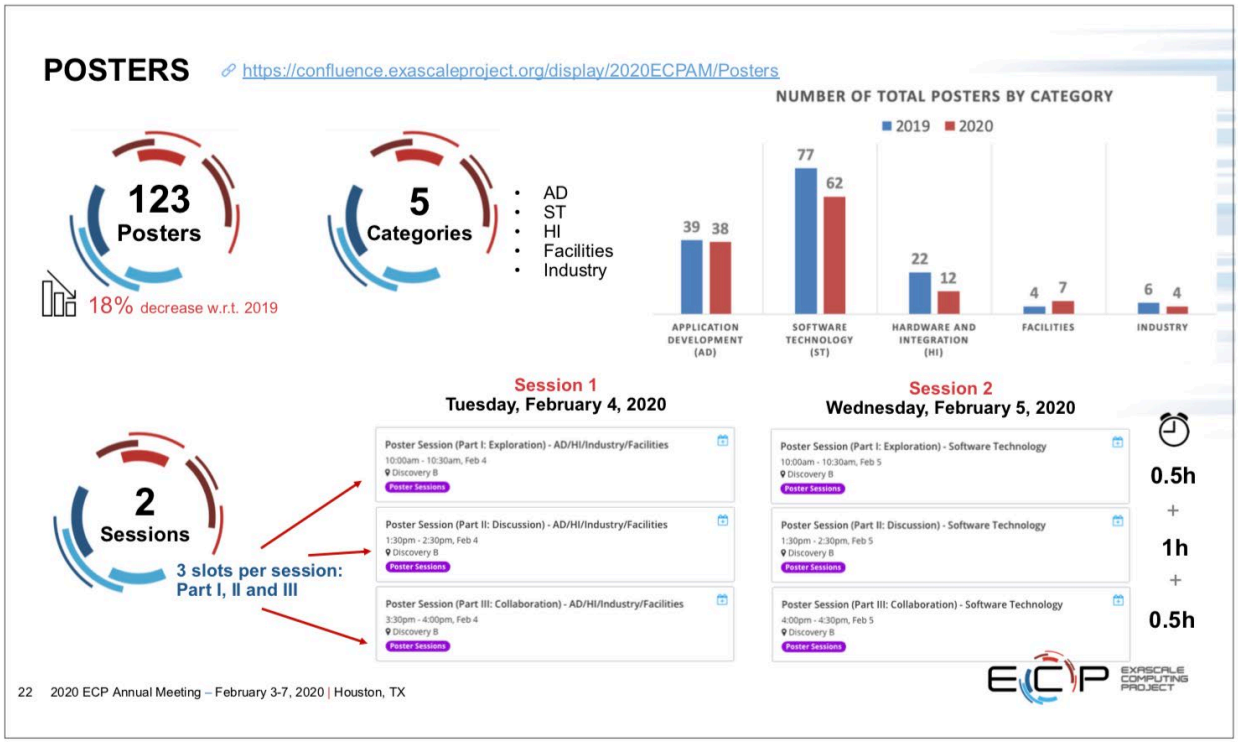

ECP-O

Posters Droid Helper!

New this year!

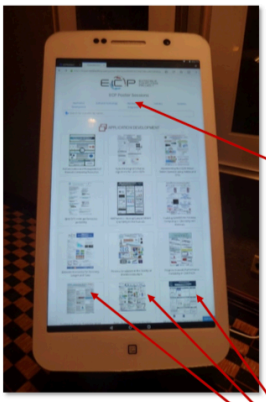

Sample of posters displayed in giant touch-screen device: ECP-O

GOALS

- Complement the poster sessions in paper/printed format with an interactive, compact and easy-to-use platform displaying all the posters of the meeting.
- Serve as a resource of additional space in a strategic location for discussions and impromptu pop-up poster presentations.

HARDWARE FEATURES

- Giant touch-screen display with zoom in/out capability.
- Device available 24/7 for all attendees.

SOFTWARE FEATURES

- New web page developed with PDF files of posters available.
- Posters organized alphabetically and displayed by categories: AD, ST, HI, Facilities, Industry.
- Search engine available to find posters based on keywords, titles or other content.
- Small thumbnails of posters added to the titles to facilitate search and encourage device/human interaction.

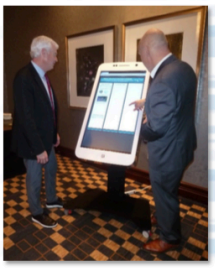

ECP Director (Dr. Doug Kothe) and DOE Office of Science Director (Dr. Chris Fall) using the interactive display

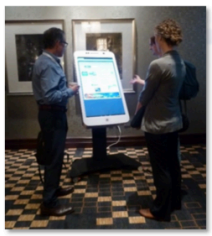

NERSC Team discussing the content of a poster

ECP-O

Samples of categories and posters display of the web page developed with posters content for ECP-O

https://ecpannualmeeting.com/poster-interface2389473231 4e23hreu823rd/

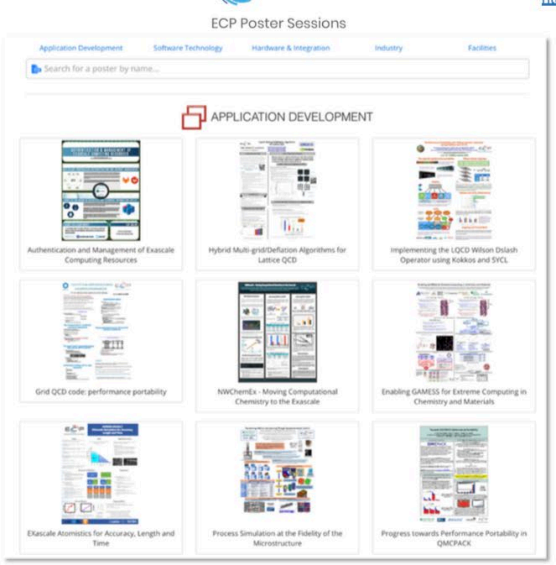

DISCLAIMER

The data contained in the following section comes from the information attendees entered during the registration process.

The project numbers identified by attendees were not validated: however, we believe these graphs provide a good representation of the audience at the 2020 Annual Meeting and we recommend their use for future planning purposes.

2020 ECP Annual Meeting

Yearly Meeting of ECP members, leadership, vendors and stakeholders.

7% increase w.r.t. 2019

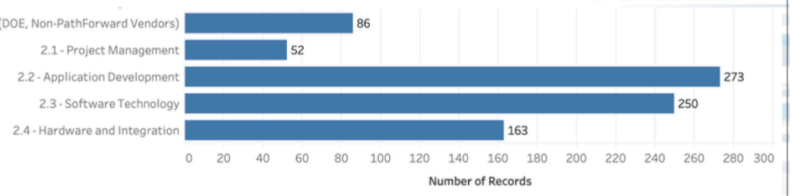

- This January was the 4th ECP Annual Meeting
  - 4 days of deep technical sessions, tutorials, posters, plenaries
- Provides a forum for ECP participants to meet/talk directly with other project members across ECP in order to:
  - Explore collaborations
  - Coordinate plans and milestones
  - Learn about resources, facilities common tools and libraries
- Industry participants
  - Allowed three people from each member's company
  - Plus two each from oil and gas

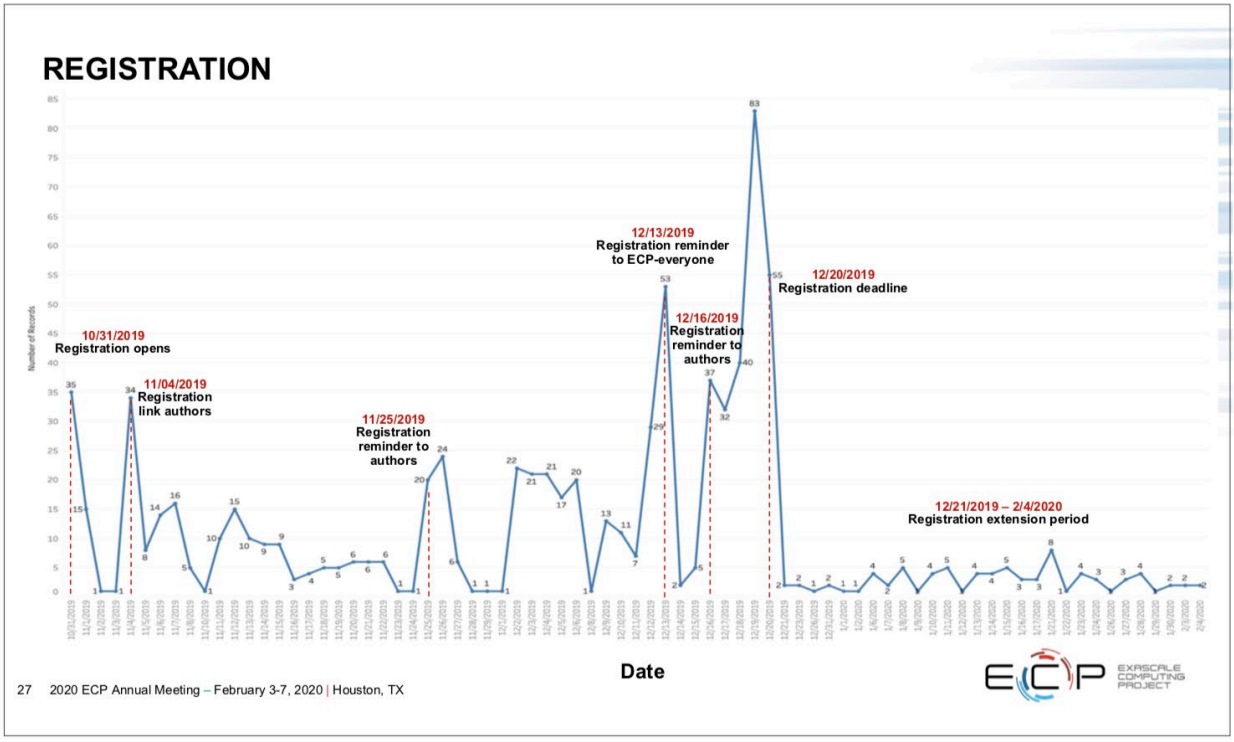

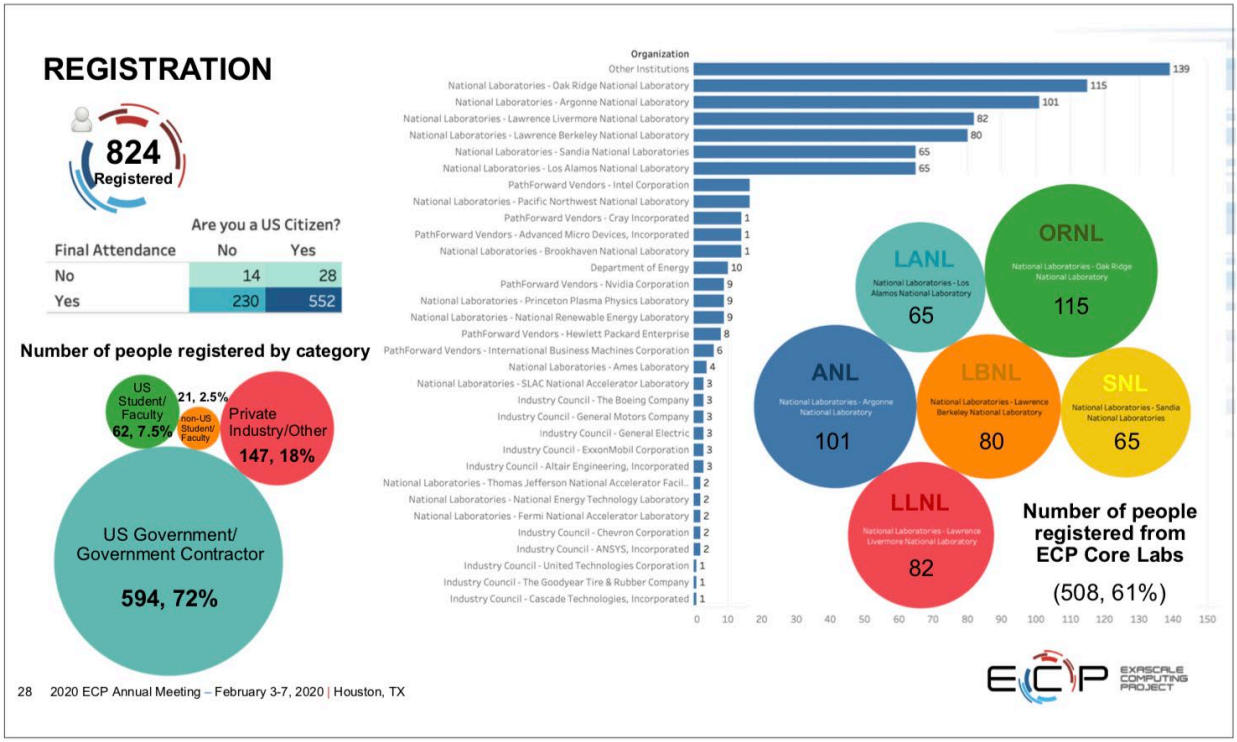

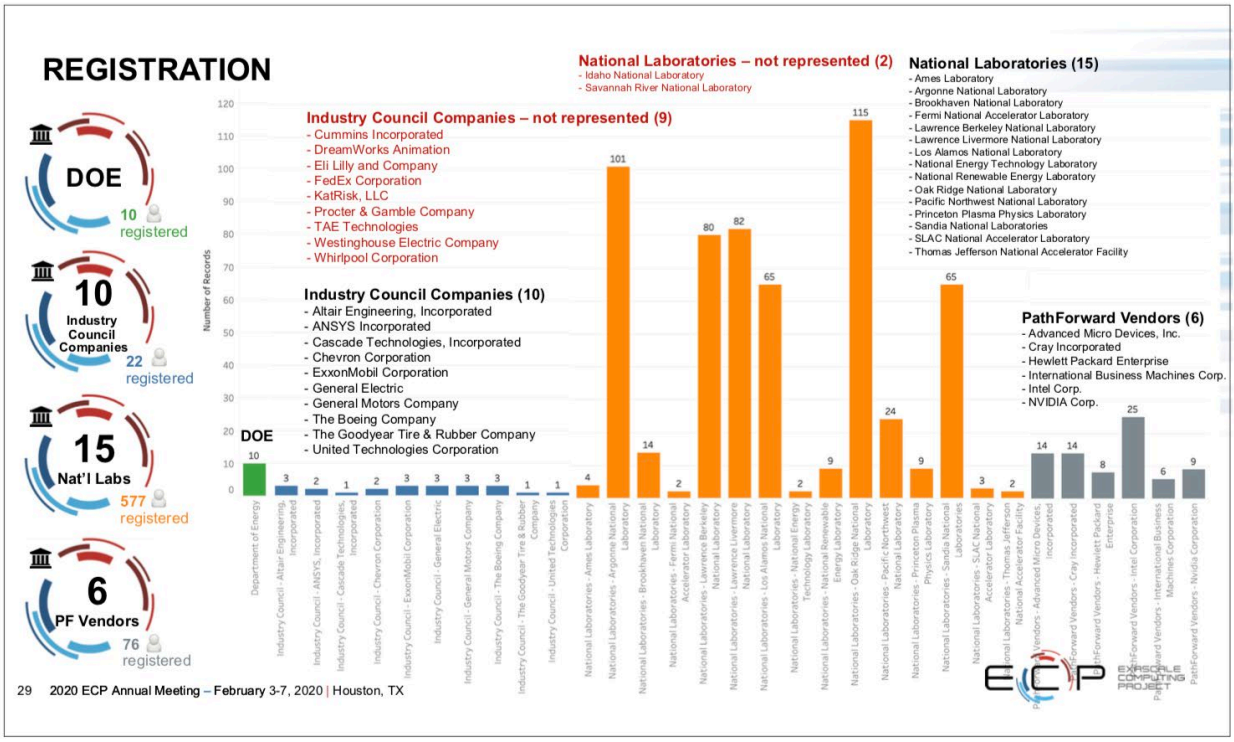

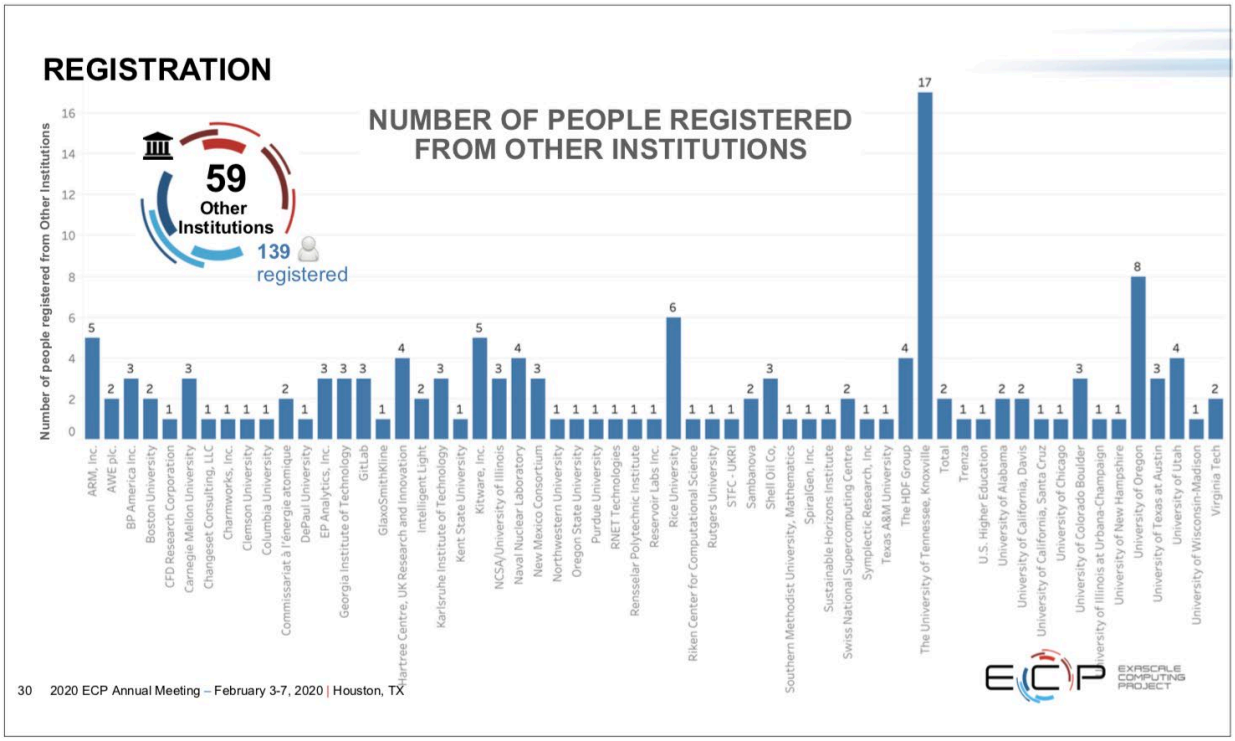

REGISTRATION

12 WBS 2.1 Projects Identified

51 WBS 2.1 people

NUMBER OF PEOPLE REGISTERED
2.1 PROJECT MANAGEMENT

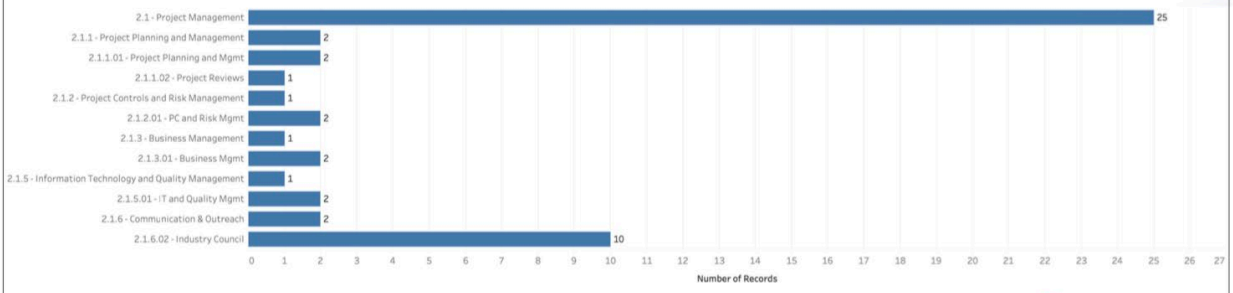

REGISTRATION NUMBER OF PEOPLE REGISTERED – 2.2 APPLICATION DEVELOPMENT

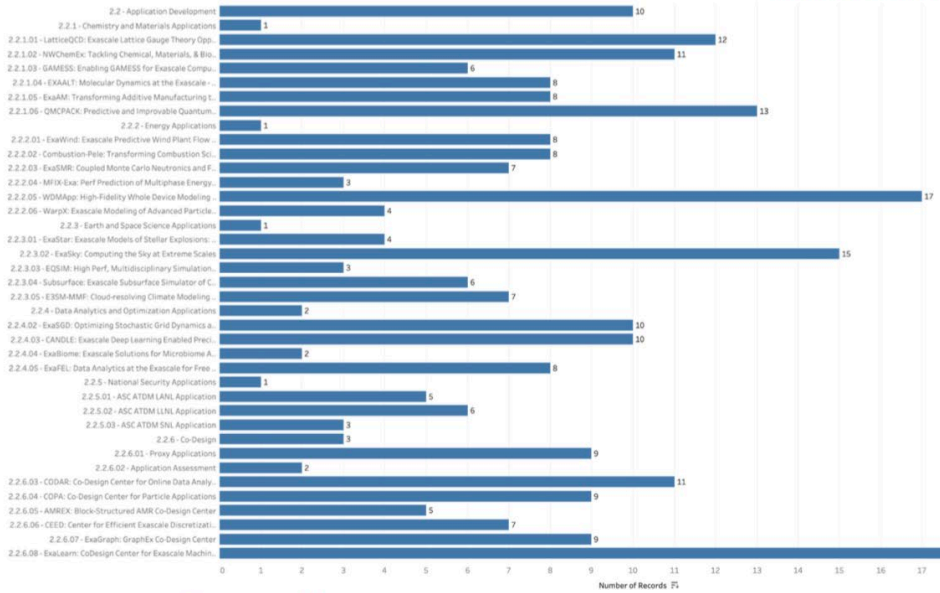

39 WBS 2.2 AD Projects Identified

273 WBS 2.2 people

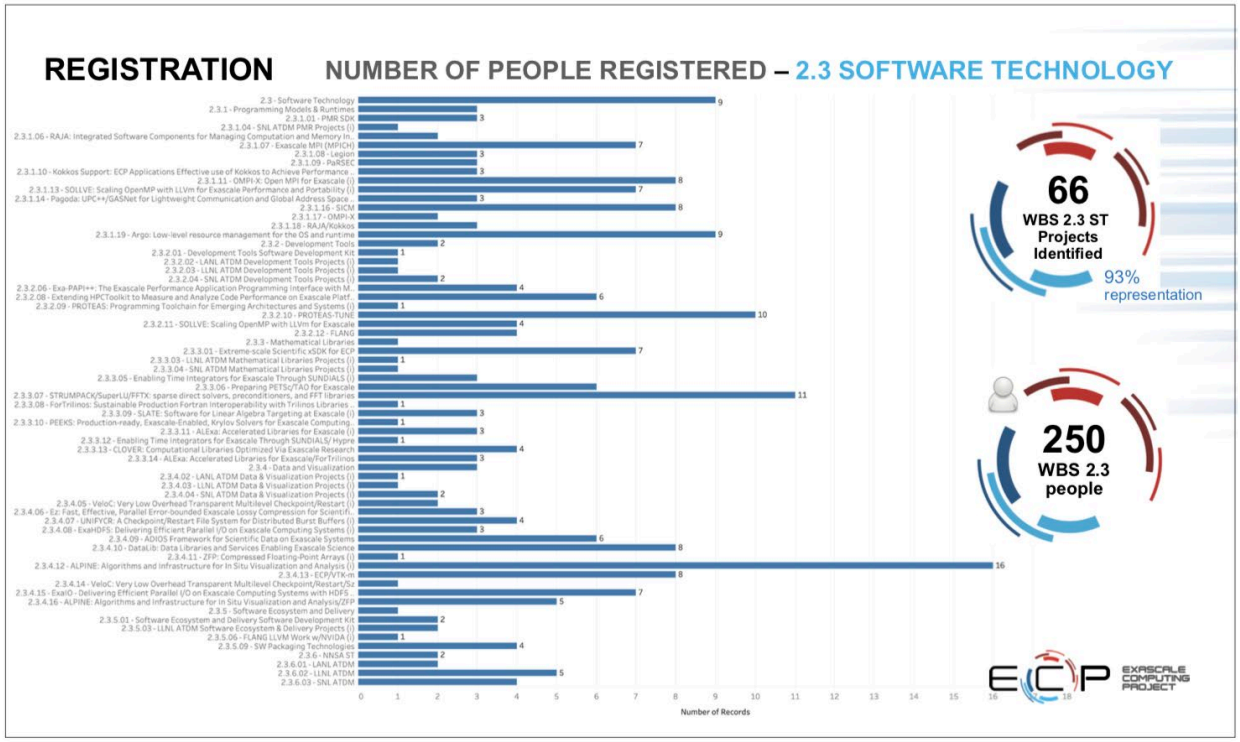

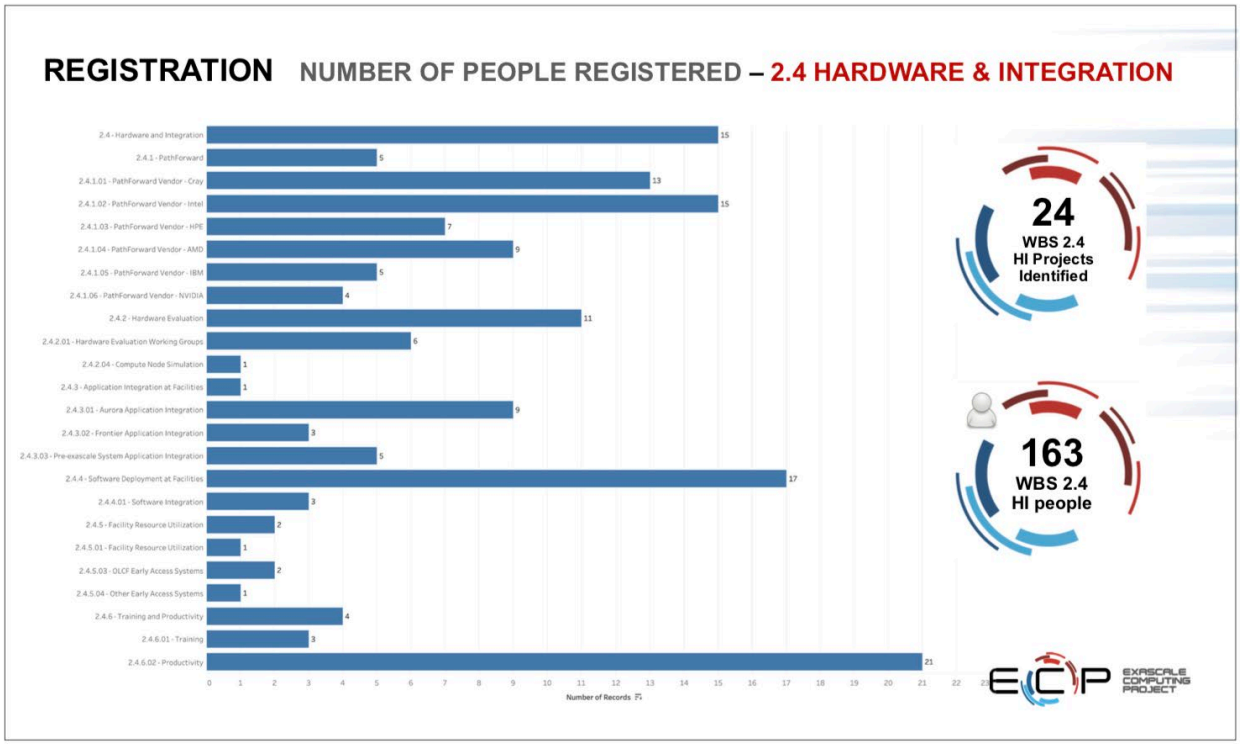

FINAL ATTENDANCE

DISCLAIMER

The data contained in the following section comes from the information gathered from Whova and other sources.

Some attendees did not check in with Whova: however, we believe these graphs provide a good representation of the audience at the 2020 Annual Meeting and we recommend their use for future planning purposes.

ECP EXASCALE COMPUTING PROJECT

FINAL ATTENDANCE

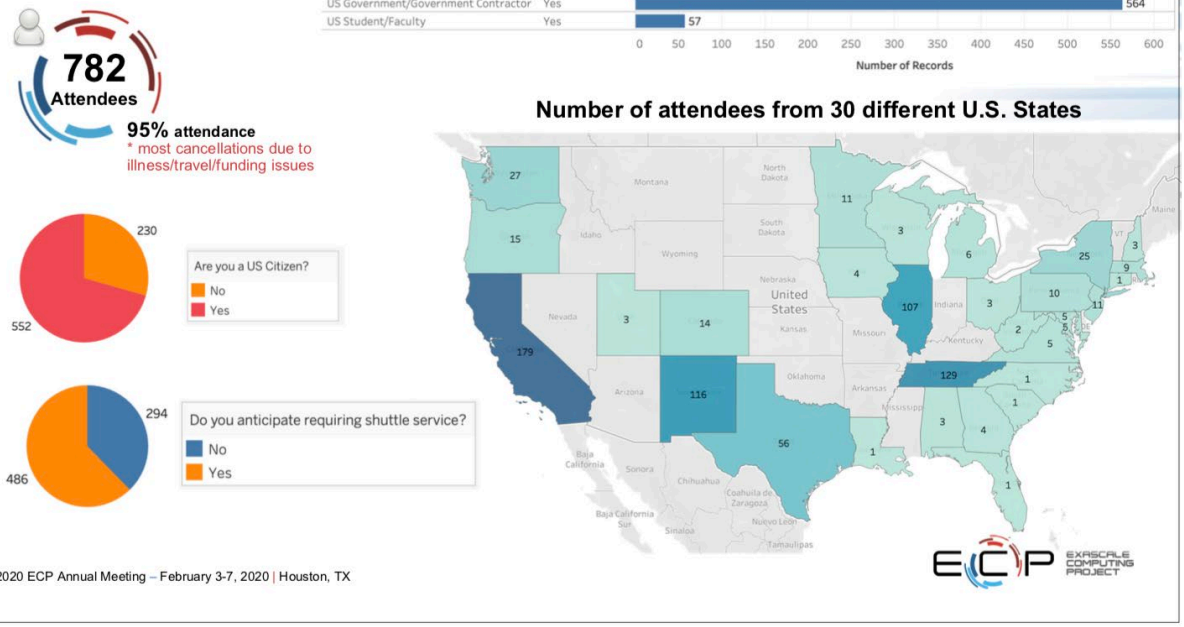

ECP EXASCALE COMPUTING PROJECT

FINAL ATTENDANCE – INDUSTRY COUNCIL (AND "FRIENDS")

13 Industry Companies Represented

- **10** Industry Council Companies
- **3** Additional Oil & Gas Firms

27 Industry Attendees

- **19** Attendees from Industry Council Companies
  - 2 Industry Council Members
  - 17 Senior Technical Staff

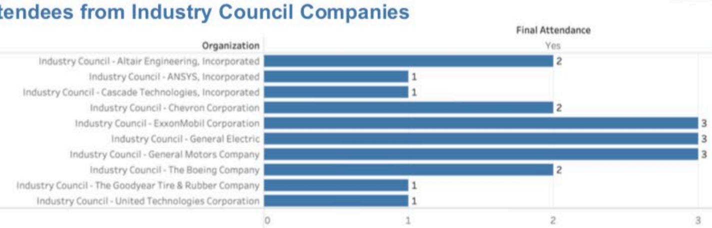

- **8** Non-member Oil & Gas Firms Representatives

INDUSTRY COUNCIL ACTIVITIES AT 2020 ANNUAL MEETING

- **Meet & Greet Reception** Tuesday evening nicely attended.
- **42** Attendees **ECP 201: Lessons Learned in Preparing Software Technology and Applications for Exascale Systems**
- Doug Kothe meeting with Dave Kepczynski, CIO GE Research/Chair ECP Industry Council and Prith Banerjee, CTO ANSYS.
- **Luncheon Plenary** ISV Panel with 4 member companies:
  - GE (moderator), Ansys, Altair, Cascade.
- 2 meetings with Mike Heroux and ST speakers presenting at the Deep Dive workshop to discuss presentation content.
- Meeting with **Hartree Centre (UK)** regarding ALCF and OLCF industry programs and ECP Industry Council.

SERVICES

- WEBSITES
- LODGING
- SHUTTLE SERVICE
- SIDE MEETINGS
- WEB/MOBILE APP
- RIBBONS FOR BADGES

WEBSITE

Changes in the website designed in 2019 with added functionalities:

- New home page image
- New Speakers page
- Hilton added as meeting hotel
- New FAQ page with interactive drop down menu

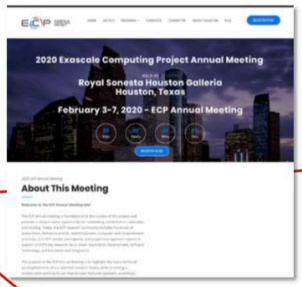

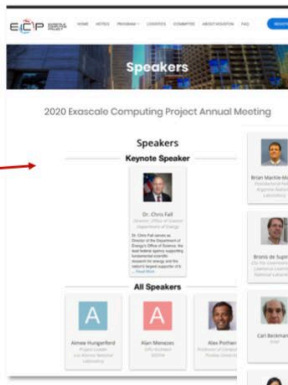

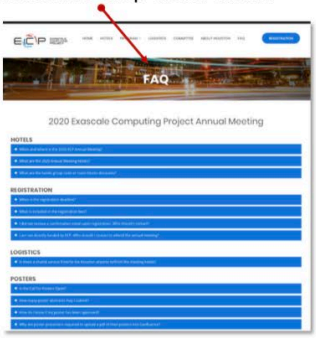

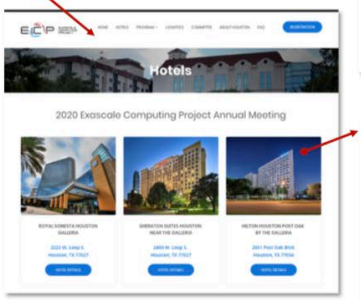

With pop-up window!

PUBLIC WEBSITE

https://www.ecpannualmeeting.com/overview.php

Material available publicly in the **Program Overview** page:

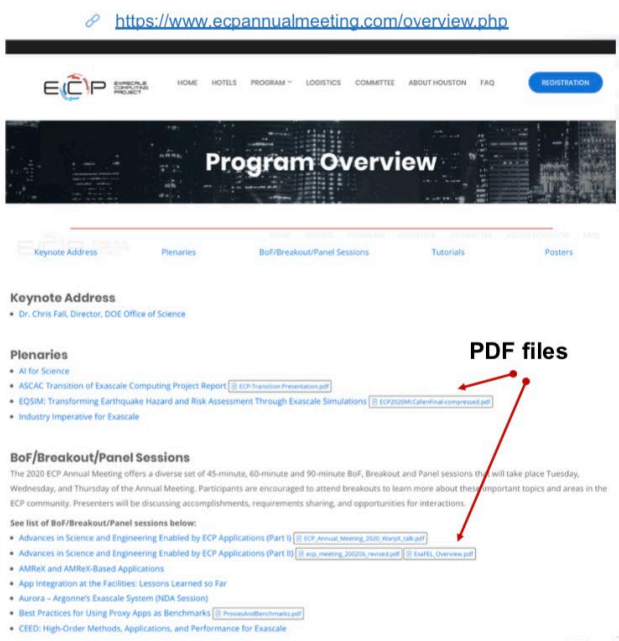

PDF files

CONFLUENCE WEBSITE

https://confluence.exascaleproject.org/display/2020ECPAM

Material available for the ECP community in the **2020 Annual Meeting** page:

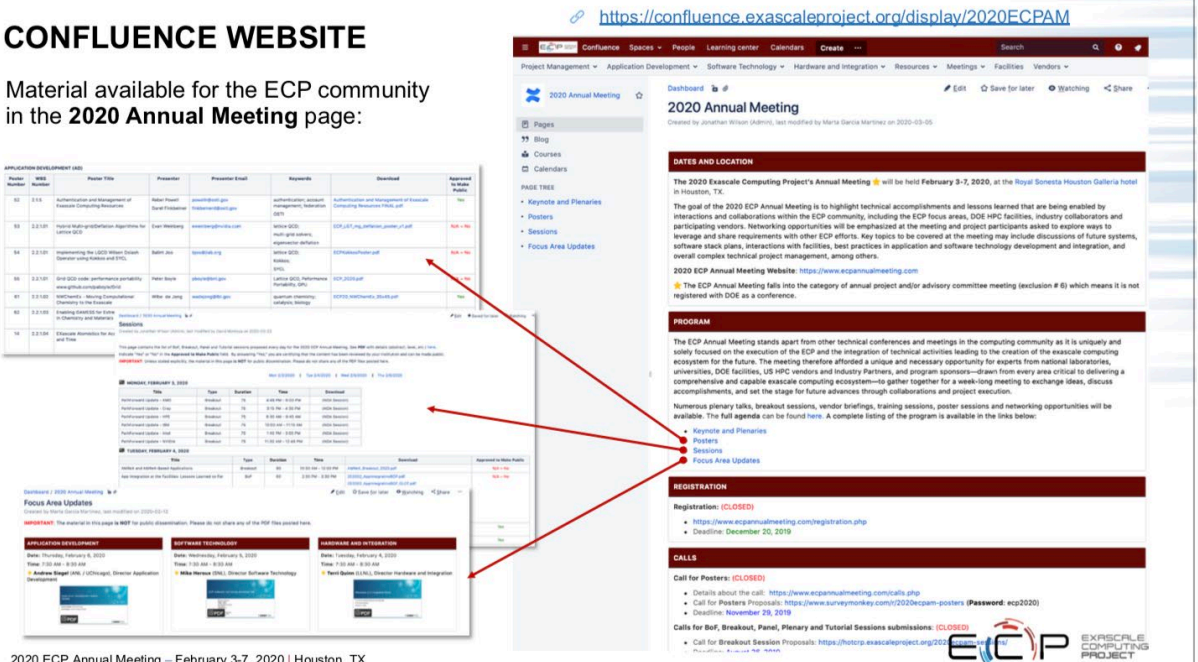

LODGING

1786 Total Rooms Contracted → 6% → 1669 Total Rooms Reserved

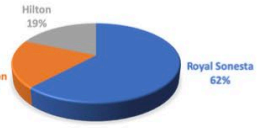

~2/3 attendees lodged at meeting location (Royal Sonesta)

Royal Sonesta Houston Galleria

1236 Rooms Contracted
1035 Rooms Reserved

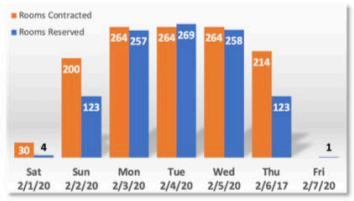

Sheraton Suites Houston Near The Galleria

275 Rooms Contracted
324 Rooms Reserved

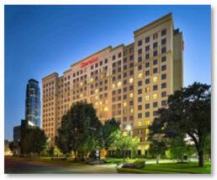

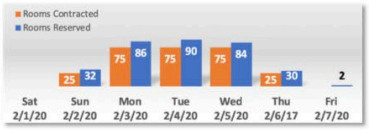

Hilton Houston Post Oak By The Galleria

275 Rooms Contracted
310 Rooms Reserved

New Meeting Hotel!

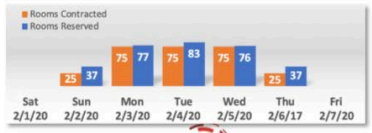

SHUTTLE SERVICE

935 Total Reservations
+27% w.r.t. 2019

488 People Served
63% of attendees!!
+30% w.r.t. 2019

- Round-trip: 447 (92%)
- Arrival only: 18
- Departure only: 23

Hotels - ARRIVAL

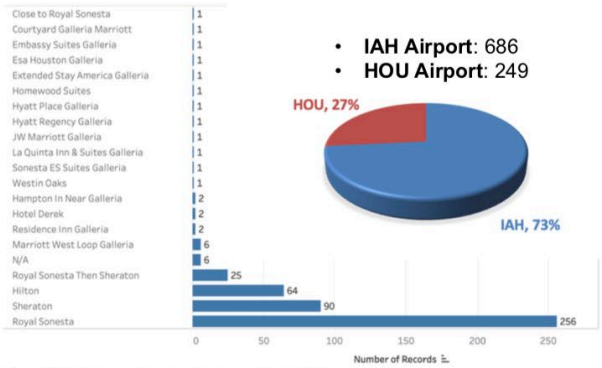

- IAH Airport: 686
- HOU Airport: 249

Shuttle Service from Houston Airports - ARRIVAL

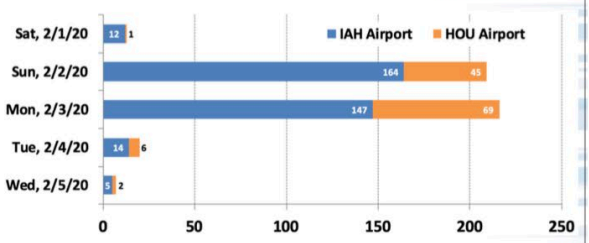

Shuttle Service to Houston Airports - DEPARTURE

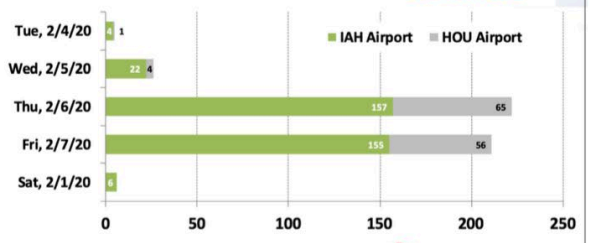

SIDE MEETINGS

Side meetings at the annual meeting are a **great added value**.

90% increase w.r.t. 2019

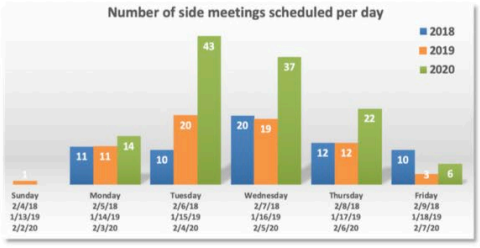

42% increase w.r.t. 2019

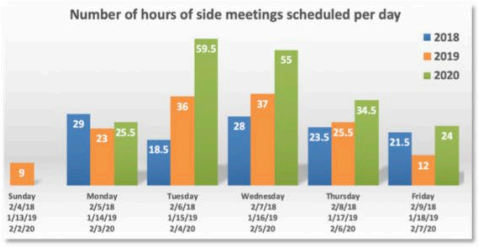

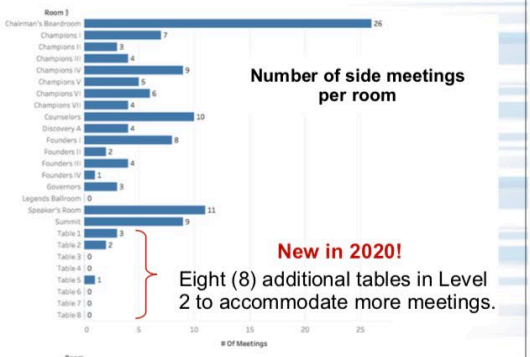

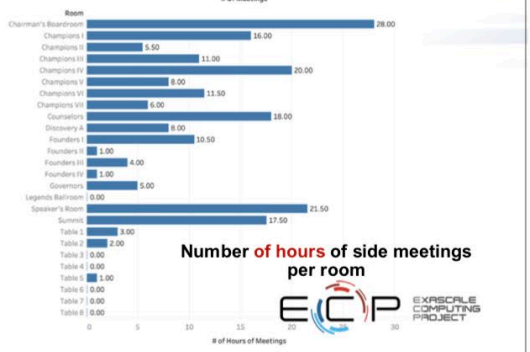

WEB/MOBILE APP

WHOVA - Statistics

- **Total Profile Views**: 2577
- **Total Messages Sent/Received**: 176
- **Total Business Cards / Contact Info Exchanged**: 10
- **Total Number of Downloads**: 651
- **Device Breakdown**: iOS (59%), Android (41%)
- **Total Agenda Webpage Views**: 9008
- **Photos Shared**: 97
- **Number of Announcements Sent**: 3
- **Community Messages**: 481
- **Active attendees in community**: 74

BREAKDOWN BY AFFILIATION
Total number of affiliations: 193

TOP 5 AFFILIATIONS:
1. Argonne National Laboratory (106)
2. Oak Ridge National Laboratory (105)
3. Lawrence Berkeley National Laboratory (80)
4. Lawrence Livermore National Laboratory (75)
5. Los Alamos National Laboratory (63)

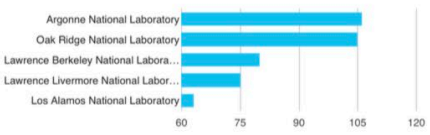

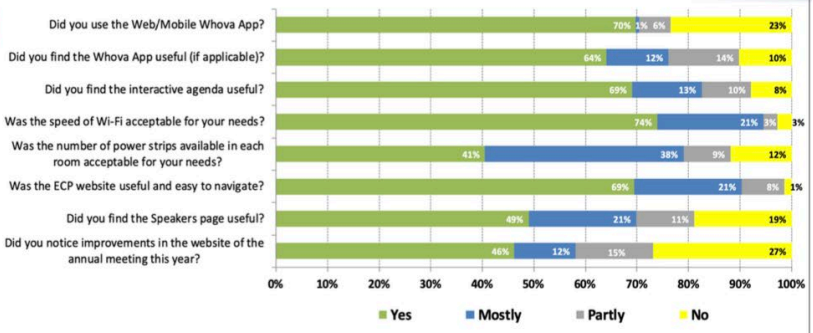

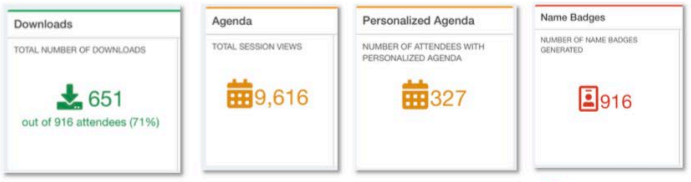

RIBBONS FOR BADGES

New this year!

POSTER PRESENTER Ribbon - Red

SPEAKER Ribbon - Royal

AUTHOR Ribbon - Navy

STAFF Ribbon - Black

- To foster interactions and discussions about posters, sessions or questions for staff, this year we provided four types of ribbons to display on the ECP badges.
- The ribbons were available for pick up anytime at the ECP registration desk.
- Colors followed ECP logo palette of blues, red and black.
- Minimal cost for the program due to high quantity orders.
- Ribbons not used are available for future events.
- Low storage requirements.
- Easy organization in ribbons container.

FINAL FEEDBACK FORM

SurveyMonkey

https://www.surveymonkey.com/r/2020ECPAM_Final_Feedback_Form

FINAL FEEDBACK FORM

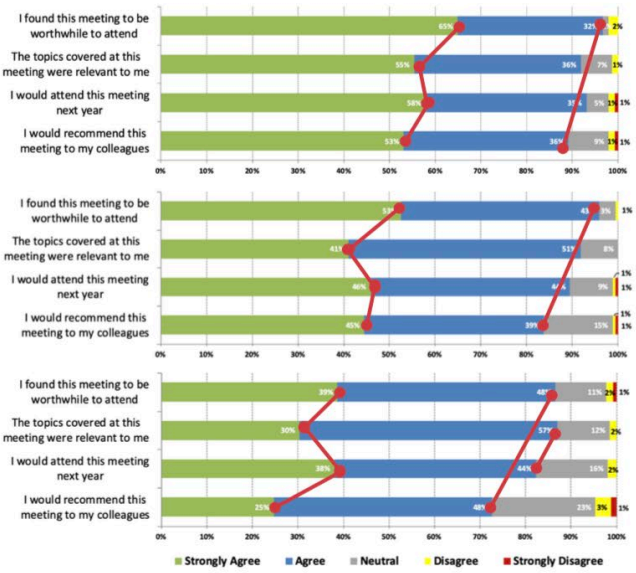

2020
151 Respondents

2019
173 Respondents

2018
261 Respondents

IMPROVEMENTS

Increase in the number of "**Strongly Agree**" in all questions.

Increase in the number of "**Strongly Agree**" and "**Agree**" for recommendation of this meeting to colleagues.

"I would recommend this meeting to my colleagues"

Strongly Agree + **Agree** = 89% (2020)
84% (2019)
73% (2018)

6% increase w.r.t. 2019
22% increase w.r.t. 2018

Increase in the number of "**Strongly Agree**" and "**Agree**" for next year attendance.

"I would attend this meeting next year"

Strongly Agree + **Agree** = 93% (2020)
90% (2019)
82% (2018)

3% increase w.r.t. 2019
13% increase w.r.t. 2018

IMPROVEMENTS IN 2020 FROM 2019 FEEDBACK

The committee took action and addressed most of the **problems mentioned in the 2019 survey**. See below:

GENERAL

- Publish the agenda much earlier to allow potential attendees to make travel plans in accordance.
- Focus all sessions in one hotel. Having a second hotel for some sessions was a disturbance for attendees and a clear reduction in attendance for presenters.
- Go paperless. The welcome package is a duplicate of what is online or even an old version of some documents due to last-minute updates. Focus on improving how people have access to that information and avoid admin time to print, prepare and ship >750 folders.
- Convenience of having meeting and hotel at the same venue.

LOGISTICS

- Improve Side Meetings scheduling and management to avoid nuisance/disturbance to meeting sessions and avoid double-booking.
- Clarify communications about NDA sessions attendance and announcements.
- Improve management of temperature in meeting rooms.
- More space/tables for impromptu meetings and side-chats.

IT

- Improve Wi-Fi connection. In 2019 we had issues with intermittency of service and quality during the first days of the meeting.
- Better communication of where material is located (PDF of slides, maps, etc.).
- Target Web/Mobile app announcements.

A/V

- Officially schedule time for speaker to test their laptops. We had some problems with some laptops in some rooms that delayed the presentations.
- Provide more dongles.
- Provide optional lavalier microphones instead of fixed ones. That allows more interaction with the audience while moving around in the room.

2020 SURVEY RESULTS – OPPORTUNITIES FOR IMPROVEMENT

GENERAL

- Meeting Hotel at location not pedestrian friendly, which impacted attendees during non-meeting lunch time slots.
- Snacks during breaks would be nice.
- Avoid talks during working meals to allow attendees to schedule their own discussions and collaboration activities.
  "This really cuts down in collaboration time and informal time."
- Consider other timezone or start/end day later since it's hard for West Coast attendees to start the day at 6:00 AM their time.
- Work harder to reduce topic conflicts in agenda.

LOGISTICS

- Bigger poster room.
- Clarify communications about NDA sessions attendance and announcements.
- Some issues with departure times or shuttle notifications:
  - Missing announcements
  - Personal information shared with others
  - Inconvenience to do not have schedule in advance
  - E.g. flight at 3:25 PM and pick up time at 10:30 AM!
- Food:
  - Avoid cleaning tables so early and give more time to pick up food.
  - Healthier desserts, more fruits instead of cakes.

IT

- Provide additional training about features of the apps (e.g. create your own agenda, connect with other audiences, updates and announcements).
- Avoid no work-related conversations or photos that distract from the focus of the meeting.
- Too many emails from the Whova App.
- Some issues with Wifi:
  - Speed for some attendees
  - Blocked IPs
  - Connection dropping
- More power strips, 1 per row would be nice and ensure all of them have power.

A/V

- Improve A/V equipment resolution to 1080 instead of 720 (if equipment allows).
- Potentially record some of the talks or sessions.
- Balance audio better for some talks.

2020 SURVEY RESULTS – OPPORTUNITIES FOR IMPROVEMENT

FOR FUTURE MEETINGS I'D LIKE TO SEE MORE OF THE FOLLOWING

- Badges on Day-1
- Q/A after meetings
- Time and space for informal/unstructured and side meetings
- Summaries of AD projects
- More lessons learned and examples of benefits of app
- New topics (some seemed repetition of 2019 content)
- Protection of personal information
- Dessert! ;-)

FOR FUTURE MEETINGS I'D LIKE TO SEE LESS OF THE FOLLOWING

- Overlap of major tutorial and meetings
- Sunday travel
- Talks during meals
- Spam from non-work issues in Whova app
- One BOF or tutorial per project would be good than multiple of them per project.

NDA SESSIONS

- Earlier notification from vendors about approval
- Less vendor advertisement

KEYNOTE/PLENARIES SESSIONS

- Make sure topics are interesting for audience
- Have those presentations outside of meal times
- Clarify actual start time

BREAKOUT SESSIONS

- Avoid simultaneous sessions
- Ensure room size is big enough for audience

TUTORIALS

- Less overlap of technical and tutorials

POSTERS

- Poster sessions too crowded
- Larger poster room and more space between aisles
- More clarification about days and poster, maybe interleave ST, AD and HI posters

2020 SURVEY RESULTS – MAIN TAKE AWAYS

- ✓ Overall consensus that the meeting was an improvement from last year.
- ✓ Very high attendance to NDA Sessions, better to have them all in 1 day.
- ✓ The meeting was considered an excellent meeting overall.
- ✓ Collaborations, side meetings, poster sessions and social events have a big impact on the value of the meeting.
- ✗ There is still potential for improvement in some areas.
  - Location
  - Poster Sessions
  - Keep improving content of session
  - Work harder to avoid overlapping of topics

SOCIAL MEDIA

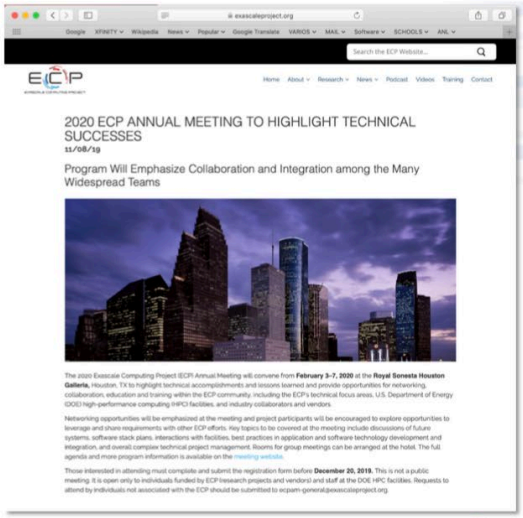

SOCIAL MEDIA

Tweet Announcements

SOCIAL MEDIA

Whova – photos shared from the App

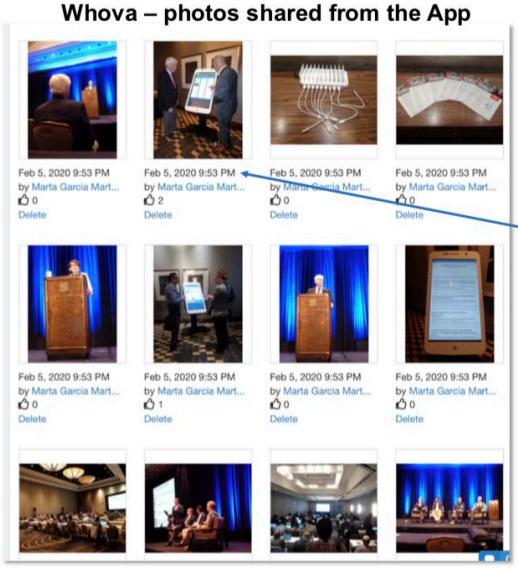

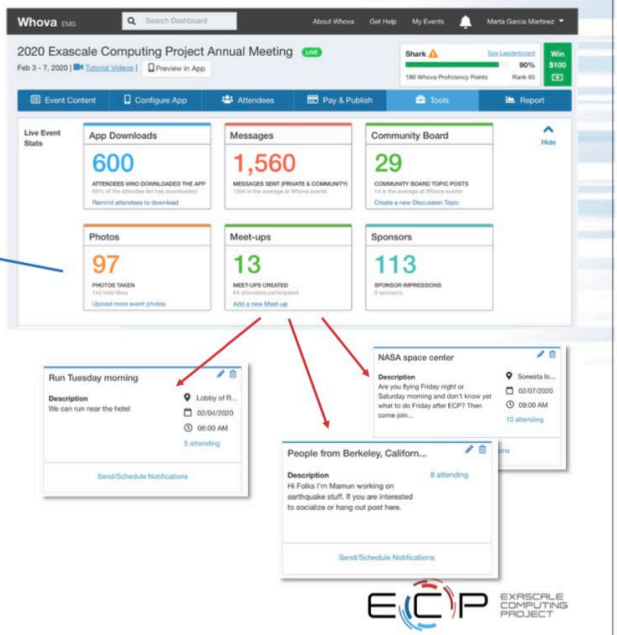

ACKNOWLEDMENTS

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two DOE organizations—the Office of Science and the National Nuclear Security Administration—responsible for the planning and preparation of a capable exascale ecosystem—including software, applications, hardware, advanced system engineering, and early testbed platforms—to support the nation's exascale computing imperative.

ECP EXASCALE COMPUTING PROJECT

ECP EXASCALE COMPUTING PROJECT

Questions?

Marta García – mgarcia@anl.gov

+1 (630) 252-0091