

SANDIA REPORT

SAND2010-6890

Unlimited Release

Printed October 2010

Expanding The Trilinos Developer Community

Michael A. Heroux

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Expanding the Trilinos Developer Community

Michael A. Heroux
Scalable Algorithms Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185

Abstract

The Trilinos Project started approximately nine years ago as a small effort to enable research, development and ongoing support of small, related solver software efforts. The “Tri” in Trilinos was intended to indicate the eventual three packages we planned to develop. In 2007 the project expanded its scope to include any package that was an enabling technology for technical computing[1]. Presently the Trilinos repository contains over 55 packages covering a broad spectrum of reusable tools for constructing full-featured scalable scientific and engineering applications. Trilinos usage is now worldwide, and many applications have an explicit dependence on Trilinos for essential capabilities. Users come from other US laboratories, universities, industry and international research groups.

Awareness and use of Trilinos is growing rapidly outside of Sandia. Members of the external research community are becoming more familiar with Trilinos, its design and collaborative nature. As a result, the Trilinos project is receiving an increasing number of requests from external community members who want to contribute to Trilinos as developers. To-date we have worked with external developers in an *ad hoc* fashion. Going forward, we want to develop a set of policies, procedures, tools and infrastructure to simplify interactions with external developers. As we go forward with multi-laboratory efforts such as CASL[2] and X-Stack[3], and international projects such as IESP[4], we will need a more streamlined and explicit process for making external developers “first-class citizens” in the Trilinos development community.

This document is intended to frame the discussion for expanding the Trilinos community to all strategically important external members, while at the same time preserving Sandia’s primary leadership role in the project.

Acknowledgements: The author would like to thank the ASC, LDRD and ASCR programs that have funded this work, and James M. Willenbring who helped form some of the proposed changes.

Table of Contents

Executive Summary	6
1. Introduction	8
2. OpenMPI: A Case Study	9
3. A More Open Trilinos Developer Environment.....	10
3.1. Copyright and Licensing	10
3.2. Trilinos Contributor Agreements	11
3.3. Project Portal	12
3.4. Project Developer Site	12
4. Summary and Conclusions	13
5. References	14

This page is intentionally left blank.

Executive Summary

In order to collaborate with external developers most effectively, the Trilinos project proposes to make progress on four topics. These topics are discussed below in detail, but we state the recommendations here for quick reference.

1. **Copyright and Licensing:** *The Trilinos Project should continue efforts to make as much of its software base available under the BSD license as possible. Future new packages should be licensed under the BSD license. All future software contributions by outside individuals and organizations must be given to Trilinos under a BSD license with external contributor copyrights in appropriate source files.*
2. **Contributor Agreements:** *The Trilinos Project should have an individual and organization contributor agreement similar to OpenMPI. These agreements should be standard forms available from our website. All contributions, outside of Sandia-funded work that is already unambiguously owned by Sandia, should be made under one of these agreements.*
3. **Project Portal:** *The Trilinos Project portal (the public face of Trilinos) should be <http://www.trilinos.org>. This site will be the first place Trilinos users and developers will go for access to Trilinos documentation, discussions and downloads. We will not eliminate SSG, or TSG. In fact, the trilinos.org website will at first be a façade for these other sites, and allow us to gradually shift the location of data and services as we go forward, to best serve our interests. We anticipate eliminating TSG within one year, but will keep SSG indefinitely.*
4. **Project Developer Site:** *The Trilinos Project should continue using SSG as the primary project developer site, but we should explore other options for hosting the Trilinos developer tools and repositories in the future. At this time, we do not see a viable alternative to using SSG, but we hope that in the future we could provide more open access to external developers.*

This page is intentionally left blank.

1. Introduction

The Trilinos Project started approximately nine years ago¹ as a small effort to enable research, development and ongoing support of related solver software projects. The “Tri” in Trilinos was originally meant to indicate the eventual three packages we intended to develop. Presently the Trilinos repository contains over 55 packages, 48 of which are part of the Version 10.4.X release.

Obviously the size and scope of Trilinos has changed significantly from the original vision. We foresee that it will continue to change, and believe that the changes are positive. Trilinos has grown because it has provided an attractive environment for collaborative development of high-quality, state-of-the-art scientific software. Growth within Sandia is not presently a constraint: any library-focused software project can in principle participate in Trilinos. However, growth in external developer collaborations is restricted by policy, infrastructure and business needs. Although Sandia should retain control over the project for its business needs, it is in our best interest to reduce the policy and infrastructure barriers to external developers as much as reasonably possible. By doing so, we can increase external use and reliance on Trilinos and increase our opportunities for greater collaboration across laboratories, universities, industry and internationally. The options for improved external developer collaborations are (1) separation of externally developed capabilities from Trilinos thereby missing opportunities to improve Trilinos capabilities and complicating package dependencies or (2) continued *ad hoc* arrangements that are already barriers to collaboration with some external organizations and may lead to confusion about ownership of contributions in the future.

In this document we discuss the current state of external collaboration mechanisms in the Trilinos Project, and what options are available to improve policies, procedures and infrastructure.

2. OpenMPI: A Case Study

The OpenMPI Project [5] is a combined effort of several organizations, including Sandia, to develop an open source implementation of the Message Passing Interface[6]. Although the external interaction needs of Trilinos are not identical, OpenMPI serves as a notable case study because it shares a number of common or desired attributes with the Trilinos Project. Four specific topics of interest are:

1. **Copyright and Licensing:** OpenMPI software is licensed under the New BSD License [7]. All contributors are listed in the copyright notice [8] and source files contain the copyright of any institution that contributed to a given file.
2. **Contributor Agreements:** OpenMPI is structured to accept contributions from any organization or individual, as long as contributors are willing to sign a

¹ The first source repository commit was made on December 14, 2001.

standard agreement prior to making contributions [9]. Furthermore, there is an institutional version of the agreement that covers all members of an organization, even as new contributors from the organization are added.

3. **Project Portal:** The public face of the OpenMPI Project is the main website: <http://www.open-mpi.org>.
4. **Project Developer Site:** OpenMPI software, systems and tools are hosted at the University of Indiana. As a result, new developers can obtain access to systems fairly easily. Specifically, foreign nationals do not need special permission above and beyond U.S. citizens.

3. A More Open Trilinos Developer Environment

In this section we discuss the four topics of interest as listed in the previous section, and some of the options that we find attractive for Trilinos. Although we propose recommended actions, the purpose of this document is to frame the discussion for eventual decisions, not to jump to any decision prematurely.

3.1. Copyright and Licensing

Present Status: From the very beginning of the Trilinos project, we have distributed our software as open source under the GNU Lesser General Public License (LGPL)². In early 2010, we received approval to move to the New Berkeley Standard Distribution License (BSD)³. With Trilinos Version 10.6 in October 2010 we plan to begin releasing a large number of Trilinos packages under the BSD license. Small portions, where copyright is co-owned with another institution, will remain under the LGPL license until some joint agreement can be made with co-owners.

Recommendation: *The Trilinos Project should continue efforts to make as much of its software base available under the BSD license as possible. Future new packages should be licensed under the BSD license. All future software contributions by outside individuals and organizations must be given to Trilinos under a BSD license with external contributor copyrights in appropriate source files.*

Discussion: More than a year ago, we started discussing the move from LGPL to BSD. With increasing collaborations, especially with industrial partners, LGPL is considered too risky for many users who would like to base their own software on Trilinos capabilities. In particular, industry lawyers are very reluctant to allow development teams to incorporate LGPL software even as third party libraries. As a result we have

² LGPL contains a “copylefting” mechanism such that any changes made to the licensed software must be given back to the open source community. However, LGPL software can be wholly integrated into another project without forcing that project to also be open source. This is in contrast to the GNU General Public License (GPL), which requires any project that incorporates a GPL project to be open source.

³ BSD has no copylefting mechanism. Its only requirement is that copyright information be displayed by the incorporating project.

strong agreement that BSD is an appropriate model for Trilinos, since it eliminates the concerns of our industrial collaborators.

Risks and Requirements: The GNU LGPL license always guaranteed, within the law, that all Trilinos software and derivative works would remain open source. The BSD license has no such guarantees. Some outside organization could take all of Trilinos, modestly change it to its needs and never give the improved product back. Although in theory this risk seems serious, in practice we believe that anyone who is interested in using and contributing to Trilinos will want to collaborate with us and provide their contributions back to the project, or their changes will be special-purpose and not appropriate for the general Trilinos distribution. For example, a hardware vendor can make optimizations appropriate to a specific platform that will benefit Trilinos users on that platform, but these changes need not be part of the general Trilinos distribution.

By going to BSD, we must be vigilant to maintain the integrity of Trilinos as a single collection of software. There must be a single “gold standard” version managed by Sandia.

3.2. Trilinos Contributor Agreements

Present status: Presently we accept external contributions to Trilinos, but the vast majority of these come from colleagues who are officially sponsored by Sandia (for example summer interns) and, as part of their contract, agree to Sandia ownership of their contributions. Other contributions are made via formal CRADA agreements. Finally, some contributions are offered as modifications that are added by a Sandia developer.

Trilinos also contains software copyrighted by other organizations via an open source license of some kind, such as GNU LGPL or BSD. Copyright is indicated by the appropriate text in the source file.

***Recommendation:** The Trilinos Project should have an individual and organization contributor agreement similar to OpenMPI. These agreements should be standard forms available from our website. All contributions, outside of Sandia-funded work that is already unambiguously owned by Sandia, should be made under one of these agreements.*

Discussion: Although we can continue to work under the current model, we would like to expand the scope of external contributions without introducing ambiguity in the ownership or licensing of Trilinos software. By adopting a formal contribution agreement, we allow almost anyone to contribute to Trilinos while still protecting our interests in controlling Trilinos content.

Risks and Requirements: If the number of external contributors grows too quickly, we can lose control of the project. However, if we continue to host Trilinos development on Sandia systems (as discussed below), we substantially reduce this risk. Even so we must make sure that contributions are coordinated. Also, on a practical level, we will need a process to inspect contributions.

3.3. Project Portal

Present status: The Trilinos Project initiated the platform called software.sandia.gov (SSG) and our original website was <http://software.sandia.gov/trilinos>. Approximately five years ago, we split the project portal (the public face of Trilinos) off and put it on a separate system called trilinos.sandia.gov (TSG). All user-oriented webpages, documentation, and download access are managed on TSG. All developer-oriented tools and documentation remain on SSG. Separating the user and developer websites is a big improvement for robustness and security. Putting users on a corporate-supported system (TSG is managed by Sandia webservices) and keeping them off of the primary developer system that is managed by our local system administrators is very effective. We want to continue with the separation of these concerns, but in order to make Trilinos a broader community effort, we want to create a new web portal called trilinos.org.

Recommendation: *The Trilinos Project portal (the public face of Trilinos) should be <http://www.trilinos.org>. This site will be the first place Trilinos users and developers will go for access to Trilinos documentation, discussions and downloads. We will not eliminate SSG, or TSG. In fact, the trilinos.org website will at first be a façade for these other sites, and allow us to gradually shift the location of data and services as we go forward, to best serve our interests. We anticipate eliminating TSG within one year, but will keep SSG indefinitely.*

Discussion: Creating the trilinos.org website is a strong symbolic move that tells users and external collaborators that we are open to external collaborators and can give their institution recognition within the Trilinos Project without explicitly pulling it underneath the Sandia banner, even though Sandia will continue to be the *de facto* hub of development.

Risks and Requirements: Removing Sandia's name from the main Trilinos portal site (going from trilinos.sandia.gov to trilinos.org) can be seen as a risk and a loss of an important connection. However, there is a growing drive to create multi-institutional library projects. Already many Trilinos developers are at other labs and universities. By moving Trilinos in a direction that fosters multi-institutional development, Trilinos can become the multi-institutional project that people are looking for. If we do not make this move, there is a chance that some other multi-institutional project could emerge, which could stunt the growth of Trilinos.

3.4. Project Developer Site

Present Status: All Trilinos developer tools and source repositories are hosted on software.sandia.gov (SSG). Although SSG is on the Sandia Open Network (SON) and provides an externally visible website, external collaborators are subject to some network access restrictions. In particular, foreign nationals require significant paperwork to receive an account.

Recommendation: *The Trilinos Project should continue using SSG as the primary project developer site, but we should explore other options for hosting the Trilinos*

developer tools and repositories in the future. At this time, we do not see a viable alternative to using SSG, but we hope that in the future we could provide more open access to external developers.

Discussion: Migrating Trilinos development tools and software repositories to an external system would make access by external developers easier, as it does for OpenMPI developers, whose systems are based at Indiana University. However, by keeping Trilinos development on Sandia systems, we maintain a strong degree of control in the midst of the above-proposed changes.

Risks and Requirements: We see little risk in maintaining the status quo, and there are no additional requirements. We should revisit this decision over the next year or so, once other issues are resolved.

4. Summary and Conclusions

The Trilinos Project has been growing and expanding for nearly a decade. Given the strong community adoption of Trilinos and the enthusiasm of external collaborators, we believe now is a good time to make the necessary structural changes to Trilinos that position us to make Trilinos a true community project. Moving in this direction comes with some risks, but we believe the opportunities for growth far outweigh these risks. In fact, the greatest risk is that some other project is developed in a way similar to what we propose above, allowing it to become the primary community libraries project, and Trilinos growth is stunted as a result.

Comments and concerns are welcome. Please contact the author.

5. References

1. Heroux, M.A., *The Changing Scope of the Trilinos Project*. 2007, Sandia National Laboratories: Albuquerque, NM, SAND2007-7775.
2. Kothe, D. *Consortium for Advanced Simulation of Light Water Reactors*. 2010; Available from: <http://www.ornl.gov/sci/casl/>.
3. *X-Stack Call for Proposals*. 2010; Available from: <http://www.grants.gov/search/search.do?mode=VIEW&oppId=51444>.
4. Dongarra, J., et al., *The International Exascale Software Project: a Call To Cooperative Action By the Global High-Performance Community*. *Int. J. High Perform. Comput. Appl.*, 2009. **23**(4): p. 309-322.
5. *Open MPI: Open Source High Performance Computing*. 2010; Available from: <http://www.open-mpi.org/>.
6. *The Message Passing Interface (MPI) Forum*. 2010; Available from: <http://www.mpi-forum.org>.
7. *Open Source Initiative OSI - The BSD License: Licensing | Open Source Initiative*. 2010; Available from: <http://www.opensource.org/licenses/bsd-license.php>.
8. *Open MPI License*. 2010; Available from: <http://www.open-mpi.org/community/license.php>.
9. *Contributing to Open MPI*. 2010; Available from: <http://www.open-mpi.org/community/contribute/>.

Distribution

1	MS 0321	R. W. Leland, 01400
1	MS 1318	B. A. Hendrickson, 01410
1	MS 1320	S. S. Collis, 01416
1	MS 1318	R. J. Hoekstra, 01414
1	MS 0380	D. E. Womble, 01540
1	MS 1318	J. R. Stewart, 01411
1	MS 1320	M. A. Heroux, 01416
1	MS 1320	J. M. Willenbring, 01414
1	MS 0899	Technical Library, 9536 (electronic copy)

This page is intentionally left blank.



Sandia National Laboratories