

xSDK Foundations: Toward an Extreme-scale Scientific Software Development Kit

Roscoe Bartlett¹, Irina Demeshko², Todd Gamblin³, Glenn Hammond¹, Michael Heroux¹, Jeffrey Johnson⁴, Alicia Klinvex¹, Xiaoye Li⁵, Lois Curfman McInnes⁶, J. David Moulton², Daniel Osei-Kuffuor³, Jason Sarich⁶, Barry Smith⁶, Jim Willenbring¹, Ulrike Meier Yang³

© The Authors 2017. This paper is published with open access at SuperFri.org

Extreme-scale computational science increasingly demands multiscale and multiphysics formulations. Combining software developed by independent groups is imperative: no single team has resources for all predictive science and decision support capabilities. Scientific libraries provide high-quality, reusable software components for constructing applications with improved robustness and portability. However, without coordination, many libraries cannot be easily composed. Namespace collisions, inconsistent arguments, lack of third-party software versioning, and additional difficulties make composition costly.

The Extreme-scale Scientific Software Development Kit (xSDK) defines community policies to improve code quality and compatibility across independently developed packages (hypre, PETSc, SuperLU, Trilinos, and Alquimia) and provides a foundation for addressing broader issues in software interoperability, performance portability, and sustainability. The xSDK provides turnkey installation of member software and seamless combination of aggregate capabilities, and it marks first steps toward extreme-scale scientific software ecosystems from which future applications can be composed rapidly with assured quality and scalability.

Keywords: xSDK, Extreme-scale scientific software development kit, numerical libraries, software interoperability, sustainability.

1. Software Challenges for Extreme-scale Science

Extreme-scale architectures provide unprecedented resources for scientific discovery. At the same time, the computational science and engineering (CSE) community faces daunting productivity and sustainability challenges for parallel application development [1, 12, 13, 21]. Difficulties include increasing complexity of algorithms and computer science techniques required by coupled multiscale and multiphysics applications. Further complications come from the imperative of portable performance in the midst of dramatic and disruptive architectural changes on the path to exascale, the realities of large legacy code bases, and human factors arising in distributed multidisciplinary research teams pursuing leading edge parallel performance. Moreover, new architectures require fundamental algorithm and software refactoring, while at the same time demand is increasing for greater reproducibility of simulation and analysis results for predictive science.

This confluence of challenges brings with it a unique opportunity to fundamentally change how scientific software is designed, developed, and sustained. The demands arising from so many challenges force the CSE community to consider a broader range of potential solutions. It is this setting that makes possible a collaborative effort to establish a scientific software ecosystem of

¹Sandia National Laboratories, USA

²Los Alamos National Laboratory, New-Mexico, USA

³Lawrence Livermore National Laboratory, Livermore, USA

⁴Salesforce, San-Francisco, USA

⁵Lawrence Berkeley National Laboratory, Berkeley, USA

⁶Argonne National Laboratory, Lemont, USA

reusable libraries and community policies to guide common adoption of practices, tools, and infrastructure. Incremental change is not a viable option, so migration to a new model for CSE software is possible.

The xSDK has emerged as a first step toward a new ecosystem, where application codes are composed via interfaces from a common base of reusable components more than they are developed from a clean slate or derived from monolithic code bases. To the extent that this compositional approach can be reliably used, new CSE applications can be created more rapidly, with greater robustness and scalability, by smaller teams of scientists, enabling them to focus more attention on obtaining science results than on the incidentals of their computing environment.

1.1. Related Work

The scientific software community has a rich tradition of defining *de facto* standards for collections of capabilities. EISPACK [11, 23], LINPACK [7], BLAS [8, 9, 16, 17], and LAPACK [2] delivered a sound foundation for numerical linear algebra in libraries and applications. Commercial entities such as the Numerical Algorithms Group (NAG) [26], the Harwell Subroutine Library (HSL) [30] and IMSL have provided high quality, unified software capabilities to users for decades.

More recently, the TOPS [14], ITAPS [5], and FASTMath [6] SciDAC institutes brought together developers of large-scale scientific software libraries. While these libraries were independently developed by distinct teams and version support lacked coordination, the collaborations sparked exchange of experiences and discussion of practices that avoided potential pitfalls and facilitated the combined use of the libraries [19] as needed by scientific teams. Prior efforts to provide interoperability between solver libraries can be found in PETSc [3], which allows users to access libraries such as hypre [25] and SuperLU [28] by using the PETSc interface, sparing users the effort to rebuild their problems through hypre’s or SuperLU’s interfaces. Trilinos [31], a collection of self-contained software packages, also provides ways for users to gain uniform access to third-party scientific libraries.

2. xSDK Vision

The complexity of application codes is steadily increasing due to more sophisticated scientific models and the continuous emergence of new high-performance computers, making it crucial to develop software libraries that provide needed capabilities and continue to adapt to new computer architectures. Each library is complex and requires different expertise. Without coordination, and in service of distinct user communities, this circumstance has led to difficulties when building application codes that use 8 or 10 different libraries, which in turn might require additional libraries or even different versions of the same libraries.

The xSDK represents a different approach to coordinating library development and deployment. Prior to the xSDK, scientific software packages were cohesive with a single team effort, but not across these efforts. The xSDK goes a step further by developing community policies followed by each independent library included in the xSDK. This policy-driven, coordinated approach enables independent development that still results in compatible and composable capabilities.

The initial xSDK project is the first step toward a comprehensive software ecosystem. As shown in Figure 1, the vision of the xSDK is to provide infrastructure for and interoperability of a collection of related and complementary software elements—developed by diverse, indepen-

dent teams throughout the high-performance computing (HPC) community—that provide the building blocks, tools, models, processes, and related artifacts for rapid and efficient development of high-quality applications. Our long-term goal is to make the xSDK a turnkey standard software ecosystem that is easily installed on common computing platforms, and can be assumed as available on any leadership computing system in the same way that BLAS and LAPACK are available today.

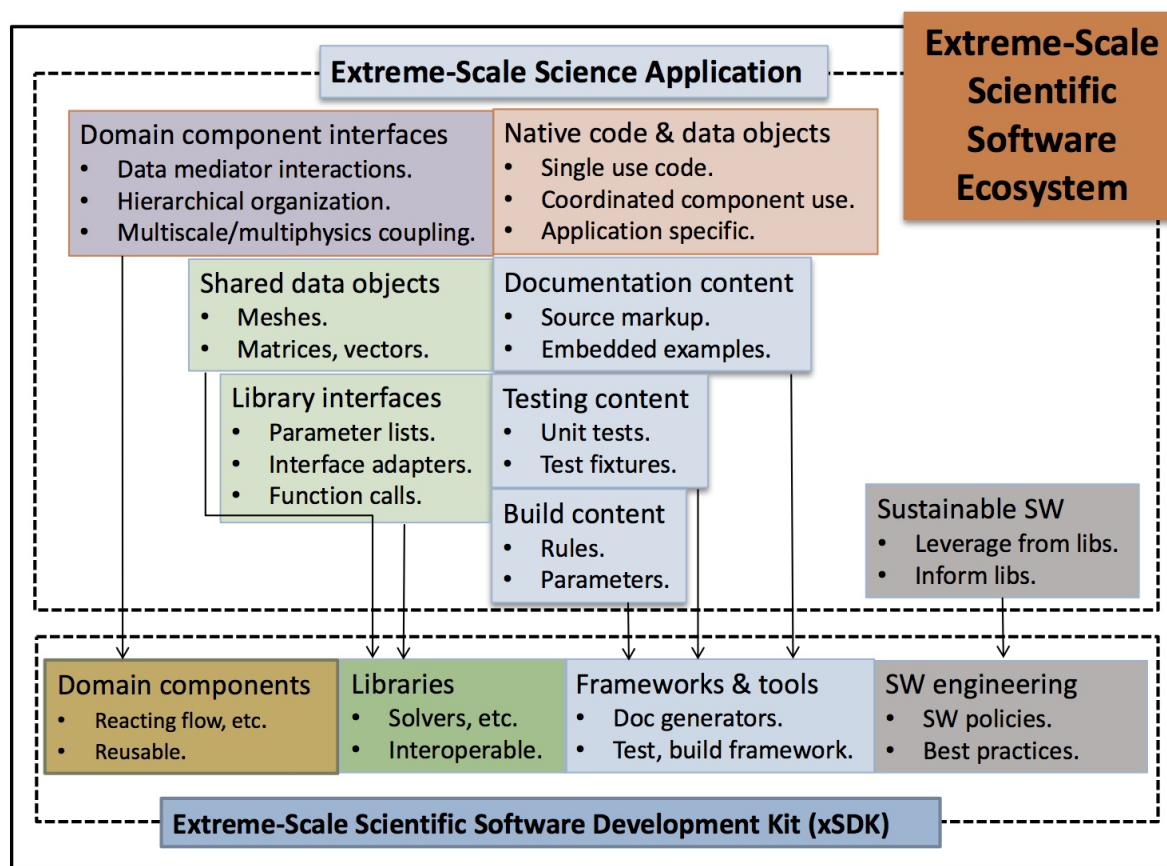


Figure 1. The xSDK intends to provide the foundation for a modern extreme-scale scientific software ecosystem, where application development is accomplished by composition of high-quality, reusable software components rather than by tangential use of libraries. Application developers produce a small portion of custom code that expresses the particular purpose of the software and then gain the bulk of functionality by parameterized use of xSDK components and libraries, which are developed by diverse, independent groups throughout the community. xSDK frameworks for documentation, testing, and code quality, as well as established software policies and best practices, can be adapted and adopted as appropriate by the application developers to provide compatible, high-quality, and sustainable software. As we move toward this new ecosystem, application development times from first concept to scalable production code should drop dramatically. Success hinges on the quality, interoperability, usability, and diversity of xSDK capabilities and our ability to deliver the xSDK to domain scientists

2.1. Elements of an Extreme-scale Scientific Software Ecosystem

Rapid, efficient production of high-quality, sustainable applications is best accomplished using a rich collection of reusable libraries, tools, lightweight frameworks, and defined software methodologies, developed by a community of scientists who are striving to identify, adapt, and adopt best practices in software engineering. Although the software engineering community has ongoing debate about the precise meaning of terms, we define the basic elements of a scientific software ecosystem to include:

- **Library:** High-quality, encapsulated, documented, tested and multi-use software that is incorporated into the application and used as native source functionality. Libraries can provide *control inversion* via abstract interfaces, call-backs, or similar techniques such that user-defined functionality can be invoked by the library, e.g., a user-defined sparse matrix multiplication routine. Libraries can also provide factories that facilitate construction of specific objects that are related by a base type and later used as an instance of the base type. Libraries can include domain-specific software components that are designed to be used by more than one application.
- **Domain component:** Reusable software that is intended for modest reuse across applications in the same domain. Although this kind of component is a library, the artifacts and processes needed to support a component are somewhat different than for a broadly reusable library.
- **Framework:** A software environment that implements specific design patterns and permits the user to insert custom content. Frameworks include documentation, build (compilation), and testing environments. These frameworks are lightweight and general purpose. Other frameworks, such as multiphysics, are considered separately, built on top of what we describe here.
- **Tool:** Software that exists outside of applications, used to improve quality, efficiency, and cost of developing and maintaining applications and libraries.
- **Software development kit (SDK):** A collection of related and complementary software elements that provide the building blocks, libraries, tools, models, processes, and related artifacts for rapid and efficient development of high-quality applications.

Given these basic elements, we define an application code as the following composition:

- **Native data and code:** Every application will have a primary routine (often a main program) and its own collection of source code and private data. Historically, applications have been primarily composed of native source and data, using libraries for a small portion of functionality, such as solvers. We foresee a decrease in the amount of native code required to develop an application by extracting and transforming useful native code into libraries and domain components, making it available to other applications.
- **Component and library function calls:** Some application functionality is provided by invoking library functions. We expect to increase usage of libraries as a part of our efforts.
- **Library interface adapters:** Advanced library integration often involves invoking the control inversion facilities of the library in order to incorporate application-specific knowledge. In the case of sensitivity analysis, embedded optimization, and related analyses, control inversion via these adapters is essential in order to permit the solver to invoke the application with specific input data.

- **Component and library parameter lists:** Libraries tend to provide a broad collection of functionality for which parameters must be set.
- **Shared component and library data:** Most libraries require the user to provide non-trivial data objects, such as meshes or sparse matrices, and may provide functions to assist the application in constructing these objects. Unlike parameter list definitions, which represent a narrow interface dependency between the application and library, application-library data interfaces can be very complicated.
- **Documentation, build, and testing content:** The application-specific text, data, and source used by the documentation, build, and testing frameworks to produce the derived software documentation, compilation, and test artifacts.

3. xSDK Approach

The xSDK approach to developing software has two distinguishing features from previous efforts in the scientific computing community:

- **Peer-to-peer interoperability:** Some previous efforts⁷ attempted to use additional abstraction layers that would hide differences in the underlying packages. The xSDK approach uses the existing extensibility features of the libraries to enable peer-to-peer access of capabilities at various levels of interoperability through the native interfaces of the packages. For example, if a user has already integrated PETSc data structures into their code, the xSDK approach preserves that approach, but permits use of capabilities in hypre, SuperLU, and Trilinos with PETSc.
- **Software policies:** Most existing scientific software efforts rely on close collaboration of a single team in order to assure that collective efforts are compatible and complementary. The xSDK relies instead on *policies* that promote compatibility and complementarity of independently developed software packages. By specifying only certain expectations for how software is designed, implemented, documented, supported, and installed, the xSDK enables independent development of separate packages, while still ensuring complementarity and composability.

The xSDK can assure interoperability and compliance with community policies because the leaders and developers of xSDK packages are members of the xSDK community. If interface changes are required in a package or a version of a third-party solver needs to be updated, these changes will be made in the member package. For example, in order for Trilinos and PETSc to use the same version of SuperLU and hypre, the Trilinos and PETSc developers commit to agreeing on changes to Trilinos and PETSc that are needed for compatibility. Similarly, changes to interfaces for interoperability and inversion of control (see the next Section 3.1) are done within the xSDK packages, and regularly tested for regressions. xSDK interoperability is possible because of the commitment of xSDK member package development teams.

⁷A notable example is the Equation Solver Interface (ESI), which defined an abstraction layer to present a common client interface to distinct software products. The challenge of this approach is that the unique features of the underlying products were difficult to access. The very use of a common abstraction reduced the usability of these products.

3.1. xSDK Library Interoperability

A fundamental objective of the xSDK project is to provide interoperability layers among hypre, PETSc, SuperLU, and Trilinos packages, as appropriate, with the ultimate goal of making all mathematically meaningful interoperabilities possible in order to fully support exascale applications.

Software library interoperability refers to the ability of two or more libraries to be used together in an application, without special effort by the user [18]. For simplicity, we discuss interoperability between two libraries; extension to three or more libraries is conceptually straightforward. Depending on application needs, various levels of interoperability can be considered:

- *Interoperability level 1*: both libraries can be used (side by side) in an application
- *Interoperability level 2*: both libraries can exchange data (or control data) with each other
- *Interoperability level 3*: each library can call the other library to perform unique computations

The simplest case (interoperability level 1) occurs when an application needs to call two distinct libraries for different functionalities (for example, an MPI library for message-passing communication and HDF5 for data output). As discussed in [19, 20], even this basic interoperability requires consistency among libraries to be used in the same application, in terms of compiler, compiler version/options, and third-party capabilities. If both libraries have a dependency on a common third party, the libraries must be able to use a single common instance of it. For example, more than one version of the popular SuperLU linear solver library exists, and interfaces have evolved. If two libraries both use SuperLU, they must be able to work with the same version of SuperLU. In practice, installing multiple independently developed packages together can be a tedious trial-and-error process. The definition and implementation of xSDK community policies standards have overcome this difficulty for xSDK-compatible packages.

Interoperability level 2 builds on level 1 by enabling conversion, or encapsulation, and exchange of data between libraries. This level can simplify use of libraries in sequence by an application. In this case, the libraries themselves are typically used without internal modification to support the interoperability. Future work on node-level resource management is essential to support this deeper level of software interoperability for emerging architectures.

Interoperability level 3 builds on level 2 by supporting the use of one library to provide functionality on behalf of another library. This *integrated execution* provides significant value to application developers because they can access capabilities of additional libraries through the familiar interfaces of the first library.

The remainder of this section discusses proposed work on integrated execution, where our guiding principles are to provide interoperability that is intuitive and easy to use, and to expose functionality of each library where feasible.

Control inversion. Interoperability between two (or more) existing library components can be achieved by one of two basic mechanisms: (i) create an abstraction layer that sits on top of both components to act as an intermediary between the user and both components or (ii) permit users to write directly to the interface of one component and provide peer-level interoperability between the two components. For example, consider the matrix construction capabilities in PETSc and Trilinos. Both libraries provide extensive support for piecewise construction of sparse matrices, as needed for building objects in applications based on finite elements/volumes/differences. It would be possible, in principle, to create a top-level abstraction

layer that could be used to build a sparse matrix or other data objects for PETSc or Trilinos, depending on an input option to select either target. Alternatively, the user can construct the data object by using the PETSc or Trilinos functions directly, and then we can create adapters in Trilinos and PETSc to wrap the respective matrix object and make it behave like one of its own.

Although the first approach may seem attractive, it is difficult to develop in a sustainable and effective way. PETSc and Trilinos data object construction processes are targeted to specific programming, language, and usage models. The differences in approach may appear small, but are very important in terms of developer productivity, code portability, and expressiveness. Any abstraction layer that would sit on top of both would discard the simplicity of one approach or the expressiveness of the other.

Peer-to-peer interoperability is much more attractive than a general abstraction layer. The xSDK libraries have mechanisms to work with or easily transform existing data objects that were built outside their own construction processes. For example, a PETSc sparse matrix can be used within Trilinos, without copying, by using an adapter class. A similar approach can work with a Trilinos matrix used by PETSc.

The hypre and SuperLU libraries do not directly support control inversion in the same way as PETSc and Trilinos, but do advertise their input data structures such that PETSc and Trilinos can construct compatible data structures that are passed to hypre and SuperLU without copying.

The current release of xSDK does not support all possible opportunities for interoperability. Level 1 interoperability is complete within the current xSDK. Level 2 interoperability is partial, with Trilinos being able to accept PETSc data structures. Level 3 interoperability is also partially available with PETSc and Trilinos able to call use hypre and SuperLU.

3.2. xSDK Community Policies

In [19, 20] various software quality engineering practices for ‘smart libraries’ are discussed that, when followed, can alleviate generation of an application executable that depends on many libraries, reduce mistakes in how to use these libraries, and provide help to users to identify and correct errors when they occur.

The first xSDK release demonstrated the impact of defining xSDK community policies, including standard GNU autoconf and CMake options to simplify the combined use, portability, and sustainability of independently developed software packages (hypre, PETSc, SuperLU, and Trilinos) and provide a foundation for addressing broader issues in software interoperability and performance portability.

xSDK community package policies [22], briefly summarized in Figure 2, are a set of minimum requirements (including topics of configuring, installing, testing, MPI usage, portability, contact and version information, open source licensing, namespacing, and repository access) that a software package must satisfy in order to be considered xSDK compatible. The designation of xSDK compatibility informs potential users that a package can be easily used with others.

xSDK community installation policies [4] help make configuration and installation of xSDK software and other HPC packages as efficient as possible on common platforms, including standard Linux distributions and Mac OS X, as well as on target machines currently available at DOE computing facilities (ALCF, NERSC, and OLCF) and eventually on new exascale platforms.

xSDK Mandatory Policies

Must:

- M1. Support xSDK community GNU Autoconf or CMake options [4].
- M2. Provide a comprehensive test suite.
- M3. Employ user-provided MPI communicator.
- M4. Give best effort at portability to key architectures.
- M5. Provide a documented, reliable way to contact the development team.
- M6. Respect system resources and settings made by other previously called packages.
- M7. Come with an open source license.
- M8. Provide a runtime API to return the current version number of the software.
- M9. Use a limited and well-defined symbol, macro, library, and include file name space.
- M10. Provide an accessible repository (not necessarily publicly available).
- M11. Have no hardwired print or IO statements.
- M12. Allow installing, building, and linking against an outside copy of external software.
- M13. Install headers and libraries under `<prefix>/include/` and `<prefix>/lib/`.
- M14. Be buildable using 64 bit pointers. 32 bit is optional.

xSDK Recommended Policies

Should:

- R1. Have a public repository.
- R2. Possible to run test suite under valgrind in order to test for memory corruption issues.
- R3. Adopt and document consistent system for error conditions/exceptions.
- R4. Free all system resources it has acquired as soon as they are no longer needed.
- R5. Provide a mechanism to export ordered list of library dependencies.

Figure 2. xSDK community policies specify expectations that any software library or framework (henceforth referred to as package) must satisfy in order to be xSDK compatible. The designation of a package being xSDK compatible informs potential users that the package can be easily used with other xSDK libraries and components and thus helps to address issues in long-term sustainability and interoperability among packages

Community policies for the xSDK promote long-term sustainability and interoperability among packages, as a foundation for supporting complex multiphysics and multiscale ECP applications. In addition, because new xSDK packages will follow the same standard, installation software and package managers (for example, Spack [10]) can easily be extended to install many packages automatically.

Figure 3 illustrates a new *Multiphysics Application C*, built from two complementary applications that can readily employ any libraries in the xSDK (hypre, PETSc, SuperLU, and Trilinos, shown in green). Application domain components are represented in orange. Of particular note is Alquimia [24], a domain-specific interface that support uniform access to multiple biogeochemistry capabilities, including PFLOTRAN [27]. The arrows among the xSDK libraries indicate current support for a package to call another to provide scalable linear solvers functionality on its behalf. For example, *Application A* could use PETSc for an implicit-explicit

time advance, which in turn could interface to SuperLU to solve the resulting linear systems with a sparse direct solver. *Application B* could use Trilinos to solve a nonlinear system, which in turn could interface to hypre to solve the resulting linear systems with algebraic multigrid. Of course, many other combinations of solver interoperability are also possible. The website <https://xsdk.info/example-usage> and [15] provide examples of xSDK usage, including interoperability among linear solvers in hypre, PETSc, SuperLU, and Trilinos.

3.3. xSDK Coordinated Software Releases and Current Status

The first critical step in producing the initial release of the xSDK in April 2016 was to define what the release should look like. On the loosely coupled end of the spectrum, one possibility would have been to certify that specific release versions of the different packages are compatible with one another and not coordinate the distribution of the release beyond that.

On the other end of the spectrum were possibilities such as a common test or even build infrastructure. The xSDK team decided on a strategy that provided a single point of distribution for all xSDK component packages, but did not force a common infrastructure on the packages, beyond the agreed upon community policies.

The chosen distribution mechanism was, at its core, the existing PETSc distribution mechanism [3]. This was a stable, well-supported option that required a relatively small amount of effort to extend for the xSDK use case, because it already supported the majority of xSDK component packages, and very little additional ongoing maintenance beyond what the PETSc team was already doing. The latter was a key consideration because, when possible, the xSDK team actively avoids solutions that create long-term maintenance beyond what is needed for the component packages.

In addition to creating and updating interfaces between xSDK packages to provide new functionality, significant work was done to make it possible to install all of the xSDK packages together. For example, PETSc and Trilinos depended on different versions of SuperLU, and no single version of SuperLU could be used with both the PETSc and Trilinos SuperLU interfaces enabled.

As the coordinated release effort began (initially work began with package-to-package compatibility and interface efforts), the xSDK installer was used to test release versions of the component packages together to avoid the churn of the various development versions. Testing was set up by multiple xSDK component package teams using their own resources. Again, sustainability was a focus. Having the individual teams manage separate test builds in which each team had a vested interest was a better choice than a centralized effort that would have no clear owner in the absence of an xSDK-level funding source. The test builds included release

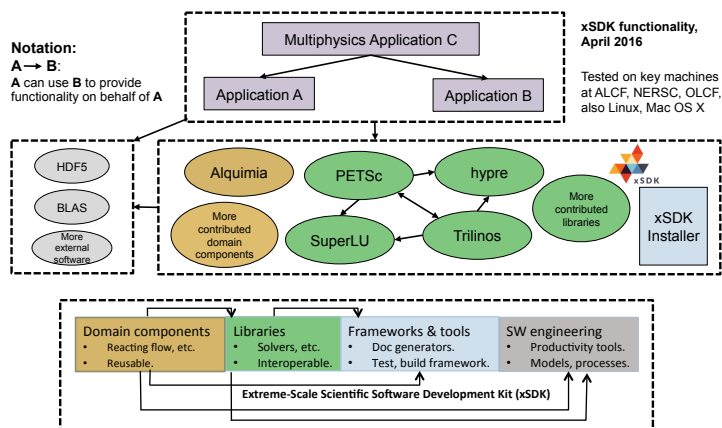


Figure 3. xSDK schematic diagram

and development versions, but leading up to the release, primary focus was given to the release testing.

The existing release processes of the various xSDK component packages varied greatly in terms of testing and overall rigor. However, since each component package had a release process that the individual development teams had determined was sufficient for their needs, the decision was made to focus on requirements outside of typical release requirements. Specifically, this involved providing the name of a branch to use for release candidate testing, setting up tests for xSDK-level build configurations most relevant to the component package, and being responsive to any issues found.

In addition to the testing conducted by each package team, the xSDK 0.1 release was ported to three target platforms at three different computing facilities: Mira at ALCF, Edison at NERSC, and Titan at OLCF. One developer was primarily responsible for each of the three porting efforts, and those people coordinated with other xSDK developers and component package team developers to resolve porting issues as necessary.

The official tag for the initial xSDK release was chosen to be v0.1.0. After the initial release, a patch release, v0.1.1, was completed. The versions of the component packages used for this subsequent release were either the same version used for the initial 0.1 release, or a patch release of the component package based on the release used for the initial release. According to xSDK release policy, only patch-level updates for component packages are to be used for xSDK patch releases, and only patch or minor release version updates for component packages are to be used for xSDK minor releases.

Prior to selecting the PETSc distribution capability for the xSDK 0.1 release, Spack [10] was also seriously considered. Spack is a multi-platform package manager that supports a variety of compilers, libraries, and applications, as well as the installation of multiple concurrent versions and software configurations. Because of the increasing popularity and robustness of Spack, and the need to expand the xSDK to include several additional component packages, the xSDK team decided to use Spack to build an alpha release version of the xSDK, to be released in early 2017. Going forward, the intent is to use Spack as the principal supported xSDK distribution capability.

4. Next Steps

The first xSDK releases focused on discovery of collaboration models and community building among four of the major open source scientific library projects in the international scientific computing community: hypre, PETSc, SuperLU and Trilinos. The xSDK project will continue over the next few years under the United States Department of Energy Exascale Computing Project (ECP) [29].

Our efforts so far have established a baseline for expanding the xSDK scope under ECP funding in several important directions:

1. **Include more libraries:** The xSDK will expand to include all library efforts under ECP funding. Specifically, widely used libraries such as SUNDIALS and Magma will become xSDK compatible, as will new efforts that address the performance challenges of exascale computing platforms.
2. **Further refine and expand community policies:** While the current xSDK community policies, summarized in Figure 2, are extremely useful as a mechanism to improve interoperability and compatibility of independently developed scientific libraries, we believe we can

further refine and expand these policies to better assure software quality and further realize the scientific software ecosystem sketched in Figure 1.

3. **Include more domain components:** As described in Section 1, the vision of the xSDK is to create a software ecosystem where new scientific applications are composed via interfaces from a common base of reusable domain components and libraries. We will work with science teams to identify opportunities for creating collections of domain components for their communities.
4. **Explore the use of community installation tools, including Spack:** While the extended PETSc installer has been very useful for establishing xSDK as a unified project, Spack [10] promises to provide a tool that serves and is supported by a larger community, making it very appealing as the principal long term installation tool for xSDK libraries.
5. **Process control transfer interfaces:** The ever-increasing use of concurrency within the top-level MPI processes requires that computational resources used by an application or library can be transferred to another library. Transfer of these resources is essential for obtaining good performance. The xSDK project will develop interfaces to support sharing and transfer of computational resources.

Conclusions

The extreme-scale scientific community faces numerous disruptive challenges in the coming decade. Fundamental limits of physics are forcing changes that dramatically impact all system layers from architecture to application software design. These disruptive changes drive us to move beyond incremental change in scientific application design and implementation. Establishing a scientific software ecosystem that focuses more on the composition of scalable, reusable components for application software development can provide an attractive alternative and the xSDK is the first step toward the ecosystem described in Figure 1.

Community policies for the xSDK promote long-term sustainability and interoperability among packages, as a foundation for supporting complex multiscale and multiphysics applications. The designation of xSDK compatibility informs potential users that a package can be easily used with other xSDK libraries and components. In addition, because new xSDK packages will follow the same standard, installxSDK and package managers can easily be extended to install many packages automatically.

Interoperability of xSDK member packages, when wrapped with adequate testing, enables a sustainable coupling of capabilities that enable applications to use xSDK member packages as a cohesive suite. The first xSDK releases demonstrate the impact of xSDK community policies, testing and examples to simplify the combined use, interoperability, and portability of independently developed software packages, establishing the first step toward realizing an extreme-scale scientific software ecosystem.

Acknowledgements

This material is based upon work funded by the U.S. Department of Energy Office of Science, Advanced Scientific Computing Research and Biological and Environmental Research programs. We thank program managers Thomas Ndousse-Fetter, Paul Bayer, and David Lesmes for their support. The work of ANL authors is supported by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357. The work of LBNL authors is partially supported

by the Director, Office of Science, Office of Advanced Scientific Computing Research of the US Department of Energy under contract no. DE-AC02-05CH11231.

Prepared by LLNL under Contract DE-AC52-07NA27344. Work of LANL authors is funded by the Department of Energy at Los Alamos National Laboratory under contract DE-AC52-06NA25396. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DEAC04-94AL85000.

This paper is distributed under the terms of the Creative Commons Attribution-Non Commercial 3.0 License which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is properly cited.

References

1. Almgren, A., DeMar, P., Vetter, J., et al.: DOE Exascale Requirements Review for Advanced Scientific Computing Research (2016), u.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, see <http://exascale.org/ascr/>, Sept 27–29, 2016
2. Anderson, E., Bai, Z., Bischof, C., Blackford, L.S., Demmel, J., Dongarra, J.J., Du Croz, J., Hammarling, S., Greenbaum, A., McKenney, A., Sorensen, D.: LAPACK Users' Guide (Third Ed.) (1999), DOI: 10.1137/1.9780898719604
3. Balay, S., Adams, M.F., Brown, J., Brune, P., Buschelman, K., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Rupp, K., Smith, B.F., Zhang, H.: PETSc Web page. <http://www.mcs.anl.gov/petsc>
4. Bartlett, R., Sarich, J., Smith, B., Gamblin, T., xSDK developers: xSDK community installation policies: GNU Autoconf and CMake options (2016), DOI: 10.6084/m9.figshare.4495133
5. Diachin (PI), L.: Interoperable Technologies for Advanced Petascale Simulations, SciDAC Institute. <http://www.scidac.gov/math/ITAPS.html>, accessed: 2017-02-15
6. Diachin (PI), L.: SciDAC-3 Frameworks, Algorithms, and Scalable Technologies for Mathematics (FASTMath) Institute. <http://www.fastmath-scidac.org/>, accessed: 2017-02-15
7. Dongarra, J.J., Bunch, J., Moler, C., Stewart, G.: LINPACK Users' Guide. SIAM Pub. (1979), DOI: 10.1137/1.9781611971811
8. Dongarra, J.J., Du Croz, J., Hammarling, S., Duff, I.: A set of level 3 basic linear algebra subprograms. ACM Trans. Math. Software 16(1), 1–17 (March), DOI: 10.1145/77626.79170
9. Dongarra, J., DuCroz, J., Hammarling, S., Hanson, R.: An extended set of Fortran basic linear algebra subprograms. ACM Trans. Math. Software 14 (1988), DOI: 10.1145/42288.42291
10. Gamblin, T., LeGendre, M.P., Collette, M.R., Lee, G.L., Moody, A., de Supinski, B.R., Futral, W.S.: The Spack Package Manager: Bringing order to HPC software chaos. In: Supercomputing 2015 (SC'15). Austin, Texas (November 15-20), LLNL-CONF-669890

11. Garbow, B.S., Boyle, J.M., Dongarra, J.J., Moler, C.B.: Matrix Eigensystem Routines – EISPACK Guide Extension, Lecture Notes in Computer Science, vol. 51. Springer–Verlag, New York (1977), DOI: 10.1007/3-540-08254-9
12. Heroux, M.A., Allen, G., et al.: Computational Science and Engineering Software Sustainability and Productivity (CSESSP) Challenges Workshop Report (September 2016), <https://www.nitrd.gov/PUBS/CSESSPWorkshopReport.pdf>
13. Johansen, H., McInnes, L.C., Bernholdt, D., Carver, J., Heroux, M., Hornung, R., Jones, P., Lucas, B., Siegel, A., Ndousse-Fetter, T.: Software Productivity for Extreme-Scale Science (2014), report on DOE Workshop, January 13-14, 2014, available via <http://www.ornl.gov/swproductivity2014/SoftwareProductivityWorkshopReport2014.pdf>
14. Keyes (PI), D.: Towards Optimal Petascale Simulations, SciDAC Institute. <http://www.scidac.gov/math/TOPS.html>, accessed: 2017-02-15
15. Klinvex, A.M.: xSDKTrilinos user manual. Tech. Rep. SAND2016-3396 O, Sandia (2016)
16. Lawson, C., Hanson, R., Kincaid, D., Krogh, F.: Algorithm 539: Basic linear algebra subprograms for Fortran usage. ACM Trans. Math. Software 5 (1979), DOI: 10.1145/355841.355848
17. Lawson, C., Hanson, R., Kincaid, D., Krogh, F.: Basic linear algebra subprograms for Fortran usage. ACM Trans. Math. Software 5 (1979), DOI: 10.1145/355841.355848
18. McInnes, L.C., Heroux, M., Li, X.S., Smith, B., Yang, with contributions from all xSDK developers, U.: What are Interoperable Software Libraries: Introducing the xSDK, version 0.1, April 25, 2016, available via <https://ideas-productivity.org/resources/howtos/>
19. Miller, M.C., Diachin, L., Balay, S., McInnes, L.C., Smith, B.: Package management practices essential for interoperability: Lessons learned and strategies developed for FASTMath (2013), DOI: 10.6084/m9.figshare.789055.v1
20. Miller, M.C., Reus, J.F., Matzke, R.P., Koziol, Q.A., Cheng, A.P.: Smart libraries: Best SQE practices for libraries with emphasis on scientific computing. In: Proceedings of the Nuclear Explosives Code Developer’s Conference (2004), available via <https://wci.llnl.gov/codes/smartlibs/UCRL-JRNL-208636.pdf>
21. Rüde, U., Willcox, K., McInnes, L.C., De Sterck, H., Biros, G., Bungartz, H., Coronas, J., Cramer, E., Crowley, J., Ghattas, O., Gunzburger, M., Hanke, M., Harrison, R., Heroux, M., Hesthaven, J., Jimack, P., Johnson, C., Jordan, K.E., Keyes, D.E., Krause, R., Kumar, V., Mayer, S., Meza, J., Mørken, K.M., Oden, J.T., Petzold, L., Raghavan, P., Shontz, S.M., Trefethen, A., Turner, P., Voevodin, V., Wohlmuth, B., Woodward, C.S.: Research and education in computational science and engineering (2016), available via <https://arxiv.org/abs/1610.02608>, submitted to *SIAM Review*
22. Smith, B., Bartlett, R., xSDK developers: xSDK community package policies (2016), version 0.3, December 2, 2016, DOI: 10.6084/m9.figshare.4495136
23. Smith, B.T., Boyle, J.M., Dongarra, J.J., Garbow, B.S., Ikebe, Y., Klema, V.C., Moler, C.B.: Matrix Eigensystem Routines – EISPACK Guide, Lecture Notes in Computer Science, vol. 6. Springer–Verlag, New York, second edn. (1976), DOI: 10.1007/3-540-07546-1

24. Alquimia Web page. <https://bitbucket.org/berklab/alquimia>, accessed: 2017-02-15
25. Hypre Web page. <http://www.llnl.gov/CASC/hypre>, accessed: 2017-02-15
26. NAG Web page. <https://www.nag.com>, accessed: 2017-02-15
27. PFLOTRAN Web page. <http://www.pflotran.org>, accessed: 2017-02-15
28. SuperLU Web page. <http://crd.lbl.gov/~xiaoye/SuperLU>, accessed: 2017-02-15
29. The Exascale Computing Project. <https://exascaleproject.org>, accessed: 2017-02-15
30. The HSL Mathematical Software Library. <http://www.hsl.rl.ac.uk>, accessed: 2017-02-15
31. Trilinos Web page. <https://trilinos.org>, accessed: 2017-02-15