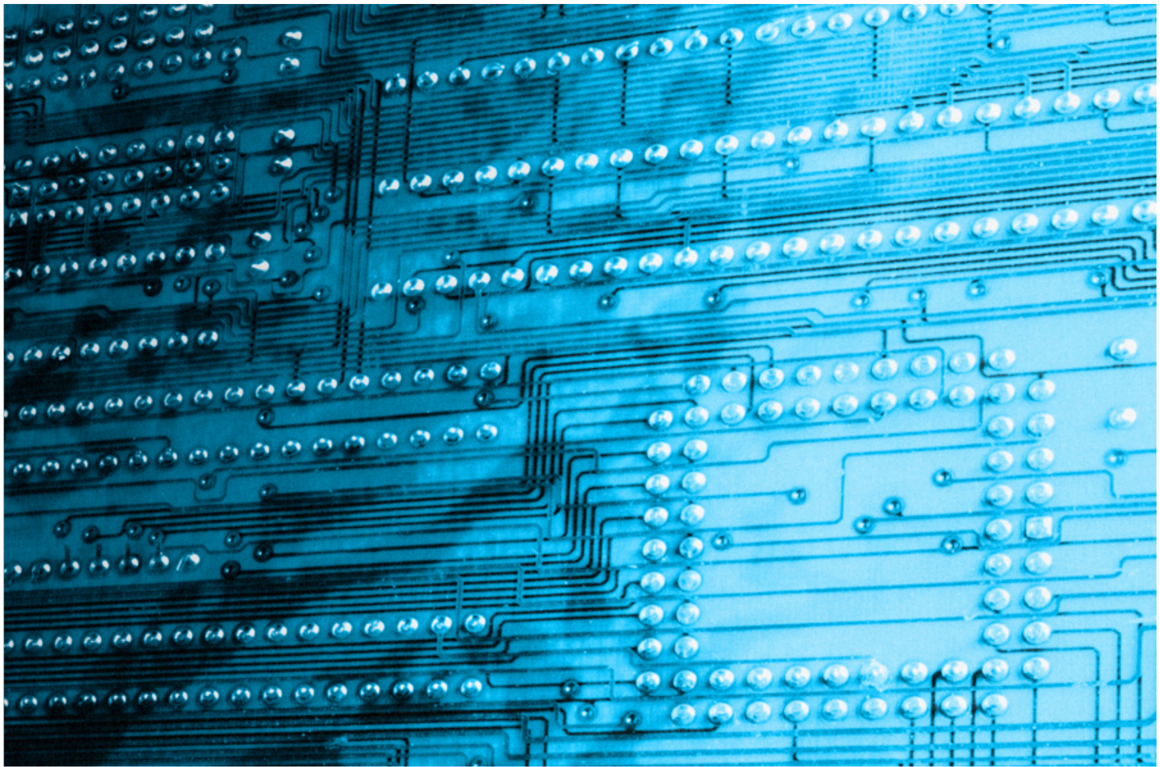

Report to the National Science Foundation
Office of Cyberinfrastructure (OCI)

9 November 2011

Report of NSF Workshop Series on Scientific Software Security Innovation Institute

http://security.ncsa.illinois.edu/s3i2/

This material is based upon work supported by the *National Science Foundation* under grant number *1043843*. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect

Table of Contents

Organizers and Editors

The workshop organizers were Randy Butler (co-Chair, U. of Illinois/NCSA), and Von Welch (co-Chair, Indiana U.). In addition Jim Basney played a key role in collecting and organizing workshop notes, and in helping to craft the workshop report, and Scott Koranda provided a valuable presentation on the relevant experiences of the LIGO project.

Organizers of the first workshop were William Barnett (Indiana University), Jim Basney (U. of Illinois/NCSA), Randy Butler (Chair, U. of Illinois/NCSA), and Doug Pearson (REN-ISAC/Indiana University)

Acknowledgements

This material is based upon work supported by the National Science Foundation under grant number 1043843. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science.

Contributors

The following individuals attended one or both of workshops in the NSF Workshop Series on Scientific Software Security Innovation Institute and contributed to the writing of this report.

Name	Project Affiliation	Workshop(s) Attended
Mine Altunay	Open Science Grid	Second
Rachana Ananthakrishnan	Globus Project	First
Matthew Arrott	Ocean Observatory Initiatve	Second
Steve Barnet	IceCube	Second
William Barnett	Organizer	First
James Basney	CILogon, MyProxy, Organizer	Both
Steve Beaty	NCAR	Second
Randal Butler	Organizer Chair	Both
Sayeed Choudhury	Data Conservancy	First
Heidi Dempsey	GENI	First
Rion Dooley	iPlant	Second
Michael Freemon	NVO	Second
Geoffrey Fox	FutureGrid	First
Kate Keahey	Nimbus	Second
Ken Klingenstein	Internet2, InCommon, Shibboleth	First
Scott Koranda	LIGO	Both
Miron Livney	Open Science Grids, Condor	First
James Marsteller	TeraGrid/XSEDE	Second
John McGee	RENCI	Second
Pascal Meunier	nanoHUB, HUBzero	First
Reagan Moore	iRods, Data Federation Network	Second
Stephen Newhouse	EGI	First
Doug Pearson	REN-ISAC, Organizer	Both
Don Petravick	DES	Second
Beth Plale	SEAD DataNet	Second
Ray Plante	LSST	Second
Scott Poole	IChass	First
Ruth Pordes	Open Science Grid	First
Mark Serville	LTER	Both
Abe Singer	LIGO	First
Adam Slagell	GENI, Blue Waters	First
Kevin Thompson	NSF	Second
Dave Vieglas	DataOne	Both
Von Welch	Organizer	Both
Patrick West	Semantic eScience Framework (SESF)	First
Alex Yahja	GroupScope	First

Executive Summary

Over the period of 2010-2011, a series of two workshops were held in response to NSF Dear Colleague Letter NSF 10-050 calling for exploratory workshops to consider requirements for Scientific Software Innovation Institutes (S2I2s). The specific topic of the workshop series was the potential benefits of a security-focused software institute that would serve the entire NSF research and development community. The first workshop was held on August 6th, 2010 in Arlington, VA and represented an initial exploration of the topic. The second workshop was held on October 26th, 2011 in Chicago, IL and its goals were to 1) Extend our understanding of relevant needs of MREFC and large NSF Projects, 2) refine outcome from first workshop with broader community input, and 3) vet concepts for a trusted cyberinfrastructure institute. Towards those goals, the participants of the 2011 workshop included greater representation from MREFC and large NSF projects, and, for the most part, did not overlap with the participants from the 2010 workshop.

A highlight of the second workshop was, at the invitation of the organizers, a presentation by Scott Koranda of the LIGO project on the history of LIGO's identity management activities and how those could have benefited from a security institute. A key analysis he presented is that, by his estimation, LIGO could have saved 2 senior FTE-years of effort by following suitable expert guidance had it existed.

The overarching finding from the workshops is that security is a critical crosscutting issue for the NSF software infrastructure and recommended a security-focused activity to address this issue broadly, for example a security software institute (S2I2) under the SI^2 program. Additionally, the 2010 workshop participants agreed to 15 key additional findings, which the 2011 workshop confirmed, with some refinement as discussed in this report. The major refinements from the 2011 workshop were:

- The NSF CI ecosystem increasing includes a number of "cloud" and other services provided by external parties. Any effort to increase the trustworthiness of that CI ecosystem must take those services into account in addition to software produced under NSF funding.
- The S2I2 must carefully balance providing a service unbiased towards any particular solution with keeping its staff suitably up-to-date through their involvement in projects, including the research and development of solutions.
- Adoption of orphaned software is a measure of last resort and the S2I2 should ideally avoid it if at all possible through planning and coordination.
- Assessment of software and services is valuable, but the S2I2 must be careful not to turn it into a purely bureaucratic function.
- A key goal of the S2I2 in providing leadership is the continuous building and distribution (through education, training and workforce development) of a body of knowledge on the topic of trustworthy CI. This includes successes and lessons learned from projects so that other projects can benefit from those.
- A key goal of the S2I2 in providing leadership is aggregating community needs and speaking on behalf of the community to external entities (e.g. InCommon, REN-ISAC).

These refinements resulted in the following set of refined key findings:

1. A security-focused S2I2 should provide NSF and the NSF research community with security leadership and guidance.
2. A security-focused S2I2 should provide documentation, training, recommendations, and consulting to NSF cyberinfrastructure projects both on software security and security software.
3. A security-focused S2I2 should provide short-term support for orphaned security software deemed critical to NSF cyberinfrastructure projects.
4. A security-focused S2I2 should perform independent software and services security assessments.
5. A security-focused S2I2 should support security design reviews of MREFC projects or smaller CI development and integration efforts.
6. The institute should independently highlight/rank security software that NSF CI relies upon.
7. The institute should provide a security auditing service that includes vulnerability analysis and overall security assessment that validates security functions within a CI.
8. The institute should not "own" software to avoid conflicts of interest between fostering adoption of that software and guiding communities to the best solution for their needs. Further the institute should take care not to allow their participation in software integration or development bias their future recommendations of that software.
9. The institute should not provide operational security services or replicate existing services.
10. The institute should be governed in an open fashion that provides venues for stakeholders to discuss priorities and influence the institute's activities as well as assures them of the institute not being influenced by any member's baises.
11. The institute should be a synthesis point for expertise but not necessarily own all the expertise in-house.
12. The institute should coordinate its efforts and seek support across federal agencies including DHS, DOE, DARPA, and NIH.
13. The institute should have well defined relationships with the CMU Software Engineering Institute, InCommon, Internet2, REN-ISAC, and the XD TAIS.
14. Funding in addition to funds supplied by NSF for a security-focused software institute should be aggressively pursued.
15. The institute must document how it would gauge its own success.

Introduction

The NSF Software Infrastructure for Sustained Innovation (SI^2) program has proposed to establish Scientific Software Innovation Institutes (S2I2s) to transform grassroots computational science and engineering software into robust and sustained software infrastructure for science and engineering. In addition to focusing on community-based software institutes, there are additional opportunities to provide cross-disciplinary cyberinfrastructure (CI), including software, practices, policies, and services, which will support all of the community-based institutes and avoid redundant efforts across institutes.

A series of two workshops was held to explore the potential for a cross-disciplinary security-focused scientific software innovation institute, to address the protection, integrity, and reliability of research processes and information. The workshops brought together representatives from NSF-funded CI projects, computational researchers, CI developers, security developers, and resource providers. The discussions covered: research security needs; existing tools, systems, processes, and organizations that secure research activities and data; outstanding issues to be addressed in research assurance; and organization and operational models for a future security institute targeting the identified security needs.

These workshops aided in the understanding of the role of cybersecurity software, practices, policies, and services in NSF research, documented the security requirements and priorities of a range of representative NSF projects and researchers, identified outstanding needs, and produced recommendations for next steps in assuring the integrity of scientific research and research data into the future.

A report from the first workshop was authored with the results of the first workshop and is available on the workshop website:

http://security.ncsa.illinois.edu/s3i2/

The second workshop was held to broaden the community input into the findings of the first workshop, particularly from MREFC and other large NSF projects. Towards those goals, the participants of the 2011 workshop included greater representation from MREFC and large NSF projects, and, for the most part, did not overlap with the participants from the 2010 workshop.

Workshop Goals

The goals of the 2010 workshop were to:

1. Document software security efforts in place in order to develop a competitive landscape of options,
2. Document security needs from the perspective of a wide variety of domain scientists, representing virtual organizations, national observatories, and small research projects,
3. Recommend whether a security-focused software institute should move forward and, if so, to identify observatory or other project partners, organizational structures, and services that would comprise this institute, and

4. Understand the parameters, including but not limited to financial, policy, and human, which influence sustainable security for research cyberinfrastructure.

The goals of the 2011 workshop were to:

1. Extend our understanding of needs of MREFC and large NSF Projects,
2. Refine outcome from first workshop with broader community input,
3. Vet concepts for a trusted cyberinfrastructure institute.

Summary of Second Workshop

The second workshop was held on October 26th at the O'Hare Hilton in Chicago IL. Nineteen participants (plus the two organizers) attended, representing 16 NSF projects (TeraGrid/XSEDE, LIGO, DES, NVO, LSST, OSG, LTER, DataOne, SEAD, iRods, Data Federation Network, CILogon, iPlant, IceCube, Nimbus, OOI) plus 2 organizations (NCAR, RENCI) and the broader security community (REN-ISAC).

The workshop started with a presentation to orient the participants on the results of the first workshop and describe the thinking of the organizers since the first workshop[1]. By the invitation of the workshop organizers, Scott Koranda from the LIGO project presented details on LIGO's efforts to transition towards a single authoritative roster of membership, which supports a single LIGO identity for each member[2]. Mr Koranda detailed the many steps, setbacks and successes as they took on this unfunded challenge as best they could. In the end, Mr Koranda estimated the potential savings for the LIGO project if they could have consulted with a security focused software institute:

"It is difficult to estimate, but I expect if a NSF S3I2 had existed and offered non-biased consulting services around IdM and cybersecurity LIGO would have saved two years of senior FTE effort (effort many smaller VOs do not have)."

The majority of the workshop was then open discussion with moderation by the organizers. The participants selected a number of topics for discussion during this time, which were:

1. What are the ways to implement the institute?
    a. Who are the experts and how to engage them?
2. What additional services should the institute consider?
    a. Orphaned security software – a revisit from the first workshop.
    b. Handling security services (cloud services Facebook/Google/EC2/Rackspace).
3. How will the institute establish and maintain relevance?
    a. If the institute doesn't develop software, will that negatively impact relevance?
    b. Risks of operating in a purely advisory role.
    c. How to build & maintain relevance/leadership in the longer term (15+ years) (and short term)?
4. Governance of institute.
    a. What are the institute's project selection criteria?

[1] http://security.ncsa.illinois.edu/s3i2/S2I3-Oct26-Workshop-slides.pptx

[2] http://security.ncsa.illinois.edu/s3i2/S3I2WorkshopChicagoOctober2011.pdf

5. Relationship with other organizations/institutes: InCommon, REN-ISAC, CERT, EDUCAUSE, Internet2 Net+ work.
6. Sustainability. Funding sources.
    a. How to avoid bias? Sources of funding.
7. How to create incentives for projects to be concerned about security? NSF requirement for security assessment plan?
8. Project Life Cycle support
    a. Engage with projects prior to award?
    b. What support should the institute offer at each project phase?
9. How to measure value of institute?

The results of these discussions are captured in the subsequent section.

Workshop Output: Key Attributes of a Security Software Institute

The key attributes of a security software institute were one of the main discussion topics for both workshops. The first workshop produced a key finding and a set of 16 additional findings. The second workshop re-affirmed these finds, adding refinement. In this section we present those findings with the refinements from the second workshop.

Key Finding

The key finding from both workshops was that security is a critical crosscutting issue for the NSF software infrastructure that must be addressed in NSF's Software Infrastructure for Sustained Innovation (SI^2) program and there is is a strong need for a "long-term community-wide hub of software excellence" focused on software security. Security is a fundamental building block that facilitates cooperation and collaboration, and without interoperable approaches, effective security is difficult to achieve. Further, security is a shared requirement and all projects could benefit by leveraging existing proven approaches and implementations. Finally, because security expertise within the NSF CI community is limited, a security-focused S2I2 would be of significant benefit to the NSF CI community.

Additional Findings

Workshop participants agreed to a number of key findings which follow, grouped by topic.

What should the institute do?

Finding #1: A security-focused S2I2 should provide NSF and the NSF research community with security leadership and guidance. There are potentially multiple customers for this service, including the NSF and the NSF CI community. With respect to NSF as a customer, the security institute could provide a good sounding board for the development of new ideas including providing non-biased opinions to NSF on the potential usefulness or need for a suggested new development activity or the list of existing related activities so that there is less duplication. Likely this would be input into the development of new solicitations and not part of the proposal review process.

The institute could additionally provide security guidance and leadership to the NSF CI community, i.e. those developing, deploying or operating CI in support of NSF funded projects. Specifically, the institute should identify common requirements and document technologies and common approaches to assist these projects in developing their approach to security. Additional information dissemination about a range of topics could be valuable, including the documentation of threats and attacks and the security landscape (new technologies and services). The participants drew analogies with the National Academy of Science model and recommended that a security-focused S2I2 could develop reports that the community takes seriously.

The second workshop confirmed this finding from the 2010 workshop. A key refinement made was that such an institute should also provide an aggregated voice for the needs of the NSF community in interactions with external communities (e.g., InCommon, REN-ISAC).

Finding #2: A security-focused S2I2 should provide documentation, training, recommendations, and consulting to NSF cyberinfrastructure projects both on software security and security software. Security training and documentation should be focused around helping CI developers, those that deploy CI to integrate and support security technologies. It might include documentation on how to perform a risk and threat analysis, how to deploy and utilize security software, how to develop or improve usable software that is secure.

In addition, the institute could provide guidance to software developers in providing well thought out diagnostics and provide documentation to assist those using the software in diagnosing problems that may arise. Further the institute might assist decision-makers and project managers in their design and deployment efforts.

The second workshop identified a key goal for the institute, which is the constant building and disseminating of cybersecurity knowledge. This would include documenting the approaches to security that each project takes on, and the successes and lessons learned, which would serve to assist other projects in making decisions.

Training materials, training events, workforce development and workshops designed around knowledge transfer were identified in the 2011 workshop as and key deliverables for the institute. Participants cited that the value of the institute is in transferring knowledge from security experts, lessons learn, and the successes and failures of each project, to other projects and project development staff. However, direct end-user training was seen to be out-of-scope for the institute; where end-user is defined as a customer of a CI project. Instead the institute should provide support to other organizations that do end-user training. Closely related to documentation and training is the ability to provide assistance to developers and those deploying security software for basic support assistance. This may be implemented as a clearinghouse to direct people to the most appropriate source of information or it may provide them with documented use case examples.

Finding #3: A security-focused S2I2 should provide short-term support for orphaned security software deemed critical to NSF cyberinfrastructure projects. It should facilitate the location of a new long-term base of support for the software or assist projects in transitioning to better-supported alternatives. There is a potential role for the institute with respect to short-term support for, or advocacy for, critical security software that is no longer supported. There may be a need for the institute to pick up support for a temporary period while working with the community to identify a longer-term support mechanism. The institute would play a key role in identifying what security software was critical to the NSF community and advocate on its behalf to NSF and other organizations that may be able to assist in its support. Further, the institute may play a role in advocating the integration of key security software into both commercial and larger open source software.

In further discussions in the 2011 workshop there was concern over the institute becoming saddled with orphaned software. In fact some participants worried that the presence of the institute may actually encourage some developers to abandon their software, hoping the institute would then pick it up. Additionally there was concern that the institute would no longer be seen as a non-biased entity once it is supporting software, and finally it was not clear that the institute would have a sufficient software development staff to support any software. Participants of the second workshop strongly encouraged that the institute take steps to mitigate the negative impacts of abandoned critical security software, short of supporting it directly. Suggestions were to first look for alternative solutions, next work with the project teams that were dependent on the software to see if any project team could pick up support within their project, if not then work collaboratively with all the projects that depend on the software for a solution and if all those efforts fail then seek out a 3rd party for support either directly or by working with NSF.

Finding #4: A security-focused S2I2 should perform independent software and services security assessment. A security institute should support the independent assessment of software from a security point of view, as is done in the Middleware Security and Testing (MIST) project at UW-Madison. The independence of the assessment is a critical attribute to lessen the potential for negative biases and assure community trust. There are a number of commercial tools for software assessment including Nessus, Coverity, and Fortify; however, there is a lack of expertise in how to interpret the results and address the issues identified. Additionally, there are other aspects of assessment that are currently human-oriented, such as architectural analysis and analysis of services for which the software is not available (e.g., cloud services) – see Finding #7. The institute should provide this service to the development teams in a responsible manner so as to encourage continued collaboration with the goal of providing more secure software.

The second workshop also identified that the NSF CI community is increasingly using cloud services (e.g., Amazon EC2) and assessment of how the use of these services affects the security risks of NSF research is increasingly important. The second workshop also identified the concern that the institute needs to take care not to turn assessment into a purely bureaucratic function (a checkbox that projects must complete at some point), but instead a meaningful part of transitioning any CI from a research phase to a production phase.

Finding #5: A security-focused S2I2 should support security design reviews of MREFC projects or smaller CI development and integration efforts. The institute could provide a list of architecture-related reviews including risk and threat analysis, policy examples, architecture, and implementation design reviews.

The second workshop expanded upon this finding suggesting that the institute could provide documentation on what a security design should include, how to develop one, and examples. Further the institute could assist projects by performing preliminary security design reviews to better prepare them for the official review.

Finding #6: The institute should independently highlight/rank security software that NSF CI relies upon. Independent assessment of value could help software owners obtain funding and help infrastructure providers make informed decisions about which software to deploy. The institute should use clearly defined metrics for the assessment of security software for the NSF community. Such metrics might include a listing of related software that has similar capabilities, dependencies on other software, usage statistics and a listing of what projects are using the software, what other software it has been successfully integrated with, what software it typically works with, and how well it meets security assessment guidelines. The institute could also provide an unbiased "weather forecast" (report on longevity/support) on security software, that goes beyond a simple ranking based on functionality and supportability into assessments based on longer-term support issues that might require a more detailed understanding of the funding landscape for software development projects.

The second workshop confirmed this finding.

Finding #7: The institute should provide a security auditing service that includes vulnerability analysis and overall security assessment that validates security functions within a CI. This is related in some sense to finding #4, however where that focuses on software assessment this recommendation is a focused assessment of the larger CI. This could involve documenting security guidelines with exemplars that projects operating CI could follow. The extent of how far to carry out a security assessment should be taken is unclear. Operational security requires an understanding of the entire CI environment and thus it is critical that any audit include an assessment of the broader CI, however this could represent a huge time commitment.

The second workshop discussed the need for supporting project teams that are transitioning from development to operations. Preparing them to document security policies and how to enforce those policies through operational practices and procedures, and the following up with reviews as described above to ensure they are conforming to their own policies and best practice.

What should the institute not do?

The workshop provided a good venue for the discussion of not only what a security-focused S2I2 might do but also on the kinds of services and activities that it should not do. The "do not do" topics covered at the workshop would each draw such an institute away from an unbiased center of excellence and would limit such an institute's ability to influence and lead the NSF CI community in security solutions for NSF CI.

The institute should take care not to allow their participation in software integration or development bias their future recommendations of that software.

Finding #8: The institute should not "own" software to avoid conflicts of interest between fostering adoption of that software and guiding communities to the best solution for their needs. Through the course of engagements the institute's development staff may participate and even lead software development efforts. Before beginning such a project the institute will document how the proposed software will be supported and further any and all software developed will be open-sourced with an identified support strategy. Acceptable support strategies include support by the community/project for which the software was developed or an open-source software development community. The NSF SI^2 solicitation identifies the need to sustain software infrastructure and the need to create anchors and leadership. Combining these can cause conflicts. One institute can't effectively do both, because the former causes biases for the latter. This discussion goes back to the earlier discussions about the key attributes of the security-focused software institute and workshop participants' feeling that the primary role of the security-focused software institute was to provide the NSF research community with unbiased leadership. There was strong feeling that the institute could not maintain an unbiased approach if it had any direct ties to software in the form of support. The institute should not directly develop specific software products unless there is a clearly identified support team outside of the institute. The institute must be non-biased in order to be able to establish an advisory role to projects and to NSF. Owning software would bias the institute towards the solution set of the participants.

The second workshop confirmed this finding from the 2010 workshop. It did discuss however a distinction between the institute being involved in research and development versus its members being involved in research and development. The participants recognized that expertise is acquired and maintained through being involved in research and development of new solutions, and hence institute team members involvement in research and development activities to maintain their expertise was seen as a positive. Governance of the institute needs to ensure that the bias of any member towards the solutions they have developed does not influence the service provided by the institute.

Additionally it was felt that the institute could serve to assist the transition of security research software, where security focused software means software developed by security researchers that has not yet been hardened for production environments. Suggestions included staying abreast of security research with an eye towards research software that shows potential to solve an existing or expected gap in NSF CI and providing either a matchmaking between a CI project in need with the security research group, or in advocating to NSF for the need to harden the code.

Finally the workshop participants felt, it would be difficult for the institute to play a role here, as it would likely translate into both development and support of software. However if an acceptable support solution is identified then the institute could participate in these activities.

The second workshop confirmed this finding from the 2010 workshop, with the same distinction between the roles of the institute versus its membership.

Recommendation #9: The institute should not provide operational security services or replicate existing services. If the security institute supported operational security services, such as an identity management service or monitoring, those services would likely influence or bias the institute to recommend that service, therefore causing the institute to lose credibility. The institute should maintain independence from security operations. In some cases there are already existing services, and the participants clearly felt that a security-focused S2I2 should not replicate existing security services. Examples of these kinds of services include but are not limited to:

- Coordination of software vulnerability handling: The institute should not duplicate the work done by the CERTs to handle software vulnerabilities, but it could provide guidance to NSF CI projects on vulnerability handling policies and procedures and assist the projects in connecting with the appropriate coordinating organizations (i.e., CERTs).

- Security incident information sharing: REN-ISAC already provides a valuable service for the sharing of security incident information, which should be leveraged and not be replicated.

- Security monitoring: The institute should not provide monitoring services for projects that have 24/7 operational services but are lacking staff to perform security monitoring and analysis of the events. It is important for the institute to remain free of operational services including security monitoring, as it would be a distraction from the software focus of the institute.

The second workshop confirmed this finding from the 2010 workshop, in that the institute should not provide operational security services, although there was emphasis on the importance of providing guidance, documentation and training on operational security to assist projects that are transitioning into production as well as those projects already in production. The institute could assist in development of policies, and procedures through best practice guides, documented examples, and focused training. There was additional discussion of leveraging the experience of distributed security teams such as the one found within the eXtreme Science and Engineering Discovery Environment (XSEDE) project.

Governance Models

For a security focused software institute to be successful, it must be an unbiased entity that provides the NSF research community with leadership and recommendations. Therefore, the institute will require a strong, well-documented governance model. This section focuses on the fourth workshop objective: *to understand the parameters, including but not limited to financial, policy, and human, which influence sustainable security for research cyberinfrastructure*.

What is governance?
Governance is the process by which stakeholders oversee the management of an operation or institute. Governance includes policies and objectives (community driven, strategic), together with staff people (managerial, tactical). The policies and objectives are meant to serve the institute's stakeholders so it is important to first understand whom those stakeholders are. The stakeholders include four primary group: 1) software developers that create software for NSF researchers, 2) NSF-funded CI projects that are deploying and supporting infrastructure, 3) the other SI^2 awardees, and 4) NSF itself. The workshop

participants felt that such an institute would not directly support NSF CI end-users. It is also unlikely that the institute would serve university campuses.

Recommendation #10: The institute should be governed in an open fashion that provides venues for stakeholders to discuss priorities and influence the institute's activities as well as assures them of the institute not being influenced by any member's biases. Stakeholders translate the goals of the institute into policies that guide management. Institute managers then make decisions, based on those policies and goals, about how to allocate resources and undertake tasks. In the case of the security-focused software institute, management would allocate time to training, software assessment, and other work of the institute. Stakeholders subsequently review the performance of institute management in terms of outcomes, efficiency, and effectiveness to ensure alignment with policies and goals. It was noted that there are a range of governance models from the highly participatory to central authority models. Highly participatory models provide for an excellent way to gather consensus, however they often suffer from a loss of focus and follow-through as priorities change. The central authority model, meanwhile, lacks the open participation in priority setting but can do an excellent job staying of task. Examples of governance approaches that fall somewhere in the middle include the one-page proposals process for LIGO and the CILogon workshops that set project priorities and goals based on the needs of their constituents. Both LIGO and CILogon governance models involve active community engagement and participation, however in the end the projects make the final priority decisions.

Example governance models such as the executive and technical advisory boards for REN-ISAC, the NEES external governance board, and the Open Science Grid's consortium-based advisory boards were discussed as viable options. These organizations all have found these "internal" advisory boards to be far more effective than external advisory boards, where board members don't have a stake.

The 2011 workshop discussed how an open governance model was also important to ensure stakeholders would feel secure in knowing the services they were receiving from the institute were not biased towards a particular solution. Acknowledging that all people have biases and that being involved with researching and developing solutions are critical to gaining and maintaining expertise, the institute's governance is important to ensure that biases of the institute's members don't influence the help it provides.

Finding #11: The institute should be a synthesis point for expertise but not necessarily own all the expertise in-house. It was felt that the institutes would not have the capability to house the experts for all areas and that they should instead draw on the combined expertise within the NSF research community. An open question was: if the institute pulled in knowledge from external contributors, what would be the incentive for external participation?

There were extended discussions on this topic in the second workshop and which focused on how to attract expertise to the institute, or rather on what might be incentives for security experts to make themselves available to the institute. While no solution was arrived at it was noted that this could be a challenge for the institute to overcome.

Key Relationships

Workshop participants were also interested in identifying if there was potential for different types of software institutes, and if so how would a security-focused software institute fit into the larger SI^2 ecosystem? The open question was whether there would be a single general purpose S2I2 with affiliated specialized S2I2s such as the security-focused software institute discussed at this workshop. Many of the workshop participants felt that there was a strong possibility of a more general S2I2 that addresses crosscutting CI software; if this were the case, clearly a security-focused institute would coordinate through the more general S2I2.

Finding #12: The institute should coordinate its efforts and seek support across federal agencies including DHS, DOE, DARPA, and NIH. Participants strongly favored linking the proposed security-focused software institute with other federal agencies such as DHS, DOE, and DARPA. Participants were interested in how this security-focused software institute might coordinate cross agency and even gain support through funding or other avenues from these other agencies.

This was lightly discussed in the second workshop and participants saw this cross agency coordination as a difficult challenge for the institute and that the goal here ought to be to look for and assist in any cross agency coordination efforts. The concern was also raised that as the institute found support from other sources, it would need to take care not to take on conflicting priorities, which could make serving the NSF community more difficult.

Finding #13: The institute should have well defined relationships with other S2I2s, the CMU Software Engineering Institute, InCommon, Internet2, REN-ISAC, and the XD TAIS. The list here is not exhaustive nor is it intended to be, rather it is meant to highlight that a security-focused S2I2 needs to establish and maintain an array of relationships with other projects, agencies, standards bodies, and organizations with related and complementary expertise. Further, the scope of these relationships is not limited to just national efforts, but the institute should recognize and establish relationships with both national and international bodies.

The second workshop confirmed this finding.

Financial Stability

Finding #14: Funding in addition to funds supplied by NSF for a security-focused software institute should be pursued. Financial stability of such an institute was discussed and the participants explored a various options. It was assumed that a security-focused software institute would begin under grant funding but that the funding model might evolve over time. It was not clear to the participants that such an institute could be sustained without NSF funding, however the group felt that other supplemental funding could potentially be developed, including institutional underwriters such as was developed by the Saki project, volunteer membership model such as with Linux, fee for service, corporate partnership/sponsorship, and support from other government agencies. It was felt that long-term funding from NSF would be critical in sustaining longer-term coordinated activities.

Participants from the second workshop were concerned about the impact on the non-biased nature of the institute if it were to actively seek additional outside funding as that may force the institute to lean one direction or the other. It was suggested that the institute should not author proposals for funding.

Gauging Success

Finding #15: The institute must document how it would gauge its own success. Initial justification for funding a security-focused S2I2 would have to come from a speculative assessment of community needs, but longer-term funding should be justified by real metrics that assess the impact that such an institute is having, such as how many projects utilize the institute. Other metrics may include measurement of the movement of NSF communities towards the institute's suggested approach, the level of participation in institute-organized workshops or training, and how many experts are engaged in supporting the institute's goals.

The second workshop discussed the need for metrics as one way to help establish and maintain relevance. The participants however did not identify specific required metrics for the institute.

Conclusions

This workshop series explored the topic of a security-focused scientific software innovation institute. There was enthusiastic agreement that such a security-focused software institute would benefit both individual NSF CI activities and projects, as well as provide a broad leadership and guidance across all NSF CI activities. A number of ideas and possibilities were discussed at a high-level during the two, daylong workshops that resulted in 15 distinct findings that NSF should consider as they develop a solicitation around the S2I2 theme. Those recommendations fell generally into five categories: 1) key activities of the institute, 2) key activities to avoid, 3) governance recommendations, 4) key relationships the institute should support, and 5) financial support recommendations.

Key recommended activities included providing leadership and guidance on security topics to NSF and the NSF research community, providing documentation, training, and consulting advice, shepherding of critical orphaned security software, software security assessments, security design reviews, ranking of security software, and security auditing. Workshop participants felt strongly that such an institute should not develop new software, integrate software, provide operational security services, or replicate existing services. However the institute could play a vital role in identifying security research software that has potential to fill NSF CI gaps, and further the institute should assist in matchmaking between security researchers, and CI projects as well as advocating to NSF to provide assistance in production hardening of this software. Such an institute should be governed in an open fashion that supports stakeholder input, and the institute should draw expertise broadly through collaborative relationships. The institute should develop metrics that could be used to assist in measuring impact and to guide the institute in setting priorities. The institute should develop and maintain relationships with other federal agencies and a number of existing community efforts and NSF-funded projects. Finally, such an institute should aggressively seek methods to supplement any NSF funding.