



Center for Trustworthy Scientific Cyberinfrastructure

Year Three Report

NSF ACI Grant # 1234408
Covering Project Year 3
October 1, 2014 - September 30, 2015

CTSC Team

Andrew Adams¹, Jared Allar¹, Jim Basney³ (co-PI), Randy Butler³ (co-PI), Robert Cowles⁶, Jeannette Dopheide³, Terry Fleury³, Randy Heiland², Elisa Heymann⁴, Craig Jackson², Scott Koranda⁵ (co-PI), Jim Marsteller¹ (co-PI), Prof. Barton Miller⁴ (Senior Personnel), Susan Sons², Amy Starzynski Coddens², Von Welch² (PI)

CTSC Students

Vineeta Sangaraju², NaLette Brodnax²

¹Carnegie Mellon University/PSC

²Indiana University/CACR

³University of Illinois/NCSA

⁴University of Wisconsin

⁵University of Wisconsin-Milwaukee

⁶Independent Consultant

This report describes work supported by the National Science Foundation under Grant Number OCI-1234408. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

For updates to this report and other reports from CTSC, please see <http://trustedci.org/reports/>

Executive Summary

The Center for Trustworthy Scientific Cyberinfrastructure (CTSC) is transforming and improving the practice of cybersecurity and hence the trustworthiness of NSF scientific cyberinfrastructure (CI) and the science it produces. CTSC is providing the NSF CI community with cybersecurity leadership, expertise, training, and the nexus of a community for sharing experiences and lessons learned. The vision of CTSC is an NSF CI community in which each project knows where it fits in a coherent cybersecurity ecosystem, has access to the tools and expertise to enact a cybersecurity program that efficiently support science, participates in the sharing of experiences and collaboration between projects and is greatly benefited by leveraging services from universities, regional and national networks (e.g., CIC, SURF, Internet2).

This report covers CTSC project year three, from October 2014 through September 2015, during which time CTSC engaged with nine NSF CI projects, organized and hosted the 2015 NSF Cybersecurity Summits for Large Facilities and Cyberinfrastructure, developed and provided training in developing cybersecurity programs, secure coding, and incident response, provided the community guidance with dealing with vulnerabilities, and authored and submitted to NSF a section on cybersecurity for the NSF Large Facilities Manual. PI Welch also joined the InCommon Identity Federation Steering Committee as an advisor for research.

Over 180 individuals, representing over a hundred NSF projects, attended one of the three CTSC-hosted NSF Summits. The 2015 Summit continued to build community around a call for participation that resulted in the broader community presenting three training sessions and six plenary sessions. Two former NSF employees presented invaluable perspective on NSF's history with cybersecurity.

Through its three years, CTSC has now engaged with 22 NSF projects (9 new in year three), and trained nearly 300 CI professionals representing over 60 NSF projects. Those numbers include a significant impact on NSF Large Facilities, who comprised 7 CTSC engagees, 15 of the projects who have attended a Summit, and 15 of the projects benefitting from CTSC training.

This report describes all CTSC's activities in detail, concluding with a set of lessons learned by CTSC over its three years. The CTSC project funding is expiring at the end of 2015 and the project plans for the remainder of 2015 and proposal to sustain itself as a NSF Cybersecurity Center of Excellence are described.

Table of Contents

1. Introduction: CTSC Overview and Vision.....	5
2. CTSC Impact on the NSF Community	6
3. Engagements.....	7
4. Year Three Engagements	7
4.1 HUBzero	7
4.2 SciGaP.....	8
4.3 Ocean Observatories Initiative	8
4.4 Large Synoptic Survey Telescope	9
4.5 Gemini.....	9
4.6 National Ecological Observatory Network.....	10
4.7 perfSONAR	12
perfSONAR Vulnerability Management Practices Review.....	11
perfSONAR Code Review.....	12
4.8 CC-NIE Peer Review (U. Cincinnati and U. Pittsburgh)	12
4.9 U. Oklahoma CC-NIE Review	14
4.10 AARC.....	14
4.11 Network Time Protocol	15
4.12 Array of Things	17
5. Feedback from Previous Engagements.....	18
5.1 CC-NIE Peer Review (Utah/PSU).....	18
5.2 Long Term Ecological Research Network Office.....	19
5.3 Laser Interferometer Gravitational-Wave Observatory (LIGO).....	20
5.4 IceCube.....	21
5.5 Pegasus.....	22
5.6 Large Synoptic Survey Telescope (LSST).....	22
5.7 Gemini	23
6. Education, Outreach, and Training.....	23
6.1 Training	23

6.2 Operational Training.....	24
6.3 Secure Software Development and Analysis Training	25
6.4 Student Interns.....	26
6.5 Outreach	27
7. Leadership of NSF CI Cybersecurity.....	27
7.1 NSF Cybersecurity Summit	27
7.2 Guide for Developing Cybersecurity Programs and Large Facilities Manual Cybersecurity.....	28
7.3 CTSC Cybersecurity Program	29
7.4 CTSC Publications	29
7.5 CTSC Collaborations	29
8. CTSC Advisory Committee.....	30
9. Lessons Learned	31
9.1 Engagements are Essential.....	31
9.2 Engagements Require Flexibility and Innovation	31
9.3 The Summit is Critical to Community Building and Outreach	31
9.4 Venues for Delivering Training are Scarce	32
9.5 Templates Partially Address the Sharing Challenge	32
9.6 Leveraging Campuses is Possible to a Degree.....	32
9.7 Cyberinfrastructure has Its Own Security Challenges	33
9.8 Strong Community Ties, Operational Security Expertise, and Diverse Backgrounds Critical to Success	33
10. Post Year 3/No Cost Extension Plans	34
11. Conclusion.....	35
12. References.....	36

1. Introduction: CTSC Overview and Vision

The Center for Trustworthy Scientific Cyberinfrastructure (CTSC) is transforming and improving the practice of cybersecurity and hence trustworthiness of NSF scientific cyberinfrastructure (CI) and the science it enables. CTSC is providing readily available cybersecurity expertise and services, as well as leadership and coordination across a range of NSF scientific CI projects via a series of engagements, best practices, online and in-person training, and the annual NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure.

As NSF pushes towards its vision of “a comprehensive, integrated, sustainable, and secure CI” as described in the Framework for 21st Century Science and Engineering¹, cybersecurity plays a key role. Yet the NSF CI community faces strong challenges in implementing cybersecurity. Projects are forced to divert their resources to develop appropriate expertise, address risks haphazardly, unknowingly reinvent basic cybersecurity solutions, and struggle with interoperability [S312]. Contributing to the challenge is the fact cybersecurity cannot be solved by a single solution. Every project has its own culture, risk tolerance, unique combination of cutting edge and legacy technologies, collaboration patterns, and timelines, making a “silver bullet” unfeasible. Even when security expertise is available within a project, the complex NSF CI ecosystem brings significant challenges in cross-project collaborations and knowledge dissemination. Lessons learned are shared haphazardly between projects. Important institutional knowledge is often lost when a project is completed or key personnel leave the community. Additionally, requiring each CI project to tackle cybersecurity independently is inefficient and often redundant, leading to multiple implementations that do not interoperate and confound the goal of scientific collaboration, data stewardship, and dissemination.

The vision of CTSC is an NSF CI community in which each project knows where it fits in a coherent cybersecurity ecosystem, has access to the tools and expertise to enact a cybersecurity program that supports science, participates in the sharing of experiences and collaboration between projects and is greatly benefited by leveraging services from universities, regional and national networks (e.g., CIC, SURA, Internet2).

Toward this vision, CTSC undertakes activities organized into three thrusts: 1) **Engagements** with specific communities to address their individual challenges and deepen CTSC’s knowledge of community requirements; 2) **Education, Outreach and Training**, providing the NSF scientific CI community with training, student education, best practice guides, and lessons learned documents; and 3) **Cybersecurity Leadership**, building towards a collaborative, coherent, interoperable cybersecurity community and ecosystem.

¹ https://www.nsf.gov/about/budget/fy2012/pdf/40_fy2012.pdf

2. CTSC Impact on the NSF Community

In this section, we present key metrics summarizing the impact of CTSC’s activities on the NSF community over its first three years. Subsequent sections of this report describe the activities in detail.

Method of Impact	Total # of NSF Projects & Facilities	Total # of NSF Large Facilities	Total # of NSF Personnel
One-on-one engagements (completed and in progress)	22	7	n/a
Training	74 individuals representing 63 projects	33 individuals representing 15 Large Facilities	15
Cybersecurity Summit Attendance	111	15	34

Table 1: NSF projects and personnel directly impacted by CTSC

Metric	Value
Training curriculum developed	5 ²
Training sessions provided	20
Number of in-person trainees	294
Online training videos	33
Number of views for online training	6,369
Number of individuals attending one or more cybersecurity summits	186
Number of views of blog posts (best practices, guidance)	13,990
Unique visitors to trustedci.org website	2,664
Number of technical reports, guidance publications, and published engagement products	41
Mentions in media, blog posts, etc.	2
Listed as a resource in an NSF solicitation	1
Invited talks	9

Table 2: Other CTSC impact metrics

² Does not include advancement of the Secure Coding tutorial developed by Prof. Miller prior to CTSC’s inception and revisions of the Cybersecurity Program development.

3. Engagements

One of CTSC's main activities is an ongoing set of engagements with NSF-funded scientific CI projects to solve cybersecurity challenges faced by those projects. During the third year, CTSC undertook new engagements with the Gemini Observatory, perfSONAR, the National Ecological Observatory Network, the Ocean Observatories Initiative, the Network Time Protocol project, the Authentication and Authorisation for Research and Collaboration, the U. Oklahoma CC-NIE project, the Array of Things, and orchestrated a peer review between two CC-NIE projects at U. Pittsburgh and U. Cincinnati. Additionally CTSC had the following engagements started in year two: Hubzero, SciGaP, and the Large Synoptic Survey Telescope.

In this section we describe each of the engagements in turn, including the resulting benefits for the engaged projects and the broader scientific community. Importantly, all CTSC engagement plans call for follow-up contact with engagement communities to assess the impact of the engagements. For all the listed engagements, plus engagements completed in prior years, we solicited and included a statement from the project regarding the engagement and its impact. Comments from the projects are included verbatim with no modification by CTSC.

4. Year Three Engagements

The following engagements were undertaken in year three. Some represent engagements started in year two.

4.1 HUBzero

As described by their website³, HUBzero is an open source software platform for building powerful Web sites, or "hubs" that support scientific discovery, learning, and collaboration. HUBzero was originally created by researchers at Purdue University in conjunction with the NSF-sponsored Network for Computational Nanotechnology to support nanoHUB.org. The HUBzero platform now supports dozens of hubs across a variety of disciplines, including cancer research, pharmaceuticals, biofuels, microelectromechanical systems, climate modeling, water quality, volcanology, and more.

In April 2014, CTSC conducted a "cybercheckup" for HUBzero, a short, focused engagement to identify gaps in an existing cybersecurity program. CTSC's broader engagement with HUBzero kicked off in September 2014 with a close review of their Web Server Security Model and Disaster Recovery Plan documents. CTSC provided HUBzero with recommendations to improve upon the existing policy and procedures to better address issues of access control and incident response. The engagement wrapped up by proposing a framework to develop a Content Management System (CMS) Access Control and Security Model to complement the Web Server Security Model and reviewing resulting documentation created by HUBzero. The process was both one of discovery – codifying ad-hoc practices into a document something that can be analyzed and improved – and analysis with the goal of producing

³ <https://hubzero.org/> (much of this paragraph is quoted from that website)

actionable recommendations for HUBzero to get the most return on their continuing efforts to secure their core software offering.

The engagement additionally generated a Vulnerability Management policies and procedures document for both the operations of HUBzero's hosted hubs and the development of HUBzero's CMS software. This work generated templates that can be applied by other projects in conjunction with CTSC's *Guide to Developing Cybersecurity Programs for NSF Science and Engineering Projects*⁴ (the Guide).

4.2 SciGaP

Our engagement with the SciGaP⁵ project continues to be longer-term and with less frequent interactions than most other CTSC engagements, with us providing consulting on an ongoing basis as needed as opposed to undertaking a pre-defined task. This style of interaction represents an experiment in engagement methodology by CTSC given SciGaP's point in their lifecycle as a software project developing their first product. The approach seems to be working and we plan to apply it in future similar situations (e.g., S12 Software Institutes).

This past year we focused on authentication and authorization architecture design. As discussed briefly in the Year 2 report, the primary challenge SciGaP faces is needing to authenticate and authorize clients with different trust models and ranging from traditional web portal-based science gateways to desktop applications and mobile apps. CTSC staff, including Heiland, Koranda, and Welch, met with SciGaP project personnel (both face-to-face and teleconference) during the past year to hear SciGaP detail use cases, explore trust models, and discuss the balance between operational and sustainability concerns with the desire for a simpler authorization architecture. CTSC helped the SciGaP architects explore and understand the details of various OAuth2 grants and the roles of the various actors involved, and helped map them onto the SciGaP use cases. One immediate and concrete result of the engagement was a joint paper accepted for the Workshop on The Science of Cyberinfrastructure: Research, Experience, Applications and Models. The paper reviewed authorization options and concluded that OAuth2 is the best framework for SciGaP to build upon given the current needs and direction of the project

4.3 Ocean Observatories Initiative

CTSC and the Ocean Observatories Initiative (OOI) began an engagement in June 2015 to assist OOI development its cybersecurity program. With this primary goal in mind, CTSC committed to working with the OOI staff on this effort on a weekly basis through September 30, 2015.

Specific objectives for the CTSC-OOI engagement included:

1. Assist OOI in addressing NSF cooperative agreement terms.
2. Advise on the security stance of the OOI system architecture.

⁴ <http://trustedci.org/guide>

⁵ <http://www.scigap.org/> - funded by NSF ACI awards 1339774, 1339649, and 1339856

As a first step, CTSC provided OOI with a set of policy templates to develop to address NSF's Large Facilities cooperative agreement terms and conditions, to align with existing institutional policies, and to be implementable within OOI's budgetary limitations.

The second phase of the engagement resulted in CTSC providing a report of observations and prioritized recommendations for advancing OOI's security posture, based on the information gathered through an in-person meeting with OOI's cyberinfrastructure team at Rutgers University, as well as emails and a preliminary conference call. The primary purpose of these recommendations was to identify potential gaps in OOI's nascent information security program where a redistribution or increase in effort/resources could more effectively reduce information security related risk to OOI's science mission. Additionally, at OOI's request, CTSC provided written feedback on a draft RFP targeted at the acquisition of security technologies.

CTSC supported OOI in its policy development efforts through the end of September 2015. OOI recently informed the CTSC team there has been some significant changes to the CI architecture since July and would like to continue working with CTSC after the current engagement ends.

4.4 Large Synoptic Survey Telescope

At the 2014 Cybersecurity Summit, Don Petravick approached CTSC to assist with developing a new Large Synoptic Survey Telescope (LSST) security plan based on CTSC's *Guide*. With a January 2015 deadline to provide a progress report to the NSF and CTSC already having other commitments, CTSC committed to meet with LSST on a weekly basis through the end of 2014 to help rework their security plan. This was another example of CTSC's experimentation with engagement methodologies, and tested the viability of NSF projects using the Guide with minimal assistance.

The effort was extended through the end of January 2015, but otherwise the engagement was successful. LSST carried out the planning effort, with CTSC acting in an advisory role. LSST provided CTSC with an opportunity to observe a project's utilization the Guide, its templates, and tool, and supplied constructive feedback for future versions. At the completion of the five month engagement, LSST had a revised cybersecurity plan that included a Master Information Security Policy, Acceptable Use Policy, Incident Response Policy, and a risk assessment based on the current and planned project environment.

4.5 Gemini

In June 2015, as a precursor to a forthcoming full engagement, Gemini Observatory and CTSC undertook a brief, but very intensive "cybercheckup" –style engagement⁶. Using Indiana University's REDCap⁶ service, a web application for secure online surveys, CTSC developed a questionnaire designed to gather key pieces of information regarding the information security program at large-scale NSF projects and facilities. Gemini personnel completed this questionnaire, and met with the CTSC engagement team on two occasions, to discuss the cybercheckup process and provide more detailed information. In early July,

⁶ <https://redcap.uits.iu.edu/>

CTSC delivered a report to Gemini with recommendations for the Gemini information security program, prioritized by CTSC's estimated cost and impact in implementing the recommendations. Following the 2015 NSF Cybersecurity Summit, the CTSC and Gemini teams met in person in Arlington to debrief. We learned that report was well-received, shared within the Gemini organization, and has been utilized to produce a prioritized cybersecurity project plan. We also learned that CTSC's framing of information security as a programmatic activity has been critical for dovetailing cybersecurity activities with Gemini's project management processes.

"I feel very fortunate to have the resources of CTSC available to Gemini Observatory as we develop a more mature, comprehensive "v2.0" cybersecurity program. The breadth and depth of knowledge and experience that the CTSC team has contributed thus far is vast, and has been key in gaining budgetary and Directorate support for cybersecurity initiatives." -- Tim Minick, Information Technology Services Manager, Gemini Observatory

Gemini and CTSC will use the results of our report and Gemini's planning efforts to structure and make the most of our Fall 2015 full engagement. CTSC and Gemini will select particularly challenging, resource intensive, and high impact projects to make maximum progress in our engagement window.

4.6 National Ecological Observatory Network

The National Ecological Observatory Network (NEON) is a nationwide network of ecological sensors and observation facilities sponsored by the National Science Foundation to gather and synthesize data on the impacts of climate change, land use change, and invasive species on natural resources and biodiversity. NEON collects data from over 80 land and water based sites across the U.S. and standardizes this data for use by scientists.

CTSC, in collaboration with NEON team, performed a cybersecurity risk assessment on the NEON network of sensors and data servers. The results of this assessment will be used to develop a cybersecurity plan for the NEON project. The engagement commenced in March 2015 and was completed in August 2015. CTSC personnel conducted this review using a combination of CTSC assessment methodologies designed to fit the scope and objectives of the review. CTSC personnel interacted closely with NEON personnel to perform this engagement.

The goals for the engagement with NEON was to:

1. Generate a list of threats, vulnerabilities, estimates for likelihood, and impacts
2. Review and prioritize lists into risks
3. Generate a high level cybersecurity plan for NEON's AOP and CI

The engagement began with a cybercheckup to get a rough assessment of the status of NEON cybersecurity. The cybercheckup was performed on NEON's (then) current cybersecurity program. NEON staff was asked to review "Securing Commodity IT in Scientific CI Projects"⁷ and see how well the

⁷ <http://trustedci.org/guide/docs/commodityIT>

recommended controls were applied to NEON's systems. The areas this checkup reviewed included policies and procedures, host protection, network security, physical security, and monitoring and logging. The result of this quick survey revealed that NEON was lacking many of the basic cybersecurity controls. This then led towards a more detailed Risk Assessment and Security Planning effort.

The formal Risk Assessment of NEON concluded that the NEON networks and system infrastructure are currently vulnerable to numerous risks. Out of the nearly sixty assets listed in the Risk Assessment Table, seven issues or concerns were ranked as having a "very high" residual risk, twelve had a "high" residual risk, and nineteen had a "medium" residual risk. The "very high" and "high" risk issues were related primarily to PII (personally identifiable information), and access to the networks located at the Observatory Sites and Domain Support Facilities that would allow an attacker access to the Denver Data Center servers. Several issues identified in the engagement can be addressed by developing NEON policies and implementing formal operational processes and procedures. Other issues can be addressed by utilizing software solutions such as monitoring and vulnerability scanning software.

Working closely with the NEON team, CTSC concluded the risk assessment, transferred the skill of performing and maintaining a risk assessment table, and assisted the NEON team in documenting recommended cybersecurity controls that, when implemented, will mitigate the current level of risks for NEON. Considering that full operation of the NEON network is planned by 2017, an effective security strategy is critical to protecting and isolating data from external and internal threats.

4.7 perfSONAR

perfSONAR⁸ provides an appliance solution for running network tests across multiple domains. It is used extensively by the network research and education community, including numerous NSF CC-NIE awardees, with over 1300 deployments as of February 2015. Due to the complexity of the perfSONAR project, CTSC and perfSONAR undertook two engagements in parallel: one team addressing perfSONAR vulnerability management practices and the other reviewing perfSONAR source code for security weaknesses.

perfSONAR Vulnerability Management Practices Review

CTSC staff met with perfSONAR's core software developers at the end of March 2015 to discuss the Vulnerability Management Review engagement. We determined during this meeting the goals of the engagement, which were to review their vulnerability management both from the perspective of the development team and also as experienced by users of perfSONAR. The specific tasks of the engagement include:

1. CTSC team to read through perfSONAR user documentation to provide feedback from "fresh eyes" on the expectations established regarding maintenance and updates.
2. CTSC team and perfSONAR to walk through perfSONAR's current vulnerability management process -- especially noting any differences between how various packages are handled -- with CTSC team to provide perfSONAR with recommendations for process improvement and automation.

⁸ <http://www.perfsonar.net/>

3. CTSC team to look at web100 kernel patch, seek a workaround to eliminate the need to use it, or (if no workaround is found) investigate the possibility of getting that patch merged into the main kernel tree so that it can benefit from the kernel team's maintenance and testing resources.

CTSC team members addressed the first task by performing an installation of a perfSONAR node while following the online documentation to determine if there were any major problems that stood out, and if the documentation encouraged the “set and forget” mentality. In addition, the CTSC team reviewed perfSONAR's current informal process for handling vulnerabilities, as described in task #2. Based on the findings of the first two tasks, the CTSC team delivered a recommendation for a more formal process.

In addition, we were able to review the architecture and its documentation for the purpose of reducing the attack surface of perfSONAR nodes and reducing the complexity in the vulnerability management process. CTSC team members were able to find specific parts of the node architecture that could benefit from better access control in order to reduce the attack surface. CTSC has reported these findings to the perfSONAR team.

perfSONAR Code Review

CTSC staff had a kick off meeting with a subset of the perfSONAR core software developers in late April to discuss the Code Review engagement. During this meeting, CTSC got a high-level overview of the key software components of perfSONAR. CTSC also got an update on the location and layout of software repositories and documentation relating to the key components. At the end of the meeting, we reached agreement on a prioritized list of software components that would benefit from a code review. The primary goal was to have CTSC perform an independent, non-biased, fairly detailed analysis of at least one critical component of perfSONAR, looking for potential design weaknesses and software vulnerabilities that could compromise security. The initial component that would be analyzed in detail was the Bandwidth Test Controller (BWCTL)⁹. The analysis had two parts: 1) apply the First Principles Vulnerability Assessment (FPVA) methodology¹⁰, and 2) use the automated tools in the Software Assurance Marketplace (SWAMP)¹¹ to perform a static analysis of the code. Each of these has been completed and we are currently reviewing the details of the SWAMP results. A final report for the Code Review is expected by the end of October.

4.8 CC-NIE Peer Review (U. Cincinnati and U. Pittsburgh)

The NSF CC-IIE program will have, with anticipated 2015 awards, over 120 projects¹², a number that CTSC cannot hope to engage with individually in any reasonable period of time. Following on the new peer review process which CTSC undertook in 2014, CTSC held another peer review between University of Cincinnati and University of Pittsburgh in Summer of 2015.

⁹ <http://software.internet2.edu/bwctl/>

¹⁰ <http://research.cs.wisc.edu/mist/includes/vuln.html>

¹¹ <https://www.mir-swamp.org/>

¹² <http://www.nsf.gov/pubs/2014/nsf14521/nsf14521.htm>

Both institutions agreed it would be best for them to hold a day long video conference meeting during which they could present and discuss their project, security policies and controls. CTSC scheduled and mediated the meeting by WebEx video conference which included two staff from CTSC, a member of the Bro Center of Expertise¹³, the PIs representing both institutions and representatives of their security groups and network engineers, a total of 11 people. During this meeting the security groups from both institutions presented their security policies and controls in their CC-IIE upgrades and engaged in discussions. Topics of discussion during the meeting included:

- Configuration management
- Data Transfer Nodes
- Dual network homing hosts
- IPv4 vs IPv6
- Log management and analysis
- MAC Address authorization
- Multifactor authentication
- Network topology
- Physical security
- Security auditing
- User adoption
- User requirements

Both institutions also shared and reviewed documents related to their network upgrade including plans, policies, technical descriptions, and network maps. Both institutions have agreed to a follow up one hour conference call to take place in early November. This engagement was a success and both institutions said it was beneficial. We note an issue for improvement is that it was challenging to schedule two large groups to meet with each other and this process took up as much time as the meeting itself.

In response to our questions about the peer review process, Bruce Burton from University of Cincinnati submitted this feedback:

- What part of the peer review was most beneficial to you?
“The most beneficial part of the review for me was seeing how another institution was actively deploying their ScienceDMZ. I enjoyed seeing what use cases Pittsburgh was trying to address with their implementation and the challenges encountered.”
- What part of the peer review as most beneficial to your team?
“The team benefited most from reviewing another institutions security approach and seeing how they tackled security for their ScienceDMZ deployment. The review helped us realize some security aspects we need to strengthen.”
- What were some of the common themes you noticed?

¹³ <https://www.bro.org/nsf/>

“A common theme was that, we both have ‘moving targets’, in that we are both learning and modifying as we further our respective deployments. There are still many unanswered questions as we move forward.”

- What were some of the contrasts you noticed?
“Our ScienceDMZ deployment is one that is geared toward off-campus collaboration, where Pittsburgh’s seem to lean toward internal high speed transfers with HPC clusters.”
- Has the peer review influenced any immediate changes in your network upgrade plan?
“Our physical architectural layout will remain the same, we will however look at implementing some of the security procedures that Pittsburgh has brought forth.”

4.9 U. Oklahoma CC-NIE Review

In late spring of 2015 the CTSC initiated an engagement with the University of Oklahoma’s OneOklahoma Friction Free Network (OFFN), an NSF CC-NIE project, to provide guidance on the security plan development process and possibly perform a risk assessment. The OFFN project is quite new; they are still in the process of implementing a cyberinfrastructure security plan and are challenged by the mix of production and testbed activities on the same network infrastructure. Due to the relative youth of their project, the engagement evolved into more of a security consultation (similar to our experiences with SciGaP). We did review with OFFN their network and hardware diagrams and talked about general security related issues and configuration suggestions. We dedicated one session to a discussion about how XSEDE manages the security relationships between sites. We shared with OFFN the XSEDE Service Providers Best Practice Guide and presented XSEDE security policies, information sharing, and incident response coordination. We plan to re-engage with OFFN after they have implemented their cybersecurity plan to assess it.

4.10 AARC

The two-year Authentication and Authorisation for Research and Collaboration (AARC) project¹⁴ started in May 2015 to “develop an integrated cross-discipline AAI framework, built on production and existing federated access services.”¹⁵ The project team consists of 20 European partners, lead by the former Trans-European Research and Education Networking Association (TERENA) now known as GÉANT.

During a presentation about AARC at the Federated Identity Management for Research Collaborations (FIM4R) meeting¹⁶ in February 2015, attendees discussed the importance of coordinating AARC activities with representatives of US research cyberinfrastructure. As a result, CTSC established an engagement with AARC to gather input from US cyberinfrastructure projects on AARC-lead activities, disseminate

¹⁴ <https://aarc-project.eu/>

¹⁵ <https://aarc-project.eu/about/>

¹⁶ <https://indico.cern.ch/event/358127/>

training and other AARC project outputs to US cyberinfrastructure projects, and facilitate EU-US pilot project activities.

Since launching in May 2015, the engagement has included the following initial activities:

- Jim Basney (CTSC) participated in the June AARC kick-off meeting.¹⁷
- Basney/Koranda (CTSC) and Romain Wartel (AARC/CERN) led a requirements-gathering break-out discussion at the August Cybersecurity Summit which identified a common need for education materials and training on federated attribute release for science collaborations. This provided confirmation for the planned AARC-CTSC work on training development and dissemination.
- Basney (CTSC) assisted with organizing the October WISE Workshop ("Wise Information Security for collaborating E-infrastructures").¹⁸
- Basney/Koranda (CTSC) participation in AARC discussions and document drafts on attribute management, guest identities, technology options (X.509, SAML, OpenID), non-web authentication and credential translation.

Upcoming activities include:

- Federated identity and incident response sessions at the Internet2 Tech Exchange meeting¹⁹ in October.
- Workshop on Information Security for collaborating E-infrastructures²⁰ in October.
- AARC training workshop²¹ in October.
- REFEDS, eduGAIN, and FIM4R sessions²² at the European Workshop on Trust & Identity (EWTI) in December.

4.11 Network Time Protocol

The Network Time Foundation²³ maintains the Network Time Protocol (NTP) reference implementation, a nearly ubiquitous foundational component of the Internet and critical to Internet and cyberinfrastructure security. In February 2015, vulnerabilities in NTP²⁴ exposed organizational problems within the project that impaired its ability to implement, test and release fixes for these vulnerabilities. Specifically, it lacked suitable technical expertise, a solid software-testing environment, as well as the documentation and developer tools needed to on-board more volunteer help.

CTSC and Indiana University's Center for Applied Cybersecurity Research together supported an effort to remediate the problems with the greatest direct impact on the NTP software's security and create a stable base which the software's community could continue to maintain. We succeeded in:

¹⁷ <http://blog.trustedci.org/2015/06/aarc.html>

¹⁸ <http://blog.trustedci.org/2015/08/wise.html>

¹⁹ <https://meetings.internet2.edu/2015-technology-exchange/>

²⁰ <https://www.terena.org/activities/ism/wise-ws/>

²¹ <https://eventr.terena.org/events/2240>

²² <https://eventr.terena.org/events/2188>

²³ <http://www.networktimefoundation.org>

²⁴ <https://ics-cert.us-cert.gov/advisories/ICSA-14-353-01C>

- Migrating NTP’s development history from a proprietary repository with severe access limitations to a publicly-accessible git repository, including reconstruction of data that was obfuscated by previous unclear migrations between source control systems.
- Modernizing NTP’s build and test infrastructure in order to make it more stable and more accessible to developers. The build system reduction in complexity was itself incredible: 31,000 lines of kludgy, brittle code were reduced to 884 lines that were clean, modern, and reliable.
- Creating documentation suitable for onboarding new developers. Previous to CTSC’s intervention, NTP’s documentation was both incomplete and years out of date, a situation that crippled NTP’s ability to bring additional developers to bear on its problems.
- Significantly increasing the maintainability and security of NTP’s aging code base. While much more remains to be done, this created a solid base from which a community-supported fork of NTP, called NTPSec²⁵, grew to continue the work of maintaining and securing NTP.
- Building relationships with Linux Foundation’s Core Infrastructure Initiative (CII), the Internet Civil Engineering Institute, and the wider open source software community. These relationships will aid CTSC in our vulnerability and threat awareness efforts, and connect us to expertise and development residing outside of CTSC so that we can better advise the NSF projects we serve.

According to Mark Atwood, who’s taken responsibility for both NTP and NTPSec at CII:

"NTP is the protocol that keeps clocks in sync across the internet. Synchronized time is critical for security, database replication, data integrity, logging and debugging. The NTP protocol is one of the oldest still operating protocols on the internet, and the widely used reference implementation has been showing it’s age, and has become known as a source of many security and integrity problems. Because of the necessity and risks, NTP came to the attention of the Linux Foundation Core Infrastructure Initiative (CII).

The CII was pleased to discover that rescue and refactoring had already begun on NTP, in the form of the NTPsec project, initially funded by the NSF CTSC. NTPsec had already done the difficult work to migrate the code repository from BitKeeper to Git, which opened up the field of developers from a small handful to literally potentially thousands of skilled developers, and had started the work of migrating the build system from an obsolete autotools system to state of the art waf, which again opened up the field of reviewers and contributors.

Because of those efforts, funded by the CTSC and the CII, NTPsec has already attracted skilled developers, has been able to refactor away over half the code as unnecessary or obsolete, and has already received or discovered several critical security vulnerabilities and then has been able to promptly fix them.

The Linux Foundation, the CII, and our member companies would like to thank the CTSC for it’s initial funding of the NTPsec project."

²⁵ <http://www.ntpsec.org/>

4.12 Array of Things

The Array of Things²⁶ (AoT) is an NSF-funded urban sensing project, a network of hundreds of interactive, modular sensor boxes that will be installed around Chicago to collect real-time data on the city's environment, infrastructure, and activity for research and public use. This initiative has the potential to allow researchers, policymakers, developers, and residents to work together to evaluate and take specific actions that will make Chicago and other cities healthier, more efficient, and more livable.

The AoT project includes a growing number of collaborators intending to deploy test configurations in nearly twenty other cities in the U.S. and globally. Because all of the data will be published openly and without charge, it will also support the development of innovative applications, such as a mobile application that allows a resident to track their exposure to certain air contaminants, or to navigate through the city based on avoiding urban heat islands, poor air quality, or excessive noise and congestion. Additionally, the nodes will evolve over time with new sensors and processing capabilities. Thus a carefully developed privacy policy is essential and must both contemplate current capabilities and plans as well as governing future decisions about new features.

At the AoT project kickoff meeting, an interdisciplinary group of university researchers, city policy makers, and private sector participants met September 2-4, 2015 in Chicago IL. A breakout group met and discussed requirements around an AoT privacy policy. Staff from CTSC led the breakout section and edited the resulting report. We will continue to advise AoT with the development of their initial privacy policies.

²⁶ arrayofthings.us

5. Feedback from Previous Engagements

We followed up with projects which we previously engaged (or completed engagements early in our third year) and asked them for updates on the impact of those engagements. For projects which responded, their responses are included here verbatim.

5.1 CC-NIE Peer Review (Utah/PSU)

Feedback from Joe Breen, U. of Utah:

How much impact have the results of the peer review had in terms of the cybersecurity of your CC-NIE project?

"I am not sure how to correctly gauge the impact. We have definitely taken the discussions into account as we have continued along our path to securing the Science DMZ. I believe the impact for the UofU was to make the existing UofU processes more rich in thought and content, and, to point out some areas such as Researcher MoUs that needed a more direct focus and fuller implementation. The UofU continues to iterate."

Have you maintained ongoing discussions with the peer as a result of the review?

"Yes, [Ken Miller, Penn State University] and I have had follow-up discussions and have touched base on different ideas that we are respectively considering. We are tentatively planning to get together in August 2015 time-frame to review continued progress and approaches."

Any impressions on the positive and negative impacts of a peer review versus a direct engagement with CTSC?

"The peer review with Penn State, along with the guidance of CTSC and the Bro Center of Expertise was a very positive engagement. I cannot compare to only a direct engagement with CTSC. I can confirm that bouncing ideas with a peer institution was very productive in understanding different approaches to handling research faculty, understanding different approaches to various technical problems, and understanding different issues that the respective environments face. The guidance of the CTSC provided a nice framework and put together a context for focused discussions. The experience of the Bro Center of Expertise provided additional input, both from another security perspective, and a tool use perspective."

Anything you would suggest doing differently with regards to the peer review process?

"For future peer reviews, I might suggest some of the following: (NOTE: Our process had some of these and some seemed to grow out of the process itself. I am trying to capture the thoughts in a cohesive manner for reference.)

- *Start with a 1-2 pg description for each respective peers that:

 - a) *described the process (i.e. 3-4 weeks of (1) 1hr sessions ea. week)*
 - b) *outlined overarching goals of process, i.e. review Science DMZ security approach both individually and in terms of the holistic campus security approach**

- c) *described necessary documentation to have available before starting, i.e. specific Science DMZ diagrams, applicable campus infrastructure, Researcher MoU templates, applicable campus security policy, Cyberinfrastructure Plan, etc.*
- d) *described applicable/suggested personnel for calls, i.e. security personnel, HPC personnel, monitoring personnel, network personnel, architects, etc. -- NOTE: doesn't necessarily need all of these but the call personnel need to be able to represent the respective campus areas as much as possible*
- e) *requested a block of time that each group could commit to and coordinate up front with the respective peer for the length of the process*
- f) *described a "starter" list of questions across architecture, design, policy, and research interactions with which each respective group could come prepared to discuss -- perhaps each session focuses on one aspect or is more ad-hoc?*
- *If peers are not already familiar with each other, a very simple MoU of mutual privacy and limited disclosure may be of use*
- *A shared whiteboard for each session might be of use, especially for technical discussions and referencing specific documents.*
- *Have someone take notes and send out action items or key discovery points/observations after each session. Send a reminder before each session with the notes/points from the previous session, plus specific documents to reference for the upcoming session.*
- *Make sure the CIO/CISO/lead research governance group are aware of the peer review. Encourage them to think of the security in terms of their campus CI plan. Also encourage the thinking of the Science DMZ security model in terms of a holistic campus model."*

5.2 Long Term Ecological Research Network Office

Mark Servilla of LNO provided this statement:

"The LTER Network Office, based on the review and recommendations of the Center for Trustworthy Scientific Cyberinfrastructure, has now implemented a process for automated system and security patches. This procedure, in place since Winter 2014, removes human dependencies in deploying security updates on all core server operating systems. As part of this procedure, all firewall intrusion prevention rules are scanned and updated as necessary on a regular basis. Revision planning of all system and security procedures is now in process, with primary focus on isolation of core servers and development of a full incident response and mitigation plan. In addition, the LTER Network Office administrators are in direct communications with the University of New Mexico Information Technology staff to coordinate network intrusion detection and issue mitigation."

5.3 Laser Interferometer Gravitational-Wave Observatory (LIGO)

Warren Anderson, LIGO Identity and Access Management Lead, provided the following feedback:

“The interaction between CTSC and LIGO has been less direct in this past year than it has in previous years. However, we continue to see value from the CTSC’s engagement at large. In particular, the CTSC has taken a leading role in promoting the needs of research VOs in the international identity management community. Von Welch’s work on the InCommon Steering Committee has kept research participating as a first-class citizen in discussions of federated identity practices in the US and internationally. This has led to LIGO being included in a number of discussions on topics it is concerned with, including the importance of attribute release for research VOs, federated security incident response, and providing an Identity Provider of Last Resort for those researchers who do not, for whatever reason, have access to a suitably federated identity provider through their home institution. Scott Koranda has taken a more grass-roots approach toward the same issues, and his in-depth knowledge of how LIGO and similar VOs operate is key to bringing the pertinent issues to light. CTSC has also directly provided seats at the table for LIGO to participate in planning and discussion related to security and trust for research VOs, which we have unfortunately not been able to fully participate in, although we hope to have more availability to do so in the future.

One aspect of CTSC that I am unable to comment on is the direct interaction between CTSC and research VOs on matters of day-to-day cybersecurity. I suspect this is largely due to the fact that LIGO is large enough to have a CSO who takes responsibility and oversight for these matters on his own, and thus we do not need these more detailed interactions. However, it seems to me there is clearly a need to have a set of cybersecurity resources available for smaller VOs who might not have in-house expertise.

Finally, it may be that we have just missed it, but I think it would be useful to have a yearly “unmeeting” in the style of Internet2’s “Advanced Camp” for research computing people related to security issues. While there is the large NSF meeting each year, I find that more structured than what I had in mind. Basically, I think it would be useful to have research VOs sit down and simply discuss what they are doing and, more importantly, what they should or want to do but don’t have resources for. I think this would help identify clearly where individual research VOs and the CTSC could most profitably spend their limited resources.”

5.4 IceCube

Steve Barnet of the IceCube project provided the following feedback:

“The IceCube engagement with CTSC has been extremely important to the ongoing development of the IceCube Cybersecurity program. In particular, by initially engaging with us face to face and learning about our facility first hand, CTSC was able to provide us with feedback and guidance that was far more relevant to our facility and research programs than a generic set of security standards or compliance templates. This level of feedback has helped us to focus our resources on the most important systems first and develop plans for continuous improvement.

The most immediate outcome of our engagement with CTSC has been an improvement in our security posture. One of the first outcomes was been to re-write and clarify our dated security policies. While not the most glamorous item in our security toolkit, the policies establish the framework within which our security controls must operate. One our other significant initiatives was to initiate routine external security scans against our external networks. This identified many mid-level vulnerabilities and enabled us to address those vulnerabilities before they were actively exploited.

In addition, the ongoing engagement and community development via the NSF Cybersecurity Summits and online collaboration tools has proven itself invaluable. The opportunity to interact with peer facilities and learn from them as well as passing along lessons learned helps provide focus and energy as well as new ideas for securing our facility while allowing our research programs to retain maximum flexibility.

If there is any opportunity for enhancement, it would be to establish some mechanism for follow-up engagements as an occasional progress check and opportunity to ensure that our activity in this area is keeping pace with the rapidly evolving environment. Security is a moving target and competes with many other priorities so occasional followup engagements can help to keep security prominent in our normal operational activities.

To summarize, IceCube has seen a great benefit from our collaboration with CTSC. The advice, guidance, and support has proven immediately useful and we look forward to continuing this collaboration for years to come.”

5.5 Pegasus

“The Pegasus Team engaged with the CTSC team in Year 1 of the project, to explore the various scenarios under which the workflows executed using Pegasus access user credentials for data staging. In particular, the team focussed on the cases where the credentials are staged to the remote worker nodes along with the jobs themselves. This is required to allow the jobs to retrieve input data from remote data staging sites to the worker nodes when they execute, and push the output data to the data staging site after completion. One of the protocols that Pegasus supports for file transfer is SCP that does the data transfers over SSH. SSH does not support security delegation, and hence the private key used for data transfers is transferred to the worker node along with the job description and inputs.

Our engagement with CTSC focussed on the problem of how to avoid the storage of SSH credentials on the local filesystem of the worker nodes for the duration of job execution. The CTSC team came up with a set of recommendations some of which we plan to incorporate in Pegasus in the near future. During this exercise, the two teams also explored various alternatives that initially looked promising and feasible, but later had to be discounted on account of potential security holes or increased complexity of the system.

Overall, we characterize our engagement with the CTSC team as a success, as it has helped us identify and formalize various solutions. The associated engagement report prepared by the CTSC team, will serve as a blueprint on how to tackle this problem. We feel that its applicability is not limited only to Pegasus but to other systems that support distributed execution of jobs.”

5.6 Large Synoptic Survey Telescope (LSST)

Don Petravick, PI Dark Energy Survey Data Management and local PI for LSST at NCSA provided us with the following comment on their engagement experience for our blog post on their engagement²⁷:

“The project was under pressure to deliver an updated Cybersecurity program. CTSC understood our situation and provided a contemporary framework that was straightforward and practical to apply to our environment. With their support we were able to meet the deadline with a revised modern Cybersecurity plan.”

²⁷ <http://blog.trustedci.org/2015/06/large-synoptic-survey-telescope-lsst.html>

5.7 Gemini

Tim Minick, Information Technology Services Manager, Gemini Observatory provided the following feedback:

"As a multinational institution with unique cybersecurity challenges, under the purview of varied laws and regulations, we very fortunate to have the resources of CTSC available to Gemini Observatory as we develop a more mature, comprehensive "v2.0" cybersecurity program. During our initial engagement the breadth and depth of knowledge and experience that the CTSC team contributed was vast, and has been key in gaining budgetary and Directorate support for cybersecurity initiatives. Recommendations from the initial engagement have provided the basis for formulating and prioritizing Gemini's 2016 cybersecurity program. The second phase of our engagement in late 2015 will focus on a remodel of our legacy cybersecurity plan and address issues related to ICS/SCADA infrastructure, with tangible results emerging by the end of 2015. The benefits realized from attending the CTSC hosted NSF cybersecurity summit and the networking opportunities with other NSF facility cybersecurity staff easily justify the permanent inclusion of the summit in Gemini's training and professional development program." -

6. Education, Outreach, and Training

A key component of our mission to achieve more trustworthy NSF scientific CI is the development of new cybersecurity expertise through the creation, dissemination, and delivery of training and educational materials, and outreach to the community to make them aware of CTSC's services and improve the understanding of cybersecurity for science. Towards this end, CTSC undertakes a set of Education, Outreach and Training (EOT) activities.

6.1 Training

Training for NSF CI professionals is a significant activity within CTSC and currently takes the form of lecture-style training materials delivered in person by CTSC staff.

CTSC presented the following training, which is described in detail subsequently:

- Barton Miller and Elisa Heymann. Secure Coding Practices. NSF Security Summit, Arlington, Virginia, August 2015.
- Bob Cowles, Craig Jackson, Jim Marsteller & Susan Sons. Developing Cybersecurity Programs for NSF Projects. NSF Security Summit, Arlington, Virginia, August 2015.
- Randy Butler. Incident Response Training. NSF Security Summit, Arlington, Virginia, August 2015.
- Barton Miller and Elisa Heymann. Hacks and Counter Hacks: How the Bad Guys Think about Your Code and Some Defensive Techniques. Lockdown 2015 (statewide conference for Wisconsin government and academic IT leaders and practitioners), Madison, Wisconsin, July 2015.
- Barton Miller and Elisa Heymann. Secure Coding Practices (and Other Good Things). International Conference on Software Quality, Long Beach, Calif., March 2015.
- Barton Miller and Elisa Heymann, Secure Coding Practices and Software Analysis. Presented to the University of Wisconsin Division of Information Technology staff, HTCondor Project staff, and Software Assurance Marketplace (SWAMP) staff, April 2015.

6.2 Operational Training

At the 2015 NSF Cybersecurity Summit, Jim Marsteller, Craig Jackson, Susan Sons, and Bob Cowles presented two interactive half day sessions on developing cybersecurity programs for NSF science and engineering projects. This was an updated and expanded version of the successful training presented at the 2014 Summit.

Morning Session. This instructional morning session was based on a cybersecurity planning guide (see, trustedci.org/guide) developed with input from the Daniel K. Inouye Solar Telescope (DKIST) project, and in use at a number of NSF facilities and projects. Some of the topics covered include:

- Building or Improving an Information Security Program
- Unique and Critical Science Requirements, Constraints, and Security Controls
- Information Security Policies and Procedures
- The Role of Project Leadership and Risk Acceptance
- Establishing a Risk Management Approach to Information Security
- Defining, Identifying, and Classifying Information Assets
- The Role of Risk Assessments within the Program Lifecycle
- Baseline Controls and Best Practices
- Topical Information Security Considerations: Third-Party Relationships, Asset Management, Access Control, Physical Security, Monitoring, Logging, and Retention
- Program Assessment and Evaluation

Afternoon Session. The afternoon sessions focused on deep dives and discussion on two challenging areas. Our descriptions follow:

1. **Cybersecurity Program Governance, Risk Acceptance, and Intra-organization Communication.** In most organizations, the people writing code, maintaining the network, and administering systems have the most information about the organization's information assets and risks thereto. Most decisions about resourcing and risk acceptance, however, are made much higher up the chain, and the greatest concentration of information security expertise likely lies somewhere in between. Meanwhile, technologists and managers often have very different ways of thinking and communicating about information security issues. In this module, we'll talk about common failure modes in organizational management and communication around information security that can cause poor decisions in organizational risk management to be made on the back of bad information.
2. **Securing Novel Technologies.** Science often relies on specialized systems, including one-of-a-kind instruments and sensors, ICS/SCADA components, and custom software. Securing these systems requires more than applying industry best practices – by definition, mature best practices don't yet exist – it calls for technical analysis and communities of practice. In this module, we'll talk about helpful resources, and ways of tackling the security of these challenging systems.

6.3 Secure Software Development and Analysis Training

Developments in Year 3 in the area of secure software development and analysis include producing self-contained presentation modules to be used for online delivery, developing an slide organization infrastructure for the hundreds of presentations slides we have developed, and adding a major new module on the use of software assurance (static analysis for now) tools.

In preparation for producing online podcast versions of our tutorial materials, we have taken each technical area and divided them into 10-20 minute units of slides. Each unit is given an Objectives section that covers the goals of the unit and the background needed to benefit from it. Next comes a Motivation section, which provides concrete examples of problems being addressed by that unit. The remainder of the technical presentation for each unit (the bulk of the slides in a unit) have been updated based on our experience and feedback teaching these units. At this point, we have 8 units ready for presentation and video recording. As a side effect, the slides we use for live presentations are benefiting from these improvements.

As we develop new materials, we now have hundreds of slides organized in a variety of ways. Each presentation we make is based on a PowerPoint file that contains the slides for that presentation. As a result, the same slides, or versions of the same slides, appear in many different files. Keeping these materials up to date has become a challenge. As there are no commercial or open source tools for supporting our needs, we have developed a simple slide tool infrastructure, based on a simple interface and plug-ins to PowerPoint. In this infrastructure, each narrow technical area (1-10 slides, typically) is in a separate file. For each presentation, we have a configuration file that names the slides to be included in the presentation by file and slide number or slide type (for example, “motivation”, “example”, “quiz”, “exercise”). The result is that slides are stored only in one place, so there is a definitive version to keep up-to-date. Presentation files are generated as needed and easily updated (generating a large presentation file, with 100 or so slides, takes only a few seconds). As we get more experience with this infrastructure, we will share it more broadly on the project and beyond.

The major new development in the secure software development and analysis area is the new tutorial module on software assurance static analysis tools. This module is structured into four sections:

1. *Conceptual Basics*: In this section, we discuss the basics of code analysis, providing technical foundation to give the student the theory behind these analysis tools. It covers ideas such as syntax vs. semantic analysis, local vs. path analysis, single file vs. whole program analysis, soundness vs. approximation, and control flow and data analysis basics.
2. *Tool Specifics*: In this section, we discuss a variety of specific assurance tools. The current set of tools includes both commercial tools (such as Coverity) and open source tools (such as FindBugs), covering C, C++, and Java. In the future, we will add modules for scripting languages as well. Note that we have received permission from the commercial tool vendors to present screenshots and examples of use of their tools.

In addition to the tools, we have created ten small programs (that will fit on a screen) that include examples of specific code weaknesses in C and C++ and ten more in Java. These example programs were extracted from the Juliet test suite; really, it is more precise to say that they were untangled from the test suite, as the structure of that suite is quite complex, building all their tests into one monolithic collection. The result is that we can present the students with

small, self-contained, separately compiled and built test programs, each with a labeled weakness in the code, and each with a false positive version of the same weakness.

We then produced slides showing each of our example tools works on each of our sample programs. From this large collection of slides, we can produce a set of slides customized to match the interests and needs of a specific audience.

3. *Tool Use in the SWAMP*: In this section, we show how to apply the previously described tools to the sample programs in the SWAMP environment. This section ties the conceptual material to the examples in the context of the SWAMP. As a result, the students can immediately try out the ideas that we present without needing to acquire, install, or configure the tools.
4. *Hands-on Exercises*: The last section presents a hands-on exercise, where the students use one or more of the tools that we have presented on software that we provide to them. This section is just starting. Currently, they can use the sample programs that they saw in Part 2 of this module. We will start development of larger example software, with the first example coming from a simple web server that has software flaws added to it.

Our planned means of delivery for the tools and sample software is to use pre-configured virtual machine appliances that will run on a variety of student computers. This means of delivery will simplify the students' access to both the tools and sample software, and reduce the amount of class time needed to conduct these exercises. Of course, it also means that we do not have to provide computers to the students.

6.4 Student Interns

CTSC supports student interns in two ways: an ongoing internships working with CTSC throughout the year; and five student scholarships to participate in the NSF Cybersecurity Summit.

Two students worked as hourly interns with CTSC staff at Indiana University. Vineeta Sangaraju, an Informatics Master's student, worked with CTSC through Fall of 2014 and Spring of 2015. Vineeta worked on the HUBzero engagement, co-led a cybersecurity training session at Penguicon 2014, and assisted the team revising CTSC's cybercheckup process. NaLette Brodnax, a Master's student with the School of Public and Environmental Affairs, worked with CTSC during the Summer of 2015 on the engagement with NTP.

Five students participated in the 2015 NSF Cybersecurity Summit through scholarships offered by CTSC. These student scholars received valuable exposure to cybersecurity discussions in the context of NSF science. The feedback we received included the following statement:

The knowledge I gained at the NSF Cybersecurity summit was crucial in helping me develop my goals as a student. I personally valued the interaction I had with professionals who provided academic and career guidance, as well as their perspectives on current issues facing information security. Moreover, the openness and transparency expressed by those who have recovered from security breaches has given me a stronger

sense of community and partnership that I look forward to becoming a member of. Ultimately, a result of this conference, I have become challenged to start thinking more outside the box in terms of strengthening our methods of defense and offense methods in information security. I'm ready for the future up ahead.
Dora Baldwin (CSU, San Bernardino)

6.5 Outreach

CTSC undertakes outreach activities both to disseminate its work and to make NSF CI projects aware of its services. CTSC's outreach mechanisms include the CTSC website (trustedci.org), an ongoing blog covering CTSC's activities (blog.trustedci.org), and a Twitter account to disseminate both the CTSC blog posts and other cybersecurity news of interest to NSF CI projects (twitter.com/trustedci).

A highlight for the year was the appointment of CTSC Director and PI Von Welch to the InCommon Steering Committee as an advisor for Research. Through this, the CTSC Federated Identity Discussion List²⁸ was created to allow NSF-funded projects discuss NSF CI projects and InCommon, and federated identity.

CTSC made 15 presentations in Year 3, all of which can be found at: <http://trustedci.org/presentations/>.

7. Leadership of NSF CI Cybersecurity

A key challenge for CTSC is being responsive to community needs, while also staying ahead of emerging problems and providing leadership in addressing them. Over the course of its day-to-day activities, CTSC needs to lead the community towards a coherent, interoperable cybersecurity ecosystem while serving each individual project well. CTSC leverages a broad understanding of the NSF CI community to actively seek opportunities to align cybersecurity solutions for interoperability to better support collaboration.

7.1 NSF Cybersecurity Summit

Since publishing CTSC's Year 2 Report, in line with CTSC's mission to provide cybersecurity leadership and education to the NSF CI community, the center has organized and executed the 2015 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure. The NSF-funded annual Cybersecurity Summits served as a valuable part of the process of securing NSF-sponsored infrastructure and supporting NSF cybersecurity community.

The 2015 Summit took place from Monday, August 17th through midday Wednesday, August 19th, at the Westin Arlington Gateway with the theme of "Understanding the Information Assets that Enable Science." Ninety (90) people attended the event, including 18 NSF personnel and 5 student awardees. On August 17th, we offered a full day of training as a continuation from last year's program, based on

²⁸ <http://trustedci.org/ctsc-email-lists/>

strong training attendance and overwhelmingly positive feedback. CTSC personnel participated as trainers in 3 of the 6 training sessions, with the other three sessions being offered by members of the broader community.

Another continuation from last year was the call for participation (CFP). This year, response from the community to the CFP was so strong that we had more proposals than we could accept given the time and space. This response represented a significant increase in the summit content sourced from the CI community. Proposals from the CFP process included presentations, breakout and training sessions, and opportunities for student scholarships.

The second and third days were in plenary, with keynotes, panels, and speakers focused on both the key cybersecurity challenges facing Large Facilities and CI projects, and the most effective responses to those challenges. This year, like last, there were a number table topic discussions over lunch covering a variety of specific information security issues.

To bring a mature perspective on information security challenges the CI community faces, the program committee recruited George Strawn as this year's keynote speaker. George Strawn had a short industrial career, a long academic career (30 years at Iowa State University) and a long government career (24 years at NSF). Having been involved in many prior NSF Cybersecurity Summits, George discussed how information security has evolved over the past several decades and what he thinks the future holds for the CI community.

As of the time of the writing this report, the report of the 2015 Summit is in an early draft phase and will be published by the end of calendar year 2015. Information regarding the event (e.g., agenda, biographies, presentations) is most readily accessible at <http://trustedci.org/2015summit/>.

7.2 Guide for Developing Cybersecurity Programs and Large Facilities Manual Cybersecurity

In Year 3, CTSC continued to socialize and utilize its *Guide to Developing Cybersecurity Programs for NSF Science and Engineering Projects*, associated templates, and tools (available at <http://trustedci.org/guide>).²⁹ LSST utilized the Guide with limited assistance from CTSC in planning and developing an information security program. We also learned that NCAR has actively used the Guide, and has publicly promoted its use in Steve Beaty's talk at CLHS'15.³⁰ The Guide also directly informed our engagements with Gemini, LSST, NEON, and OOI. At the 2015 NSF Cybersecurity Summit, CTSC personnel updated and delivered a full day of training based largely on the Guide.

At the 2014 NSF Cybersecurity Summit, the NSF Large Facilities Office (LFO) approached CTSC regarding efforts to revise the Large Facilities Manual. The LFO asked CTSC produce an initial draft for an entirely new information security section of the manual. The new section would be based on the structure of the

²⁹ For more on history of the Guide, developed in the context of CTSC's engagement with DKIST, see our Year 2 report, Section 5.1 available at <http://hdl.handle.net/2022/20030>.

³⁰ <https://commons.lbl.gov/display/CLHS/CLHS+2015>

Guide, but tailored to the needs, capabilities, and processes of Large Facilities. CTSC delivered this draft, and it is presently under review by the NSF LFO.

7.3 CTSC Cybersecurity Program

Since the inception of CTSC the center has followed its own guidance and developed its own cybersecurity program. CTSC makes its program publicly available, along with supporting documentation, in order to both provide an example to the community and help establish the trust of potential engagees that their information will be appropriately protected.

Each year, CTSC reviews this cybersecurity program and is currently in the process of making minor updates to keep it up-to-date with changes in CTSC and the threat landscape. This program will be finalized and published by the end of December 2015.

7.4 CTSC Publications

CTSC's leadership efforts include the publication of papers providing both guidance to the community and stating opinions of direction to unify community approach. CTSC's contribution in project year three were:

- Guide to Developing Cybersecurity Programs for NSF Science and Engineering Projects, v1, Center for Trustworthy Scientific Cyberinfrastructure, August 2014. This guide along with supporting templates and training materials is available at <http://trustedci.org/guide>.
- Subsection 5.3, Information Security, NSF Large Facilities Manual. This draft section is under review by the NSF Large Facilities Office at the time of this report.
- "Securing Commodity IT in Scientific CI Projects: Baseline Controls and Best Practices." <http://trustedci.org/guide/docs/commodityIT>
- Randy Heiland, Scott Koranda, Suresh Marru, Marlon Pierce, and Von Welch. 2015. Authentication and Authorization Considerations for a Multi-tenant Service. In Proceedings of the 1st Workshop on The Science of Cyberinfrastructure: Research, Experience, Applications and Models (SCREAM '15). ACM, New York, NY, USA, 29-35. DOI=10.1145/2753524.2753534 <http://doi.acm.org/10.1145/2753524.2753534>
- Elisa Heymann, Barton P. Miller, James A. Kupsch and Vamshi Basupalli, "Bad and Good News about Using Software Assurance Tools", September 2015. Available as technical report. *Submitted for journal publication.*

7.5 CTSC Collaborations

CTSC continues to coordinate training with the NSF-funded Bro Center of Expertise (<https://www.bro.org/nsf/>) who provided training at the 2015 Cybersecurity Summit and participated in CTSC's Peer Review process between U. Pittsburgh and U. Cincinnati. The two projects continue monthly phone calls to coordinate and explore opportunities for collaboration.

Fostering interoperability between NSF CI and the global research computing ecosystem is a key goal in CTSC efforts. To that end, CTSC is maintaining strong ties with DOE, through key invitations to the NSF Cybersecurity Summit, and DHS activities, through Welch and Basney's role as co-PIs in the DHS

Software Assurance Marketplace. CTSC utilized the DHS SWAMP for their software analysis of perfSONAR.

We maintain good ties with the Internet2 and the Higher Ed community through Basney's participation on the InCommon Technical Advisory Committee and Welch serving as an advisor for research to the InCommon Steering Committee. Additionally our advisory committee, discussed in the next section, provides us with additional ties to these communities.

CTSC continues to work with the REN-ISAC³¹ to develop a membership class suitable for NSF projects.

CTSC further has a connection to the National Security Higher Education Advisory Board (NSHEAB) through the Cybersecurity Subcommittee, of which co-PI Butler is a member. This advisory board is hosted by the FBI, NHS, and CIA, and includes eight members of higher education that work with the agencies to discuss and develop methods for broader sharing of actionable cybersecurity intelligence to the higher education community.

8. CTSC Advisory Committee

To make sure CTSC is well aligned with the needs of the NSF CI community, and in touch with the broader CI and cybersecurity communities, it is guided by an advisory committee. The committee was formed at the start of the project and meets twice a year, remotely in May (via teleconference) and in-person in November (co-located with the Supercomputing conference³²).

The CTSC advisory committee members are:

- Tom Barton is senior director for architecture, integration and chief information security officer at the University of Chicago.
- Neil Chue Hong is director of the Software Sustainability Institute (SSI), the UK national facility for cultivating world-class research through software.
- Don E. Middleton leads the Visualization and Enabling Technologies Section in NCAR's Computational and Information Systems Laboratory and currently serves as PI or co-PI on a number of projects, including the Earth System Grid, the Earth System Curator, the Virtual Solar Terrestrial Observatory, the North American Regional Climate Change Assessment Program, the Cooperative Arctic Data and Information Service, and NCAR's Cyberinfrastructure Strategic Initiative.
- Nicholas J. Multari is the senior project manager for research in cybersecurity at the Pacific Northwest National Lab (PNNL) in Richland, Washington.
- Nancy Wilkins-Diehr of the San Diego Supercomputing Center has a breadth of experience in community engagement. She is currently director of XSEDE's Extended Collaborative Support for Communities program, which includes Science Gateway initiatives. She is also the PI on a Science Gateway Institute conceptualization grant.

For full bios, please see <http://trustedci.org/advisory-committee/>.

³¹ <http://www.ren-isac.net/>

³² <http://supercomputing.org/>

9. Lessons Learned

CTSC continues to evolve its lessons learned, as first reported in its year one report, refined in its year two report and further refined in this report. The lessons follow (order is not meaningful).

9.1 Engagements are Essential

In addition to direct impact, CTSC's direct, typically one-on-one, engagements with NSF projects have proven essential for CTSC's maturation. While CTSC consists of cybersecurity professionals who have undertaken many risk assessments and developed numerous cybersecurity plans over their careers, engagements provide an opportunity to perform those tasks with a frequency and with a breadth of projects that would typically be impossible. This work provides an opportunity to experiment with different techniques and determine which approaches best serve the broader NSF CI community. It also keeps CTSC involved "on the ground" and prevents the project's work from veering toward the purely theoretical. We find that having at least one in-person meeting early in an engagement is critical to establishing effective teamwork.

9.2 Engagements Require Flexibility and Innovation

Having completed more than a dozen engagements, CTSC has begun to discern the factors that substantially impact the best form for an engagement, including the following:

- at what point is the project in its lifecycle;
- is the project focused on a specific scientific problem or domain, or is it providing general purpose infrastructure;
- is the project developing software, operating infrastructure, or both;
- does the project have an existing cybersecurity program;
- how large and complex is the project;

CTSC has learned to try different engagement models (e.g., peer reviews, "cyber checkups") in order to adapt to different types of projects. As these models prove useful, we then work to institutionalize them in CTSC with well-defined processes so we can execute them efficiently.

Even when a project and engagement approach is well understood, unexpected events (e.g., events that require the engaged project to re-prioritize temporarily) require flexibility in managing the engagement. To adapt to unexpected events, we recognize that our engagement teams will sometimes have spare effort due to being blocked, as well as the need for additional effort. To allow for flexibility, CTSC maintains an ongoing task to develop training materials, best practices and other deliverables with flexible deadlines. This allows staff to be applied to or from those deliverables and time-sensitive engagement tasks.

9.3 The Summit is Critical to Community Building and Outreach

CTSC has now hosted two NSF Cybersecurity Summits for Large Facilities and Cyberinfrastructure. These events have been invaluable both in terms of building a community among NSF projects y, and making the NSF community aware of CTSC. Relationships formed at and around the summits have resulted in

several of CTSC's engagements. As discussed in the following lessons, the summits have also been a valuable venue for CTSC to deliver training.

9.4 Venues for Delivering Training are Scarce

There are not many venues that offer opportunities either to provide or receive cybersecurity training targeted to the needs of our community. Many venues face a challenge in making time for specialized topics such as cybersecurity. While CTSC has had some success with Supercomputing and XSEDE (primarily with Secure Coding), the Summit remains the main venue for CTSC delivering training.

The training at the Summit and other venues has been well received. This leads to the consideration that an event for delivering training to CI professionals by CTSC and other projects across a range of specialized topics (e.g., data management, software engineering) could be well received by the community.

9.5 Templates Partially Address the Sharing Challenge

CTSC seeks to have as broad an impact as possible by sharing the work products of its engagements with the whole NSF CI community. However, projects are sometimes reluctant to allow this. We have had some success in the past year with the paradigm of developing a project-neutral template to address a relevant cybersecurity issue and then using that to complete the engagement objectives with a project. A template, while not a complete replacement for example cybersecurity plans, does serve as a valuable, easily shared resource.

9.6 Leveraging Campuses is Possible to a Degree

As we described previously, every NSF CI project with which we have worked is embedded in and leverages varying degrees of the commodity IT infrastructure, cybersecurity infrastructure and cybersecurity policies of the university or organization that hosts it. CTSC has been trying to answer the questions regarding the degree and circumstances in which projects can leverage this existing campus policy and infrastructure. While still not completely understood, some facets of the answers are starting to emerge:

- Commodity services such as vulnerability scanning and licenses for static analysis tools are sufficiently generic to be readily used by projects.
- Campus security offices tend to understand compliance-based security, so a project with HIPAA-covered data or social security numbers will likely find policies or infrastructure they can leverage.
- Due in part to the NSF CC-NIE/IIE program, networks tuned for science (e.g., Science DMZs) are increasingly available and may be of benefit to projects with large data movement needs.
- In general, campuses are not well positioned to provide comprehensive information security plans and programs for complex, large scale, often multi-institutional science projects.

9.7 Cyberinfrastructure has Its Own Security Challenges

In applying best practices from the broader cybersecurity community (e.g., NIST), CTSC continues to identify challenges specific to the NSF CI community, from unique assets such as scientific data and instruments, to challenges such as a close relationship to institutions of higher education and research. In particular, CI has a threat model which is not clear at this point given the community's unique assets and complex institutional and infrastructural relationships. A common misconception that CTSC witnesses is projects that have no data confidentiality requirements assume this means they have no need for cybersecurity. Counters to this assumption are that project data may still have integrity requirements, their project reputation could be hard by compromises to the point it impacts the reputation of their science, and their infrastructure could be used to attack others.

9.8 Strong Community Ties, Operational Security Expertise, and Diverse Backgrounds Critical to Success

Since its inception, the CTSC team has represented a wealth of operation security experience, strong connections to NSF and other major science projects, and a variety of practical experiences in related domains (e.g., law, risk management) and communities (e.g., software development, scientific, military, corporate, government). With two years behind us, these differing connections and backgrounds have proven invaluable in being able to initiate and establish relationships needed to form engagements with diverse scientific communities represented by different NSF projects, as well as bring broader information security best practices to bear.

10. Post Year 3/No Cost Extension Plans

September 30, 2015 will mark the end of the CTSC grant. We have requested and received a no cost extension from NSF through September 30, 2016. Under that no cost extension, we plan to continue CTSC activities until the end of the 2015 calendar year.

After 2015, we hope to be continue serving the NSF community as a NSF Cybersecurity Center of Excellence under the Cybersecurity Innovation for Cyberinfrastructure (CICI)³³ and have submitted a proposal to that effect.

CTSC's planned activities under the no cost extension (Oct 1, 2015 - Sep 30, 2016, with most of the activity being before the end of the 2015 calendar year) are:

- We are in discussions with the NSF-funded Image Based Ecological Information System (IBEIS)³⁴ regarding a possible engagement regarding confidentiality issues in their animal data.
- Continuing our previously described engagements with the AARC, Array of Things, perfSONAR, SciGaP, U. Pittsburgh and U. Cincinnati.
- Attend a planned NSF Software Sustainability Workshop to discuss security issues around sustainability.
- Presentation of our Guide for Developing Cybersecurity Programs and overall approach toward improving trustworthy science at WISE.³⁵
- Publish the report for the 2015 NSF Cybersecurity Summit. Consider lessons learned for the 2016 NSF Cybersecurity Summit, which will be handed off to the NSF Cybersecurity Center of Excellence when awarded.
- Respond to any NSF feedback regarding our cybersecurity section for the Large Facilities Manual.
- Provide training in Automated Software Assessment Tools at the International Conference on Software Engineering and Data Engineering, San Diego, Calif., October 2015.
- Hold our advisory committee meeting on November 19th, co-located with the Supercomputing 2015 conference. We will focus this meeting on receiving feedback on CTSC's lessons learned and providing guidance to an anticipated NSF Cybersecurity Center of Excellence.

³³ <http://www.nsf.gov/pubs/2015/nsf15549/nsf15549.htm>

³⁴ <http://www.ibeis.org/>

³⁵ <https://www.terena.org/activities/ism/wise-ws/>

11. Conclusion

This report covers CTSC's successful third year, during which time CTSC initiated nine engagements NSF CI projects directly (bringing its total over three years to 22), organized the 2015 Cybersecurity Summit for Large Facilities and Cyberinfrastructure, authored a cybersecurity section for the NSF Large Facilities Office's Manual, and provided training in secure coding, incident response and developing a cybersecurity program. CTSC impact on the NSF CI community has been impressive, with over 180 individuals, representing over 110 projects, attending one of three Summits, over nearly 300 CI professionals representing over 60 projects attending CTSC-led training. Those numbers include a significant impact on NSF Large Facilities, who comprised 7 CTSC engagees, 15 of the projects who have attended a Summit and benefitted from CTSC training. Seven students were exposed to cybersecurity and NSF science, two working directly with CTSC for multiple months. CTSC outreach through presentations, social media, publications, and leadership in InCommon continued to broadly inform and impact the community.

12. References

- [S3I2] Butler, R., V. Welch, J. Basney, S. Koranda, W.K. Barnett and D. Pearson. Report of NSF Workshop Series on Scientific Software Security Innovation Institute. 2011. Available from: <http://hdl.handle.net/2022/14174> [cited 12 Feb 2012]
- [Pegasus] Ewa Deelman, Gurmeet Singh, Mei-Hui Su, James Blythe, Yolanda Gil, Carl Kesselman, Gaurang Mehta, Karan Vahi, G. Bruce Berriman, John Good, Anastasia Laity, Joseph C. Jacob, Daniel S. Katz. Pegasus: a Framework for Mapping Complex Scientific Workflows onto Distributed Systems - Scientific Programming Journal, Vol 13(3), 2005, Pages 219-237.
- [Pegasus-CTSC] R.W. Heiland, S. Koranda, V.S. Welch, "Pegasus-CTSC Engagement Final Report," Center for Trustworthy Scientific Cyberinfrastructure, trustedci.org, May 2013. Available: <http://hdl.handle.net/2022/15562>
- [LIGO-CTSC] Basney, J., Koranda, S. Center for Trustworthy Scientific Cyberinfrastructure Engagement Plan: Final Report for LIGO Engagement. July, 2013. <http://hdl.handle.net/2022/16689>
- [LIGO-Three] Basney, J., Koranda, S. A Study of Three Approaches to International Identity Federation for the LIGO Project. July, 2013. <http://hdl.handle.net/2022/16760>
- [LIGO-eduGAIN] Basney, Jim; Koranda, Scott. InCommon Membership in eduGAIN: the LIGO Perspective. May, 2013. <http://hdl.handle.net/2022/16690>
- [Wilkins-Diehr] Wilkins-Diehr, N. A History of the TeraGrid Science Gateway Program: A Personal View. Proceedings of the 2011 ACM Workshop on Gateway Computing Environments (GCE '11). Nov 2011. doi:10.1145/2110486.2110488
- [Saltzer1975] Saltzer and Schroeder. The Protection of Information in Computer Systems, 1975. <http://web.mit.edu/saltzer/www/publications/protection/>
- [Saltzer2009] Saltzer and Kaashoek. Principles of Computer Design, 2009. <http://books.google.com/books?id=I-NOcVMGWSUC>
- [Smith2012] R.E. Smith, A Contemporary Look at Saltzer and Schroeder's 1975 Design Principles, 2012. <http://cryptosmith.com/node/365>