E4S: An Introduction to the Extreme-scale Scientific Software Stack



Michael A. Heroux Director of Software Technology, US Exascale Computing Project Senior Scientist, Center for Computing Research, Sandia National Laboratories Scientist in Residence, St. John's University, MN

September 22, 2019





09:00am-10:00am **E4S: An Introduction to the Extreme-scale Scientific Software Stack** Michael Heroux, Sandia National Laboratories

10:00am-10:30am **Supercontainers for HPC** Andrew J. Younge, Sandia National Laboratories

10:30am-11:00am Break

11:00am-11:30am **OpenHPC Update and Future Contributions** Robert Wisniewski, Intel

11:30am-12:00pm **An Overview of the ECP Software Ecosystem** Todd Munson, Argonne National Laboratory

12:00pm-12:30pm Hartree Centre Experience in Extreme Scale Computing Addressing Industrial, Societal and Scientific Challenges Vassil Alexandrov, Hartree Centre-STFC

12:30pm-02:00pm Break

02:00pm-02:30pm Experiences with Extreme-scale Scientific Software Stack (E4S) and Spack on SDSC systems Mahidhar Tatineni, SDSC

02:30pm-03:00pm **Kokkos: Performance Portability for C++ Applications** Christian Trott, Sandia National Laboratories

03:00pm-03:30pm **xSDK: a Community of Diverse Numerical HPC Software Packages** Ulrike Meier Yang, Lawrence Livermore National Laboratory

03:30pm-04:00pm Break

04:00pm-04:30pm **Toward interoperable and flexible scientific computing libraries: Lessons learned from xSDK** Keita Teranishi, Sandia National Laboratories

04:30pm-05:00pm **Data Management and Visualization Approaches for the U.S. Exascale Program** James Ahrens, Los Alamos National Laboratory

05:00pm-05:30pm **Leveraging and Expanding the Capabilities of the LLVM Compiler Infrastructure for Exascale Computing** Patrick McCormick, Los Alamos National Laboratory

05:30pm-06:00pm **Experiences in Designing, Developing, Packaging, and Deploying the MVAPICH2 Libraries** Hari Subramoni, Ohio State University

E4S Forum Agenda

E4S Overview





Software Technology Ecosystem



ECP ST Open Product Integration Architecture

Extreme-scale Scientific Software Stack (E4S)

- <u>E4S</u>: A Spack-based distribution of ECP ST and related and dependent software tested for interoperability and portability to multiple architectures
- Provides distinction between SDK usability / general quality / community and deployment / testing goals
- Will leverage and enhance SDK interoperability thrust
- Oct: E4S 0.1 24 full, 24 partial release products
- Jan: E4S 0.2 <u>37 full</u>, 10 partial release products



<u>e4s.io</u>

Lead: Sameer Shende (U Oregon)



Sample of E4S Release and Installed Packages

- Adios
- Bolt
- Caliper
- Darshan
- Gasnet
- GEOPM
- GlobalArrays
- Gotcha
- HDF5
- HPCToolkit
- Hypre
- Jupyter
- Kokkos
- Legion

- Libquo
- Magma
- MFEM
- MPICH
- OpenMPI
- PAPI
- Papyrus
- Parallel netCDF
- ParaView
- PETSc/TAO
- Program
 Database
 - Toolkit (PDT)

- Qthreads
- Raja
- SCR
- Spack
- Strumpack
- Sundials
- SuperLU
- Swift/T
- SZ
- Tasmanian
- TAU
- Trilinos
- VTKm
- Umpire

				5. pm				
linux-centos7	x86_64 / acce4.8.	5						
autoconf#2,69	cudal 9.1.85	gmpi6.1.2	kokkos#2.83.00	libxm12#2.9.4	hpich#1.2.1	openssl#1.0.2n	readline#7.8	
automake#1.15.1	flex#2.6.4	help2man#1.47.4	libpciaccess@0.13.5	m4#1.4.18	ncurses=6.0	papi 85.5.1	tor=1,29	
bisonP3.0.4	gcc#7.3.8	hwloc#1.11.9	libsigsegvez.11	magna 2.4.8	numact1=2.0	11 pdt#5.75	util-mocros#1.19.1	
bzip2+1.0.6	gdbm 1.14.1	hwlock2.0.1	libtoole2.4.6	mpc#1.1.0	openblas 0.7	2.20 per105.24.1	xz65.2.3	
cmake=3.11.1	gettext 0, 19.2.1	is100.19	libunwind 1.1	mpfr#4.8.1	openmpi#3.4	1 pkgconf#1.4.0	zlib#1.2.11	
linux-centos7-	x86_64 / gcc#7.3							
adios 1.13.1	freetype 2	7.4	json-c=0,13,1	libxfixes 5.	0.2	popi85,5.1	py-mccabe=0.6.1	sqlite@3.22.0
odlbx 0.8.0	gasnetel.3	3,18	kbprotor1.0.7	libxm1242.9.	4	popyrusidevelop	py-mock#2,0,0	stci 0.7.3
adlbx#8.8.8	gasnet 1.3	1.9	kokkosee, 01-00	libxshmfence	193 E.	paraview 5.4.1	py-mpi4py#3.0.0	strumpack13.1.1
ant=1.9.9	gdb 8 8.1		kvtree 1 0.2	libxt 1.1.5		parmetis ⁴⁶ .0.3	py-natsort#5.2.0	suite-sparse 5.2
autoconf#Z.69	gdbm 1.14.1	L	lcms#2.8	libxv#1.0.10	1	patch#2.7.6	py-nose#1.3.7	sundials@3.1.0
automake 1.14	geopm Q.4.	All and a second se	legion#17.10.0	libxvmc@1.0.	9	pcne##/41	py-numexpred, 6.1	superlu 5.2.1
automake=1.15.1	gettexted.	19.8.1	leveldb#1.20	libyogrt#1.2	18-6	pcne/8.41	py+numpy#1.13.3	superlu-dist@5.2.
ax1#0.1.1	git#2.15.1		libarchive 3.3.2	lmod#7.7.13		pdsh#2,31	py-pandasild,21,1	swig#3.0.12
binutils=2.27	glib#2:56.6	1	Libbsd90.8.6	lua 5, 1.4		pdt#3.25	py-pbr#3,1.1	sz#1,4.12.3
binutils=2.29.1	g1#90,9.7.1		libcircle90.2.1-rc.)	1 lua-luafiles	system 1_6_3	per1#5.24.1	py-pillow=5.2.0	tar#1,29
bison#3.8.4	globalarray	/5//5.7	libedit#3,1-20170325	lua-luaposix	#33.4.8	petsc#3.8.4	py-pkgconfig#1,2_2	tasmanian=6.0
bolt#1.0b1	glproto 1.	4,17	libffi#3,2.1	lwgrp#1.0.2		pflotran#ksdk-0.3.0	py-py=1.4.33	tau#2.28
boost#1.66.0	gmp16.1.2		libice=1.0.9	1z4#1.4.1.2		pixmanP0.34.0	py-pycodestyle=2.1.1	tcl#8.6.8
boost#1.66.0	gobject-int	trospection#1.49.2	libiconvel.15	1zma=4.32.7		pkgconf=1.4.0	py-pyflakes@1.6.8	texinfor6.5
boost#1.68.0	got.challd.0	2	libjpeg-turbow1.5.3	1zo#2.89		presentproto 1.0	py-pyparsing 2.2.0	tk#8.5.8
bzipZel 0.6	gotchalldeve	ring	libmng=Z.0.3	m491.4.18		protobuf#3.5.1.1	py-pytables@3.3.0	trilinos 12.12.1
c-blosc#1.12.1	gperf#3.0.		libpcioccessed. 13.5	matio 1.5.9		py-argparse=1.4.0	py-pytest@3.6.0	turbine 1.0.0
cairo?1.14.12	harfbuzz	4.6	libpfm404.X.0	metis#5.1.0		py-bobel#2.4.0	py-pytz=2017.7	turbine#1.0.8
caliper#1.8.0	hdf501.8.15		libpng 1.6.34	nfen#3.3.2		py-bottleneck#1.0.0	py-scipy=1.0.0	umpiremoster
cmake#3.11.1	hdf501.8.1		libpthread-stubs	minicondo2	.3.30	py-configparser=3.5.	<pre>0 py-setuptools=39.0.1</pre>	unifycreaster
conduit=master	hdf5e1.10.1		libquo 1.1	minicondo3	1.3.30	py-cycler=0.10.0	py-six#1,11.0	util-macros 1.19.
curl#7.59.0	hdf5#1_10,1		libsigsegv=2.11	mpich#3.2.1		py-cython 0, 28_1	py-subprocess32#1.2.7	veloce1.0
damageproto 1 2.1	hdf5#1_10_1		libsm01.2.2	mumps#5.1.1		py-dateutile2.5.2	python 2.7.14	videoproto#2,3.3
darshan-runtime	1.0 hdf5#1.10.1		libtiffe4.0.6	nasm 2.13.03		py-enum34/1.1.6	ghull#2015.2	vtkminuster
darshan-util@3.1.	6 help2man#1	47.4	libtiffes.0.8	ncurses#6.0		py-flake803.5.0	gthreads#1.12	vtkm#1.1.0
doxygen 1.8.12	hoctoolkit	2017.06	libtool 2.4	netcdfe4.4.1	- A	py-funcsios 4	r#3.4.3	xcb-proto=1.13
dtcmp#1.1.8	hpctoolkit-	externals@2017.06	libtool#2.4.7	netlib-scale	pocket . n. 2	py-functools32=3.2.3	2 raja 0.5.3	xextproto 7.3.0
er#0.0.3	hwloc#1.11	9	libtool#2.4.6	nettles.3		py-hSpy#2.7.1	rankstr#0.8.2	xproto#7.0.31
exmcutils#0.5.3	hwlock2.0.1		libunwind 1.1	ninja 1.8.2		py-hypothesis#3.7.0	readline#7.0	xtrans#1.3.5
expot=2.2.2	hypre 2,13	8	libx11=1.6.5	numoctle2.0.	11	py-11010202.9.6	redset#0.0.3	xz#5.2.3
fftw=3.3.7	hypnes2,13	0	libxou=1.0.8	openblase0.2	1.20	py-kiwisolver#1.0.1	nubye2.2.0	zfp-0.5.0
fixesproto-5.0	1cu4cm60.1		libxcb=1.13	openmpi#3.0	1	py-lited.5.0	ruby-ronning, 7, 3	zlib#1.2.11
flex#2.6.4	inputproto	2.3.2	libxdamage#1.1.4	openssl#1.0	20	py-mako#1.0.4	scr#1.2.2	zshi 5.4.2
font-util#1.3.1	intel-tbb	5.819	libsdmcp=1.1.2	otf2+2.1		py-markupsofe=1.0	shuffiles0.0.3	zstd#1.5.8
fontconfig 7 12	idk#8u141-1	15	libsest#1.5.3	panao 1.41.0	1	py-matplotlibs7.2.2	snoppy=1,1,7	

Packages installed using Spack

- UnifyCR
- Veloc
- xSDK
- Zfp

All ST products will be released through E4S



Exascale Computing Project (ECP) Research, Development & Delivery (RD&D) Overview





DOE Exascale Program: The Exascale Computing Initiative (ECI)

Three Major Components of the ECI



ECP's three technical areas have the necessary components to meet national goals

$\langle -$	Perform	ant	mission and sc	ience applicati	ons @) scale	
N	Aggressive RD&D Mission apps & Project integrated S/W stack		ission apps & rated S/W stack	Deployment to DOE HPC Facilities		Hardware tech advances	
Application Development (AD)			Software Technology (ST)		Hardware and Integration (HI)		
Develop and enhance the predictive capability of applications critical to the DOE 24 applications including national security, to energy, earth systems, economic security, materials, and data		Deliver expande integrated soft achieve full poter comp 67 unique soft spanning progra and run times, data and vit	ed and vertically sware stack to ntial of exascale uting ware products amming models math libraries, sualization	l proc le 6 L e de a	ntegrated delivery of ECP ducts on targeted systems at eading DOE HPC facilities JS HPC vendors focused on exascale node and system esign; application integration nd software deployment to facilities		

ECP software technologies overview

National Nuclear Security Administration

Programming Models & Runtimes

- Enhance and get ready for exascale the widely used MPI and OpenMP programming models (hybrid programming models, deep memory copies) • Development of performance portability tools (e.g. Kokkos and Raja) • Support alternate models for potential benefits and risk mitigation: PGAS (UPC++/GASNet)
- (UPC++/GASNet) ,task-based models (Legion, PaRSEC)
- •Libraries for deep memory hierarchy and power management

Development Tools

 Continued, multifaceted capabilities in portable, opensource LLVM compiler ecosystem to support expected ECP architectures, including support for F18
 Performance analysis tools that accommodate

new architectures, programming models, e.g., PAPI, Tau

Math Libraries

 Linear algebra. iterative linear solvers, direct linear solvers, integrators and nonlinear solvers. optimization, FFTs, etc •Performance on new node architectures: extreme strong scalability Advanced algorithms for multiphysics, multiscale simulation and outer-loop analysis Increasing quality, interoperability. complementarity of math libraries

Data and Visualization

- I/O via the HDF5
 API
- Insightful, memory-efficient in-situ visualization and analysis – Data reduction via scientific data compression
- Checkpoint restart

Software Ecosystem

Develop features in Spack necessary to support all ST products in E4S, and the AD projects that adopt it
Development of Spack stacks for reproducible turnkey deployment of large collections of software

Optimization and interoperability of containers on HPC systems
Regular E4S releases of the ST software stack and SDKs with regular

integration of new

ST products

NNSA ST

- Open source
 NNSA Software
 projects
- Projects that have both mission role and open science role
- Major technical areas: New programming abstractions, math libraries, data and viz libraries
- Cover most ST technology areas
- Subject to the same planning, reporting and review processes

10

ECP Software Technology Leadership Team

Mike Heroux, Software Technology Director

Mike has been involved in scientific software R&D for 30 years. His first 10 were at Cray in the LIBSCI and scalable apps groups. At Sandia he started the Trilinos and Mantevo projects, is author of the HPCG benchmark for TOP500, and leads productivity and sustainability efforts for DOE.

Jonathan Carter, Software Technology Deputy Director

Jonathan has been involved in the support and development of HPC applications for chemistry, the procurement of HPC systems, and the evaluation of novel computing hardware for over 25 years. He currently a senior manager in Computing Sciences at Berkeley Lab.

Rajeev Thakur, Programming Models and Runtimes (2.3.1)

Rajeev is a senior computer scientist at ANL and most recently led the ECP Software Technology focus area. His research interests are in parallel programming models, runtime systems, communication libraries, and scalable parallel I/O. He has been involved in the development of open source software for large-scale HPC systems for over 20 years.

Jeff Vetter, Development Tools (2.3.2)

Jeff is a computer scientist at ORNL, where he leads the Future Technologies Group. He has been involved in research and development of architectures and software for emerging technologies, such as heterogeneous computing and nonvolatile memory, for HPC for over 15 years.

Lois Curfman McInnes, Math Libraries (2.3.3)

Lois is a senior computational scientist in the Mathematics and Computer Science Division of ANL. She has over 20 years of experience in highperformance numerical software, including development of PETSc and leadership of multi-institutional work toward sustainable scientific software ecosystems.

Jim Ahrens, Data and Visualization (2.3.4)

Jim is a senior research scientist at the Los Álamos National Laboratory (LANL) and an expert in data science at scale. He started and actively contributes to many open-source data science packages including ParaView and Cinema.

Todd Munson, Software Ecosystem and Delivery (2.3.5)

Todd is a computational scientist in the Math and Computer Science Division of ANL. He has nearly 20 years of experience in high-performance numerical software, including development of PETSc/TAO and project management leadership in the ECP CODAR project.

Rob Neely, NNSA ST (2.3.6)

Rob is an Associate Division Leader in the Center for Applied Scientific Computing (CASC) at LLNL, chair of the Weapons Simulation & Computing Research Council, and lead for the Sierra Center of Excellence. His efforts span applications, CS research, platforms, and vendor interactions.

We work on products applications need now and into the future

Key themes:

- Exploration/development of new algorithms/software for emerging HPC capabilities:
- High-concurrency node architectures and advanced memory & storage technologies.
- Enabling access and use via standard APIs.

Software categories:

- The next generation of well-known and widely used HPC products (e.g., MPICH, OpenMPI, PETSc)
- Some lesser used but known products that address key new requirements (e.g., Kokkos, RAJA, Spack)
- New products that enable exploration of emerging HPC requirements (e.g., SICM, zfp, UnifyCR)

Example Products	Engagement
MPI – Backbone of HPC apps	Explore/develop MPICH and OpenMPI new features & standards.
OpenMP/OpenACC –On-node parallelism	Explore/develop new features and standards.
Performance Portability Libraries	Lightweight APIs for compile-time polymorphisms.
LLVM/Vendor compilers	Injecting HPC features, testing/feedback to vendors.
Perf Tools - PAPI, TAU, HPCToolkit	Explore/develop new features.
Math Libraries: BLAS, sparse solvers, etc.	Scalable algorithms and software, critical enabling technologies.
IO: HDF5, MPI-IO, ADIOS	Standard and next-gen IO, leveraging non-volatile storage.
Viz/Data Analysis	ParaView-related product development, node concurrency.

	WBS	WBS Name	CAM/PI	PC
I ECP ST	2.3	Software Technology	Heroux, Mike, Carter, J.	_
	2.3.1	Programming Models & Runtimes	Thakur, Rajeev	_
Drojocte	2.3.1.01	PMR SDK	Shende, Sameer	Shende, Sameer
	2.3.1.07	Exascale MPI (MPICH)	Balaji, Pavan	Bayyapu, Neelima
-	2.3.1.08	Legion	McCormick, Pat	McCormick, Pat
	2.3.1.09	PaRSEC	Bosilica, George	Carr, Earl
- WRS	2.3.1.14	Pagoda: UPC++/GASNet for Lightweight Communication and Global Address Space Support	Baden, Scott	Hargrove, Paul (and PI)
	2.3.1.16	SICM	Lang, Michael	Vigil, Brittney
I - Name	2.3.1.17	OMPI-X	Bernholdt, David	Alicia Grundhoffer
	2.3.1.18	RAJA/Kokkos	Trott, Christian Robert	Trott, Christian
- PIS	2.3.1.19	Argo: Low-level resource management for the OS and runtime	Beckman, Pete	Gupta, Rinku
	2.3.2	Development Tools	Vetter, Jeff	_
- F C 5	2.3.2.01	Development Tools Software Development Kit	Miller, Barton	Tim Haines
	2.3.2.06	Exa-PAPI++: The Exascale Performance Application Programming Interface with Modern C++	Dongarra, Jack	Jagode, Heike
	2.3.2.08	Extending HPCToolkit to Measure and Analyze Code Performance on Exascale Platforms	Mellor-Crummey, John	Mellor-Crummey, John
	2.3.2.10	PROTEAS-TUNE	Vetter, Jeff	Glassbrook, Dick
	2.3.2.11	SOLLVE: Scaling OpenMP with LLVm for Exascale	Chapman, Barbara	Kong, Martin
	2.3.2.12	FLANG	McCormick, Pat	Perry-Holby, Alexis
FCP ST	2.3.3	Mathematical Libraries	McInnes, Lois	_
	2.3.3.01	Extreme-scale Scientific xSDK for ECP	Yang, Ulrike	Yang, Ulrike
Stata	2.3.3.06	Preparing PETSc/TAO for Exascale	Smith, Barry	Munson, Todd
Jais	2.3.3.07	STRUMPACK/SuperLU/FFTX: sparse direct solvers, preconditioners, and FFT libraries	Li, Xiaoye	Li, Xiaoye
	2.3.3.12	Enabling Exascale Simulations with SUNDIALS and hypre	Woodward, Carol	Woodward, Carol
	2.3.3.13	CLOVER: Computational Libraries Optimized Via Exascale Research	Dongarra, Jack	Carr, Earl
	2.3.3.14	ALExa: Accelerated Libraries for Exascale/ForTrilinos	Turner, John	Alicia Grundhoffer
= 33 L4	2.3.4	Data and Visualization	Ahrens, James	_
nrojects	2.3.4.01	Data and Visualization Software Development Kit	Atkins, Chuck	Atkins, Chuck
	2.3.4.09	ADIOS Framework for Scientific Data on Exascale Systems	Klasky, Scott	Alicia Grundhoffer
- 11 PI/PC	2.3.4.10	DataLib: Data Libraries and Services Enabling Exascale Science	Ross, Rob	Ross, Rob
	2.3.4.13	ECP/VTK-m	Moreland, Kenneth	Moreland, Kenneth
same	2.3.4.14	VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart/Sz	Cappello, Franck	Ehling, Scott
	2.3.4.15	ExalO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	/Byna, Suren	Bagha, Neelam
- 22 1 1/1 0	2.3.4.16	ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis/ZFP	Ahrens, James	Turton, Terry
different	2.3.5	Software Ecosystem and Delivery	Munson, Todd	_
	2.3.5.01	Software Ecosystem and Delivery Software Development Kit	Willenbring, James M	Willenbring, James M
	2.3.5.09	SW Packaging Technologies	Gamblin, Todd	Gamblin, Todd
	2.3.6	NNSA ST	Neely, Rob	-
	2.3.6.01	LANL ATDM	Mike Lang	Vandenbusch, Tanya Marie
	2.3.6.02	LLNL ATDM	Becky Springmeyer	Gamblin, Todd
	2.3.6.03	SNL ATDM	Jim Stewart	Trujillo, Gabrielle

ECP ST Product Dictionary List – Actively managed

Widely-recognized product names. Enables mapping between AD & Facilities dependencies and ST development efforts.

- MPI MPICH, OpenMP
- C++/C/Fortran LLVM
- Fortran Flang
- hypre hypre
- Actively managed at L2 level.

	EXASCALE COMPUTING PROJECT
--	----------------------------------

ADIOS	Flux		
	Lortron	OpenMP	SUNDIALS
AIVIL	Foluan	PAPI	SuperLU
Ascent	GASNet	Panyrus	SYCL
BLAS	Ginkgo	Doroviou	
С	HDF5	Palaview	
C ++	HPCToolkit	PaRSEC	IASMANIAN
Colinon	hump	PETSc/TAO	TAU
Camper		PnetCDF	Trilinos
Catalyst	Kokkos	PowerStack	UMap
CHAI	KokkosKernels	Ρ ΔΙΔ	Umnire
Cinema	LAPACK		Unifi
CUDA	Legion		UIIIY
 Darshan	libEnsemble	ScaLAPACK	UPC++
	MorES	SCR	VeloC
		SICM	Visit
Dyninst	MFEM	Spack	VTK-m
E4S	MPI	SPOT	vSDK
FFT	OpenACC		
FleCSI	OpenCL		ΔΓΓ

14

Product	URL	Description	Deployment Scope	Technical Area	POC
1. ADIOS	https://github.com/or nladios/ADIOS2	I/O and data management library for storage I/O, in-memory code coupling and online data analysis and visualization workflows.	Broad	Data & Viz	<u>Scott Klasky</u>
2. AID		Advanced Infrastructure for Debugging (AID) - a toolkit of debugging and correctness tools for extreme scale and GPU architectures	Experimental	Tools	Dong Ahn
AID: STAT	https://github.com/L LNL/stat	Stack Trace Analysis Tool - lightweight debugging at extreme scale	Broad	Tools	
AID: Archer	https://github.com/P RUNERS/archer	Data race detection tool for OpenMP applications	Broad	Tools	
AID: FLIT	https://github.com/P RUNERS/FLiT	Quickly detect discrepancies in floating point computation across hardware, compilers, libraries and software	Moderate	Tools	
AID: ReMPI	<u>https://github.com/P</u> <u>RUNERS/ReMPI</u>	A lightweight record-and-replay tool for MPI+OpenMP applications	Moderate	Tools	
AID: FPChecker	https://github.com/L LNL/FPChecker	Floating point exception trapping for NVIDIA GPUs	Experimental	Tools	

Partial List of Jira Dependency Issues

Se	arch	Save as							< Share 🏦	Export	• O Tools •
Pr	oject: All 🐱	Dependency 👻 Status: All 🛩	Assignee: All 🗸 Conta	ins text	More	Search	Advanced				= *
1-5	0 of 814(2										Columns 🛩
т	Кеу	Summary	Assignee	Reporter	Status ↑	Resolution	Created	Due	Baseline end date	Links	Development
	INT-691	PETSc/TAO ↔ Spack	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
0	INT-686	HDF5 ↔ PETSc/TAO	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
0	INT-690	PETSc/TAO ↔ xSDK	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
0	INT-687	MPI-IO ↔ PETSc/TAO	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
0	INT-689	OpenMP ↔ PETSc/TAO	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
0	INT-677	libEnsemble ↔ xSDK	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				
Ø	INT-681	BLAS ↔ PETSc/TAO	Unassigned	Todd Munson	APPROVED	Approved	2019-09-10				

Jira Dependency Issue Type

- Official ST and AD Product lists enable rigorous dependency management.
- New Jira Dependency issue type.
- Explicit dependency connections between AD/CD/ST products

Integration					
Dependents by Pro	ducer				
Dependen	ts by Producer				
ST Producers -	AD Consumers, ST Consumers -	Draft, Approved +			Euro
Reset scale			Critical Dependents	Important Depender	9 De
55					
50					10 Trige
4 5			C Drill thro	ugh Issue	ing:
40 40	1.00			49110000	
ue 35	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1		-		Optional It
e 30			1 A A		-
0 20 0 20					11 Linke
g 15	1 - A A		and the second second		
z 10					
5					
0				╸╴┫╴┫╶╸╴┫╶┛	A++
NOS IN	and of cot act st with	of the the set	tar aikit why the ast of	tay as is pre ct	ALL
AD ALL	in hi preto U. V	J. C3. 40	octor voren i si	or the winds. 4	
			8		12
				Producer	

Reporter* ECP Support							
	Start typing to get a list of possible matches.						
6 Producer*	Programming Models and Runtimes \$	Legion	\$				
	Select the application code or product name	that the Consumer depends on.					
7 Consumer*	Data Analytics and Optimization 🗘	CANDLE	\$				
	Select the application code or ST product tha	t depends on the Producer.					
Dependency Level*	Critical						
8	Important						
	Interested						
	Critical: The team is entirely dependent on th	ne producer for this functionality, and th	ere are no alternatives				
	available.						
	Important: The team believes this producer is the best source for this functionality, but alternative sources						
	exist.						
	work.	the functionality that they are likely to t	ry to adopt it into their				
9 Punctionality	ENTER BRIEF DESCRIPTION HERE						
Description	Enter a brief description of the functionality th	hat this Producer provides this Consum	er.				
10 Trigger Event*	Unknown 🗘						
	Provide the event/quarter you expect this dep	pendency will be needed by.					
THE REAL PROPERTY OF	and the second se						
Optional Items							
11 LINKED ISSUES	DIOCKS Ŧ						
Issue			- +				
	Begin typing to search for issues to link. If you	ı leave it blank, no link will be made.					
Attachment	C Drop files to attached	ach, or browse.					
12 Labels	auto-generated ×		-				
	Begin typing to find and create labels or press down to select a suggested label.						

Delivering a modular, interoperable, and deployable software stack

ECP ST Planning Process: Hierarchical, three-phase, cyclical

Baseline

- FY20–23 Baseline Plan High level Definitions
- Q2 FY19 start
- FY20 Base plan
- FY21–23 planning packages

Annual Refinement

FY Refine Baseline Plan As Needed Basic activity definitions

- 6 months prior to FY
- 4-6 P6 Activities/year
- Each activity:
 - % annual budget
 - Baseline start/end
 - High level description

Per Activity

Detailed Plan Complete activity definitions

- 8 weeks prior to start
- High-fidelity description
- Execution strategy
- Completion criteria
- Personnel details

Two-level
changeControlChanges to Cost, Scope,
and ScheduleMinorMajorLightweight
Review in
Jira, L3 and
L2 leadsChange
Control
Board
Review, ECP
leadership

KPP-3 Definition

KPP ID	Description of Scope	Threshold KPP	Objective KPP	Verification Action/Evidence
KPP-1	Performance of scientific and national security applications relative to today's performance	50% of selected applications achieve Figure of merit* improvement ≥50	100% of selected applications achieve Figure of merit improvement stretch goal	Independent assessment of measured results and report that threshold goal is met
KPP-2	Broaden the reach of exascale science and mission capability	50% of selected applications can execute their challenge problem*	100% of selected applications can execute their challenge problem stretch goal	Independent assessment of mission application readiness
KPP-3	Productive and Sustainable Software Ecosystem	Software teams meet 50% of their weighted impact goals*	Software teams meet 100% of their weighted impact stretch goals	Independent assessment verifying threshold goal is met
КРР-4	Enrich the HPC Hardware Ecosystem	Vendors meet 80% of all the PathForward milestones	Vendors meet 100% of all the PathForward milestones	Independent assessment of the impact and timeliness of PathForward milestones

KPP-3: Focus on capability integration

- **Capability:** Any significant product functionality, including existing features adapted to the preexascale and exascale environments, that can be integrated into a client environment.
- Capability Integration: Complete, sustainable integration of a significant product capability into a client environment in a pre-exascale environment (tentative score) and in an exascale environment (confirmed score).
- **Product:** Capabilities are integrated into primary products chosen from a product dictionary.

• Primary products examples:

- MPI is commonly known by users. MPICH and OpenMPI both provide implementations of that product.
- Fortran is a product. Flang is a particular Fortran product. LLVM is a backend for some Fortran compilers.
- FFT is a product. FFTX, FFT-ECP provide FFT capabilities through interchangeable interfaces.
- C++ is a product. Clacc provides capabilities for Clang, as does LLVM.
- Details: See Confluence page export.

Delivering a Modular, Interoperable, and Deployable Software Stack

Objectives

- Lower barrier to use ST products
- Lower barrier to enable Facilities to install all or parts
- Enable interoperability between ST products
- Enable uniform APIs where possible
- Improve KPP-3 success

Challenges

- Large diverse group of ST products
- Different project management styles
- Lack of initial drivers for integration

https://spack.io

488

2014

2013

🔰 @spackpm

Spack is used worldwide!

Over 2,800 software packages

Over 150,000 downloads in the past year

Over **300** contributors from labs, academia, industry

Support scientific stacks with multiple languages

- Flexibility:
 - Build packages many different ways
 - Change compilers and flags in builds

Inspired by Homebrew, Nix, some others

- Swap implementations of libraries (MPI, BLAS, etc.)

Spack

 Run on laptops, Linux clusters, and the largest supercomputers in the world

A flexible package manager for HPC

Easy installation

- \$ git clone <u>https://github.com/spack/spack</u>
- \$. spack/share/spack/setup-env.sh
- \$ spack install hdf5

Easy customization

- \$ spack install mpileaks@3.3
- \$ spack install mpileaks@3.3 %gcc@4.7.3 +threads
- \$ spack install mpileaks@3.3 cppflags="-03 -g3"
- \$ spack install mpileaks@3.3 target=haswell
- \$ spack install mpileaks@3.3 ^mpich@3.2

2015

2016

2018

Software Development Kits are a key delivery vehicle for ECP

 A collection of related software products (called packages) where coordination across package teams will improve usability and practices and foster community growth among teams that develop similar and complementary capabilities

• Attributes

- **Domain scope:** Collection makes functional sense
- Interaction model: How packages interact; compatible, complementary, interoperable
- Community policies: Value statements; serve as criteria for membership
- Meta-infrastructure: Invokes build of all packages (Spack), shared test suites
- Coordinated plans: Inter-package planning. Augments autonomous package planning
- Community outreach: Coordinated, combined tutorials, documentation, best practices
- Overarching goal: Unity in essentials, otherwise diversity

ECP ST SDKs: Grouping Similar Products for collaboration & usability					
Prog Models & Runtimes Core	Math Libraries (xSDK)				
Tools & Technologies	Viz Analysis and Reduction				
Compilers & Support	Data mgmt., I/O services & Checkpoint/Restart				

Horizonal (vs Vertical) Coupling

- Common substrate
- Similar function and purpose
 - •e.g., compiler frameworks, math libraries

Horizontal grouping:

- Assures X=Y.
- Protects against regressions.
- Transforms code coupling from heroic effort to turnkey.
- Potential benefit from common Community Policies
 - •Best practices in software design and development and customer support
- Used together, but not in the long vertical dependency chain sense
- Support for (and design of) common interfaces
 - Commonly an aspiration, not yet reality

There are 5 L4 projects to define and/or enhance SDKs

- Each L3 area has an L4 project devoted to SDK definition and coordination across the portfolio
- Software ecosystem L4 project focuses on packaging
 - Spack
 - Containers
- Strong coordination with HI Software Deployment projects
- Drives milestone-based planning

ECP ST SDKs will span all technology areas

Motivation: Properly chosen cross-team interactions will build relationships that support interoperability, usability, sustainability, quality, and productivity within ECP ST.

Action Plan: Identify product groupings where coordination across development teams will improve usability and practices, and foster community growth among teams that develop similar and complementary capabilities.

PMR Core (17)	Compilers and Support (7)	Tools and Technology (11)	xSDK (16)	Visualization Analy and Reduction (9)	vsis Data mgmt, I/O Services, Checkpoint restart (12)	Ecosystem/E4S at-large (12)
QUO	openarc	TAU	hypre	ParaView	SCR	mpiFileUtils
Papyrus	Kitsune	HPCToolkit	FleSCI	Catalyst	FAODEL	TriBITS
SICM	LLVM	Dyninst Binary Tools	MFEM	VTK-m	ROMIO	MarFS
Legion	CHiLL autotuning comp	Gotcha	Kokkoskernels	SZ	Mercury (Mochi suite)	GUFI
Kokkos (support)	LLVM openMP comp	Caliper	Trilinos	zfp	HDF5	Intel GEOPM
RAJA	OpenMP V & V	PAPI	SUNDIALS	Vislt	Parallel netCDF	BEE
CHAI	Flang/LLVM Fortran comp	Program Database Toolkit	PETSc/TAO	ASCENT	ADIOS	FSEFI
PaRSEC*		Search (random forests)	libEnsemble	Cinema	Darshan	Kitten Lightweight Kernel
DARMA		Siboka	STRUMPACK	ROVER	UnifyCR	COOLR
GASNet-EX		C2C	SuperLU		VeloC	NRM
Qthreads		Sonar	ForTrilinos		IOSS	ArgoContainers
BOLT			SLATE		HXHIM	Spack
UPC++			MAGMA	DM	D	
MPICH			DTK			
Open MPI			Tasmanian	Mat	h Libraries	
Umpire			TuckerMPI	Dat	a and Vis	
AML				Ecc	osystems and delivery	

AML

SDK Summary

- SDKs will help reduce complexity of delivery:
 - Hierarchical build targets.
 - Distribution of software integration responsibilities.
- New Effort: Started in April 2018, fully established in August 2018.
- Extending the SDK approach to all ECP ST domains.
 - SDKs create a horizontal coupling of software products, teams.
 - Create opportunities for better, faster, cheaper pick all three.
- First concrete effort: Spack target to build all packages in an SDK.
 - Decide on good groupings.
 - Not necessarily trivial: Version compatibility issues, Coordination of common dependencies.
- Longer term:
 - Establish community policies, enhance best practices sharing.
 - Provide a mechanism for shared infrastructure, testing, training, etc.
 - Enable community expansion beyond ECP.

ST works with HI and Facilities to produce a stack, build, and container infrastructure

E4S/Software

- Regular additions of software technologies, integration testing, and releases of E4S
- Certified Spack recipes and binaries for software technologies
- Currently includes 37 software technologies
- Broad adoption of reliable, performant, exascale-ready software

Packaging Tech/Facilities

- Turnkey deployment of large collections of software technologies at facilities
- Coordinate with and leverage facility personnel and resources, such a continuous integration capabilities
- Spack recipes for over 2,800 software packages
- Broad deployment of exascaleready software technologies

Packaging Tech/Applications

- Optimize and improve interoperability of container technologies
- Enable entire application to be packaged into reproducible container images
- Accelerate application development and deployment workflows

Continuous Integration (CI) is critical for delivery of high-quality SW Fundamental value: Minimize time between fault injection & detection

•	Work with ST/F	[:] acilities on a	consistent	deployment m	odel for ST
	products				

- Coordinate a support model for ST products
- Provide linkages to Facility "site stacks" and vendor offerings
- Build common infrastructure for testing and deployment
- Training, knowledge transfer, promotion of ST products

HI

Deplo

Continuous Integration (CI) is critical for delivery of high-quality SW Fundamental value: Minimize time between fault injection & detection

ECP Goal: To establish a cross-site Continuous Integration testing infrastructure that:

- Provides for account authentication and access to CI test resources across multiple sites
- Provides unique and targeted HPC test resources to support software development teams
- Establishes a standard process across the DOE sites for software development testing
- Allows for the verification of development efforts through automated building and testing across sites to better identify errors and improve code efficiency
- Continuous Integration/Continuous Delivery pipelines are enabled through a combination of the web
 application and project configuration files and are executed by selected runners

Preparing for Exascale Platforms

Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of accelerated node systems supporting DOE's mission

Hardware realities are forcing new thinking of algorithmic implementations and the move to new algorithms

Algorithmic Implementations

- Reduced communication/data movement
 - Sparse linear algebra, Linpack, etc.
- Much greater locality awareness
 - Likely must be exposed by programming model
- Much higher cost of global synchronization
 - Favor maxim asynchrony where physics allows
- Value to mixed precision where possible
 - Huge role in AI, harder to pin down for PDEs
- Fault resilience?
 - Likely handled outside of applications

New Algorithms

- Adopting Monte Carlo vs. Deterministic approaches
- Exchanging on-the-fly recomputation vs. data table lookup (e.g. neutron cross sections)
- Moving to higher-order methods (e.g. CFD)
- Particle algorithms that favor collecting similar events together rather than parallelism though individual histories

Exascale Application Development Challenges Overall

1) Porting to accelerator-based architectures

2) Exposing additional parallelism

3) Coupling codes to create new multiphysics capability

4) Adopting new mathematical approaches

5) Algorithmic or model improvements

6) Leveraging optimized libraries

Topology-Aware Collectives

Scope and objectives

- Exascale applications rely heavily on MPI collectives to efficiently implement group communication.
- MPI collectives can be more efficiently performed with knowledge of both the application and machine topology.
- Design and implement topology-aware support for applications and machine topologies for use at exascale.

Impact

It is imperative for communication libraries such as Exascale MPI to efficiently perform collective operations on both physical and virtual topologies. Many ECP applications rely on MPI collectives for performance and efficiency.

Project accomplishment

- We have prototyped a suite of new collective algorithms intended to perform well on large multicore machines, as well as network topologies such as a dragonfly.
- We have prototyped neighborhood collectives using sub-communicators to exploit the virtual topology of many ECP applications, i.e. nearest neighbor exchange.
- We have designed and prototyped a flexible framework for selecting the best collective algorithm for a given set of parameters.

Topology-Aware Collective Algorithms and Selection

ECP WBS

PI

k-ary tree (k=3)

MPI_Neighbor_ allgather

2.3.1.07 STPM09

Pavan Balaji, ANL

Software

K-ary, K-nomial, Recursive Exchange collective algorithms: http://www.mpich.org/static/downloads/3.3.1/mpich-3.3.1.tar.gz Intelligent Collective Selection: http://www.mpich.org/static/downloads/3.3.1/mpich-3.3.1.tar.gz

Deliverables

Message-combining Optimizations for Neighborhood Collectives: <u>https://github.com/pmodels/mpich/pull/3892</u>

Enhanced LLVM OpenMP compiler

ECP WBS SOLLVE

PI Barbara Chapman, BNL

Members BNL, ANL, Rice, ORNL, LLNL

Scope and objectives		Cool image				
OpenMP, the primary programming model for on-node currency in ECP applications				$LUD^4 \qquad \bigcirc \ \ \times \qquad \qquad$	Needle \odot \bigcirc 1.6	
Improve the performance and functionality of LLVM's OpenMP implementation, especially on GPUs		OpenMP for GPU "target region" optimizations.	1.2 - × × × ∞		© - 1.4 × dnpe	
Improve the interaction between OpenMP and loop-nest optimizations, and the ability to express loop-nest optimizations in OpenMP	v3.1 on a NVIDIA K40 and V100. Speedup vs. Input size.					
Implement new OpenMP features in LLVM, specifically features for mapping complex data structures in accelerators			$\begin{array}{c} 1 & 0 & -\frac{1}{2} & 0 & -\frac{1}{2} & -\frac{1}{2} \\ 2^5 & 2^6 & 2^7 & 2^8 & 2^9 & 2^{10} & 2^{11} & 2^{12} & 2^{13} & 2^{14} \\ \end{array}$	$1 = \frac{2}{5} \frac{2^{6}}{2^{6}} \frac{2^{7}}{2^{8}} \frac{2^{9}}{2^{9}} \frac{2^{10}}{2^{11}} \frac{2^{12}}{2^{12}} \frac{2^{13}}{2^{14}} \frac{2^{14}}{2^{14}}$		
pact		Project accomplishment				
OpenMP is the de facto standard multithreading even outside ECP. Any impact within ECP translates into impact outside of it	•	 Implemented improved general infrastructure allowing for the optimization o OpenMP target regions, with specific optimization enablement for GPUs. 				

- Most ECP applications rely on OpenMP for intranode and MPI for internode programming. Improvements in OpenMP performance on GPUs will translate directly to performance on exascale architectures
- OpenMP features implemented in LLVM are often incorporated into LLVMbased vendor toolchains
- Designed and prototyped new OpenMP loop-optimization pragmas.
 Demonstrated that these pragmas can replace complex source-level transformations yielding performance and productivity improvements.
- Implemented initial support for OpenMP user-defined mapping directives.

Deliverables

Fulfillment of NVIDIA GPU Performance Counter and Power Management Support

ECP WBSExa-PAPIPIDongarra, UTKMembersHeike Jagode (Co-PI)
Anthony Danalis (Co-PI)
Tony Castaldo

Scope and Objectives

- PAPI support for NVML, CUPTI, NVLINK performance events, and power consumption management on GPUs.
- Added PAPI support for TESLA V100 GPUs and NVLINK.
- Increase options and info for multi-objective optimization; e.g. reducing energy usage when GPUs are not the speed bottleneck, or a speedy result is not paramount.

Power Reading and Capping on TESLA V100 GPUs

Impact

- **PAPI "nvml":** Monitoring of GPU power consumption, fan speed, temperature, and power capping support: allows developers to change run profiles to reduce energy cost.
- **PAPI "cuda":** Monitoring of GPU and NVLINK performance counters aids developers in producing more efficient code by profiling the utilization of the latest GPU resources and diagnosing performance bottlenecks.

Project Accomplishment

- Production interface of PAPI to CUPTI and NVML, with well-documented example code, proving the availability and efficacy of the controls in a standardized way.
- We demonstrate PAPI NVIDIA GPU power reading and control, and the use of performance counters across multiple GPUs

Deliverables git clone https://bitbucket.org/icl/papi.git. In papi/src/components/nvml/tests: nvmlcap_plot.cu, nvml_power_limiting_test.cu. In papi/src/components/cuda/tests: simpleMultiGPU.cu. Corresponding Makefile in each directory, commentary in code.

Hierarchical Off-Diagonal Low-Rank Matrix Compression in the STRUMPACK Preconditioner

Scope and objectives

- Sparse direct solvers and factorization based preconditioners. This milestone focuses on dense matrix compression for use in the preconditioner
- Need to improve efficiency and rate of compression for frontal matrices (fill-in in multifrontal sparse factorization)
- Accurate and fast compression improves preconditioning for numerically hard problems

Impact

- Asymptotically reduce memory and computation in sparse solvers. This allows ECP and other users, (also through MFEM or PETSc) to solve larger problems, faster
- Provide robust, purely algebraic, parallel solvers and preconditioners for numerically challenging ECP applications, for instance, high frequency Helmholtz

Deliverables	https://github.com/liuyangzhuan/ButterflyPACK
	https://github.com/pghysels/STRUMPACK
	 Summit (OLCF) was used for the evaluation

HODLR and Butterfly data-sparse formats

Off diagonal blocks are low-rank/butterfly compressed. Butterfly multi-level low-rank compression is based on ideas from FFT.

Project accomplishment

- HODLR + butterfly in ButterflyPACK (Fortran/C), and integrated in STRUMPACK (C++) preconditioner
- Evaluation on MFEM indefinite Maxwell test code, typically solved with direct solvers
- Observed much reduced memory usage compared to direct solver: 20x compression rate on largest fronts using butterfly

Cross-platform implementation for parallel incomplete factorizations

ECP WBS STMS11-PEEKS Dongarra (UTK) ΡΙ Boman (SNL) Hoemmen (SNL) Members

Anzt, Yamazaki (UTK)

Scope and objectives

- MAGMA-sparse release containing implementations for parallel incomplete factorizations (ParILU) for CPU & GPU architectures.
- MAGMA-sparse release containing implementations for parallel incomplete factorizations based on thresholding (ParILUT) for CPU & GPU architectures.
- MAGMA-sparse release containing implementations for parallel sparse triangular solves for CPU & GPU architectures.

Impact

- Cross-Platform implementations for parallel ILU preconditioning, parallel threshold ILU preconditioning, and sparse triangular solves available in MAGMA-sparse part of the xSDK software ecosystem.
- Presented the performance results of the first threshold ILU preconditioner (ParILUT) at the IPDPS 2019 conference and the JLESC 2019 meeting in Knoxville.

Performance results for cross-platform ParILUT

Current efforts

- Exploring the use of faster preconditioner pattern generation via • Machine Learning.
- Looking for ECP application projects amenable to threshold-ILU preconditioning.
- Integrating ParILU / ParILUT into the Ginkgo software ecosystem featuring Continuous Benchmarking (CB).

Deliverables The ParILU and ParILUT preconditioners for architectures supporting OMP or CUDA are available in the MAGMA repository: https://bitbucket.org/icl/magma/src Performance results on the SUMMIT architecture were presented at the IPDPS conference in Rio de Janeiro, Brazil: https://hartwiganzt.github.io/slides/IPDPS ParILUT.pdf

Data Model Support in ADIOS for Visualization and Analysis

ECP WBS 2.3.4.09 STDM11-ADIOS2

PI Scott Klasky, ORNL

Members ORNL, Kitware, Georgia Tech, Rutgers, LBNL

Scope and objectives

- Describe variables in ADIOS using VTK XML format
- Develop the VTK::IOADIOS2 module to enable seamless visualization in Paraview
- Support ECP products:
 - MEUMAPPS (ExaAM) for ImageData
 - MFEM (CEED) and UnstructuredGrid

Impact

- Applications using ADIOS can visualize their data in parallel using a simple string descriptor base on the well-documented VTK File data models
- Leverage current scalable I/O capabilities in ADIOS with VTK (and potentially VTK-m) for insitu visualization using a pluggable ecosystem delivered by the ECP data analysis and visualization software stack.

Visualization of ADIOS BP files

Seamless visualization in Paraview using the VTK::IOADIOS2 module and VTK XML description in BP Files

Project accomplishment

- MEUMAPPS (ExaAM) and MFEM (CEED) users can now visualize their data by adding simple XML descriptors either during the data generation or after
- The VTK::IOADIOS2 provides a reusable architecture that expands these capabilities and take advantage or the rich VTK ecosystem

Deliverables VTK::IOADIOS2 module merge request: https://gitlab.kitware.com/vtk/vtk/merge_requests/5613

ECP ST Review Focus

- Describe your overall approach followed by specific activities in preparation for exascale platforms.
 - Include a description of your overall strategy for your L4 project.
 - Include description of pre-exascale environments you are using.
 - Include status and any results from pre-exascale environments (GPU porting, use of Summit) that illustrate your strategy.
 - Include discussion of your major challenges.
- Describe your overall approach and describe the status of your client integration efforts.
 - Include a description of your overall strategy for your L4 project.
 - Include specific client interactions that illustrate your strategy.
 - Include recent integration progress.
 - Include discussion of your major challenges.

Detailed instructions for each L3 area: PMR

- Perlmutter, Aurora and Frontier represent very different node programming environments. As appropriate, how is your
 project addressing this challenge, including node performance, performance portability, performance of MPI+X and
 preparation for further heterogeneity?
- What is the path for your capabilities to realize sustainability in the software ecosystem?
- How is your work impacting vendor capabilities?

E4S Collaborations

How E4S and OpenHPC Work Together

• OpenHPC

- Provides installation and base components for installing and running cluster to supercomputer
- Provides generic binaries for all systems (working on targeted

•E4S

- Built from scratch for your system
- Focus on upper layers of system software stack (libraries, runtimes, etc.)
- Targeted for capability-class machines

1st Workshop on NSF and DOE High Performance Computing Tools

Dates: July 10-11, 2019 Location: 312 Lillis, University of Oregon, Eugene, OR 97403

NSF Collaborations

Workshop A	Agenda	
Wednesday, Ju	uly 10, 2019: 312 Lillis, University of Oregon	
8:00am - 8:45am:	Registration	
8:45am - 9:15am:	Welcome, introductions, David Conover (VPRI, UO), Vipin Chaudhary (Program Director, CISE/OAC, NSF), Jo Science, LBL)[slides] [video]	nathan Carter (Deputy Director, Software Technology, ECP, and Deputy for
9:15am - 10:30am:	Spack tutorial - Todd Gamblin and Greg Becker (LLNL) [slides] [video]	
10:30am - 11:00am:	Break	
11:00am - 12:30pm:	Spack tutorial - Todd Gamblin (LLNL) [video]	https://oaciss.uoregon.edu/NSFDOE19/agenda.html
12:30pm - 2:00pm:	Lunch break: 440 Lillis	
2:00pm - 3:30pm:	Spack tutorial - Todd Gamblin and Greg Becker (LLNL) [video]	
3:30 pm - 4:00pm:	Break	
4:00pm - 5:30pm:	Spack tutorial - Todd Gamblin and Greg Becker (LLNL) [video 1] [video 2] [video 3] [video 4] [video 5]	
6:30pm - 9:30pm:	Working dinner at the Jordan Schnitzer Museum of Art, UO	
Thursday, July	/ 11, 2019: 312 Lillis, University of Oregon	
9:00am - 9:30am:	Unifying Software Distribution in ECP - Todd Gamblin (LLNL) [slides] [video]	
9:00am - 9:30am:	Containers for HPC - Andrew Younge (Sandia National Laboratories, NM) [slides] [video]	
9:30am - 10:00am	n: Software deployment at DOE facilities - David Montoya (Los Alamos National Laboratory, NM) [slides] [video	
10:00am - 10:30ai	m: E4S - Sameer Shende (University of Oregon) [slides] [video]	
10:30am - 11:00ar	m: Break	
11:00am - 11:30ar	m: Overview of software architecture - Todd Gamblin, LLNL	
11:30am - 12:30pr	m: Hands-on, applying Spack to applications.	
12:30pm - 2:00pm	n: Lunch: 440 Lillis	
2:00pm - 3:30pm:	Hands-on, applying Spack to applications.	
3:30pm - 4:00pm:	Break	
4:00pm - 5:00pm:	Hands-on, Spack and E4S.	
5:00pm - 6:15pm:	Closing remarks, planning for the next workshop - Jonathan Carter (Lawrence Berkeley National Laboratory)	
6:30pm:	Dinner at Excelsion Inn	

Extending Collaborations

- E4S/SDK collaborations make sense with:
 - HPC open source development projects:
 - deal.II (NSF math library),
 - PHIST (DLR Germany library).
 - Application teams in search of open source software foundation:
 - NSF science teams.
 - NOAA, others.
 - Commercial open source packagers/distributors
 - OpenHPC.
 - HPC systems vendors.
 - HPC systems facilities:
 - SDSC, TACC, others.

E4S: Building on top of previous efforts

- E4S did not emerge from nothing.
- Leveraging the work of many others.
- HPC Linux: Work done at U of Oregon, and at ParaTools.
- IDEAS-Classic: xSDK the original SDK continuing under ECP.
- Spack Pre-dates E4S.

E4S Wrap Up

Software Technology Ecosystem

ECP ST Open Product Integration Architecture

E4S Basics

What E4S is not

• A closed system taking contributions only from DOE software development teams.

• A monolithic, take-it-or-leave-it software behemoth.

• A commercial product.

• A simple packaging of existing software.

What E4S is

- Extensible, open architecture software ecosystem accepting contributions from US and international teams.
- Framework for collaborative open-source product integration.
- A full collection if compatible software capabilities and
- A manifest of a la carte selectable software capabilities.
- Vehicle for delivering high-quality reusable software products in collaboration with others.
- The conduit for future leading edge HPC software targeting scalable next-generation computing platforms.
- A hierarchical software framework to enhance (via SDKs) software interoperability and quality expectations.

Final Remarks

- Software Development Kits (SDKs) provide a new software aggregation mechanism:
 - Do not significantly increase the number of software products in the HPC ecosystem.
 - Provide intermediate build, install and test targets for E4S (reducing complexity).
 - Enable development teams to improve and standardize look-and-feel, best practices, policies.
 - Improve interoperability, interchangeability of similar products.
 - Fosters communities of developers in a cooperative/competitive environment.
 - Provides integration point for SDK-compatible, non-ECP products.
- The Extreme-scale Scientific Software Stack (E4S) provides a complete HPC software stack:
 - Does not significantly increase the number of software products in the HPC ecosystem.
 - Provides a coordinated approach to building, installing and testing HPC software.
 - Tests some products with a subset of configurations on a subset of platforms.
 - Improves stability of the ST stack so that any subset is more stable.
 - Provides a complete software stack, including non-ECP products.
- The SDK/E4S architecture is open:
 - Enables light-weight coordinated collaboration among open source software teams.
 - ECP seeks collaboration: libraries/tools and SDKs, facilities and E4S.