Report of the 2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

Large Facility Cybersecurity Challenges and Responses

August 26 - August 28

Westin Arlington Gateway - Arlington, VA

http://trustedci.org/2014summit

December 19, 2014

For Public Distribution

Craig Jackson, James Marsteller, Von Welch

Acknowledgements

The organizers wish to thank all those who attended the summit. Special gratitude goes to all those who responded to the CFP, spoke, provided training, and actively participated, including the 2014 Program Committee (highlighted in Section 3), without whom the event would not have been the same. Our sincere thanks goes to the National Science Foundation and Indiana University's Center for Applied Cybersecurity Research for making this community event possible.

The authors wish to extend particular thanks to the following individuals for their detailed assistance with this report: Ardoth Hassler, Drew Simshaw, Jim Basney

This event was supported in part by the National Science Foundation under Grant Number 1234408. Any opinions, findings, and conclusions or recommendations expressed at the event or in this report are those of the authors and do not necessarily reflect the views of the National Science Foundation.

About this Report

Drafts of this report were circulated for comment to the Program Committee (November 5, 2014) and summit participants (December 3, 2014).

Citing this Report

Please cite as: Craig Jackson, James Marsteller, Von Welch. Report of the 2014 NSF Cybersecurity Summit for Cyberinfrastructure and Large Facilities: *Large Facility Cybersecurity Challenges and Responses*. http://hdl.handle.net/2022/19244

For the latest information on the Summit

Please see, http://trustedci.org/summit/

Table of Contents

Executive Summary 3

1 Background: Prior Summits, the Evolving Cybersecurity Landscape, and Advancing the Community 5

2 The Summit's Purpose, Scope, and Theme 6

3 The Organizing and Program Committees 7

4 The Call for Participation and Program 8

5 Participants 9

6 Attendee Evaluations 13

6.1 Attendee Survey 13

6.2 Tutorial Evaluation 14

7 Findings and Recommendations 15

8 Closing Thoughts from the Organizers 19

Appendix A: The Agenda

Appendix B: Biographies for Speakers, Program Committee, and Organizers

Appendix C: Call for Participation

Appendix D: White Papers Submitted in Response to the CFP

- Saurabh Bagchi, Fahad Ali Arshad, Gaspar Modelo-Howard, *Cybersecurity Operations in a Multi-Institutional Academic Setting: The NEES Story*.
- Lee Liming, Steven Tuecke, Ian Foster, and Kyle Chard, *Case Study: XSEDE Leverages Globus Nexus for Identity and Group Management*.
- Pascal Meunier and Michael McLennan, *Curbing Abusive Behavior of Science Gateways*.
- LTC David Raymond, *Cyber Threats: From APT to Non-malicious Insiders*.
- Stephen Schwab and John Wroclawski, *Managing Security Policies for Federated Cyberinfrastructure*.

Appendix E: Training Descriptions

Appendix F: Table Topic Summaries

Appendix G: Listing of Attendees and Organizations

Appendix H: Attendee Survey summary report

Appendix I: Tutorial Evaluation survey summary report

Executive Summary

The 2014 NSF Cybersecurity Summit for Large Facilities served two related goals: building a trusting, collaborative community, and seriously addressing that community's core cybersecurity challenges. It build on the success, findings, and lessons learned from the 2013 event, around the theme, *Large Facility Cybersecurity Challenges and Responses*. The Program Committee and community members drove the agenda[1], with responses to a Call for Participation resulting in 4 case study presentations, 3 panel topics, and all training sessions. The program included keynotes from the cybersecurity community at large, and presentations from key leaders from within the NSF community.

The 2014 summit took place in Arlington, VA, August 26th through midday August 28th. On August 26th, it offered a full day of training. The second and third days followed a workshop format designed to address the key cybersecurity challenges facing Large Facilities and the most effective responses to those challenges.

One hundred seventeen (117) individuals attended the summit, with 64 individuals -- over one half of all registrants -- participating in planning, speaking, providing training, co-authoring a CFP submission, and/or leading a lunch "table talk." These individuals represented 69 NSF-funded projects or facilities, including 14 Large Facilities. Attendee evaluations and feedback were overwhelmingly positive and constructive.

The following Recommendations derive from the summit's Findings, and reflect the successful processes implemented in 2014. They will drive planning (already in progress) for a 2015 summit and the Center for Trustworthy Scientific Cyberinfrastructure's leadership efforts. More detail is in Section 7.

Recommendation 1: The NSF CI and Large Facility community should define its own best practices for cybersecurity rather than anticipating detailed direction from NSF. Clearly setting our own standards will help protect us from compliance directives that are not as well suited to our community.

Recommendation 2: The NSF CI and Large Facility community should implement a risk-based approach to cybersecurity that leverages broader best practices as much as possible, while addressing and balancing the community's particular needs around unique scientific instruments, data, openness, multi-organizational relationships, and project lifespans.

[1] See, Appendix A or http://trustedci.org/2014summit/

Recommendation 3: The NSF CI and Large Facility community should identify and share best practices for how to successfully integrate security throughout and across project organizations.

Recommendation 4: The NSF CI and Large Facility community should develop a common understanding of how risk responsibility and acceptance practices are most efficiently and appropriately distributed among project personnel and stakeholders.

Recommendation 5: The NSF CI and Large Facility community should explore ways of collaboratively developing and maintaining cybersecurity programs, such as sharing materials, services, policies, practices, lessons learned, and collaborative/peer reviews.

Recommendation 6: The NSF CI and Large Facility community should continue to find ways of sharing real-time data in order to foster continuity of expertise and gain as much of an advantage as possible in defending ourselves. Existing cross-organizational mechanisms (*e.g.*, REN-ISAC, EDUCAUSE, Internet2) should be evaluated in terms of how they could be leveraged.

Recommendation 7: We recommend the NSF CI and Large Facility community undertake or support a research effort to increase understanding and communicate that knowledge or know-how for each of the following open questions:

A. What is the threat profile for our community, and can insights into threat actors and their motivations positively impact the efficiency and effectiveness of our cybersecurity programs and risk management processes?
B. When and how does privacy intersect with NSF CI cybersecurity efforts in terms of (i) legal and regulatory requirements; (ii) our community's norms, values, and stakeholder relationships; and (iii) being a barrier to and/or enabler of science?
C. How do we include and meaningfully address software assurance, quality, or supply chain in the context of the project cybersecurity programs, and the summit itself?

1 Background: Prior Summits, the Evolving Cybersecurity Landscape, and Advancing the Community

Cybersecurity is a fast-developing and challenging field for all organizations in our contemporary world. The challenge is amplified by the intersection of myriad factors, including rapidly changing technology; ever-evolving and diverse threats; lagging workforce development; economic challenges; asymmetries in the cost and difficulty of attack and defense; and the nascent state of cybersecurity practice in general.

The NSF CI and Large Facilities community has a unique opportunity to develop information security practices tailored to these needs, as well as to break new ground on efficient, effective ways to protect information assets while supporting science.

NSF awardees face distinct questions when initiating information security programs due to their projects' unusual, and often unique, combination of attributes: distributed, collaborative organizational structures and relationships with other entities (*e.g.*, campus); unique, costly scientific instruments; limited resources, talent availability, and timelines; diversity in communities and missions; open, yet irreplaceable scientific data with an unclear threat model; and the need for reproducibility and maintaining public trust in their resulting science.

Recognizing the diversity of projects and the evolving understanding of how best to comprehensively, but efficiently, address information security, NSF sets out its information security requirements for Large Facilities and FFRDCs in fewer than 250 words[2]. These terms describe a dialogue between awardees and program officers around appropriate information security programs for NSF projects, and lay out the rough contours of policies, procedures, and practices such programs should include. These terms also represent an opportunity for the community to chart its own course, but do little alone to guide awardees to specific plans or best practices.

Best practices are evolving, both with the NSF and the broader community. For example, NIST's recent publication of *Framework for Improving Critical Infrastructure Cybersecurity*[3] and work on the National Strategy for Trusted Identities in Cyberspace (NSTIC)[4] propose important new approaches for cybersecurity programs and identity management. However, best practices for the federal government, commercial companies, and even research labs and institutions of higher education, do not directly translate to scientific communities and computing

[2] http://www.nsf.gov/pubs/policydocs/cafatc/cafatc_lf212.pdf (Item 56)
[3] http://www.nist.gov/cyberframework/
[4] http://www.nist.gov/nstic/

infrastructure.

In addition to the cybersecurity efforts and experiences of individual NSF projects, and the research advances of the NSF Secure and Trustworthy Cyberspace (SaTC) community, NSF has recently funded cybersecurity resources for the NSF community in the form of the Center for Trustworthy Scientific Cyberinfrastructure (CTSC)[5] and the Bro Center of Expertise[6]. These resources provide focal points for aggregating experiences, and translating the work from the broader world into cybersecurity practices effective for NSF scientific computing.

As one of CTSC's major leadership initiatives, it has reestablished the NSF cybersecurity summits as a step toward reinvigorating the NSF cybersecurity community. Spanning six years from 2004 to 2009 and then re-instated in 2013, the annual NSF Cybersecurity Summits serve as a valuable part of the process of securing NSF scientific cyberinfrastructure (CI) by providing a forum for education, sharing experiences, and building community. For many attendees, the summits are unique opportunities to come together with their colleagues, to benchmark and debate cybersecurity best practices, and to receive practical, relevant training.

2 The Summit's Purpose, Scope, and Theme

The 2013 summit[7] was well received both as an educational opportunity and community-reviving event after a four-year hiatus. However, we organizers believed the summits could (and still can) go further, and support measurable progress on the following goals: establishing reasonable community norms for the scope, metrics, resources, and processes for developing and implementing cybersecurity programs; providing pragmatic levels of information security; and supporting scientific discovery.

Two findings of the 2013 summit served as overarching drivers for the 2014 event:

> *Finding 5*. The community should consider the cybersecurity needs of and relationship between Large Facilities and smaller cyberinfrastructure projects, as well as how (and if) the summit can effectively address both.
>
> *Finding 6*. The community needs to develop a better understanding of the expectations for their cybersecurity programs and how to meet those expectations.

[5] http://trustedci.org/

[6] https://www.bro.org/nsf/

[7] See the 2013 summit report, agenda, and more at http://trustedci.org/2013-nsf-cybersecurity-summit/

As such, we set out the dual purposes of the proposed 2014 summit and anticipated future summits as: (a) to support the development of a trusting, collaborative community; and (b) to substantially address that community's core cybersecurity challenges. For 2014, we determined to focus efforts around the theme, ***Large Facility Cybersecurity Challenges and Responses***.

Large Facilities were a natural focus for 2014, representing a massive investment of national resources which entail the production, maintenance, and use of valuable (and sometimes one-of-a-kind) information systems and data. At the same time, in many cases, Large Facilities' resources have enabled more mature, multi-faceted cybersecurity programs, with personnel experienced and expert in information security.

The 2014 summit took place Tuesday, August 26th through midday Thursday, August 28th, at the Westin Arlington Gateway near NSF. On August 26th, the summit offered a full day of training in response to 2013's strong training attendance and overwhelmingly positive feedback. The second and third days followed a workshop format designed to identify both the key cybersecurity challenges facing Large Facilities and the most effective responses to those challenges. The event brought together leaders in NSF CI and cybersecurity to continue the processes initiated in 2013: building a trusting, collaborative community, and seriously addressing that community's core cybersecurity challenges.

The remainder of this report outlines the summit's organizational process, the resultant program, details on attendance and participation, and results of attendees' evaluations of the event. The report concludes with Findings and Recommendations, and closing thoughts of the organizers.

3 The Organizing and Program Committees

The 2014 summit was funded by a supplemental grant to the CTSC project, and three members of that project (Craig Jackson, James Marsteller, and Von Welch) served as an organizing committee. We recruited a Program Committee (PC) made up of key stakeholders, including leaders from the NSF and broader cybersecurity community and leads from large NSF CI projects. The PC was to be responsible for setting the specific agenda and inviting speakers, selecting white papers and training programs for presentation at the summit, extending invitations to expert presenters, participating actively in the event itself, and laying the framework for successful post-summit evaluation and community support. Jim Marsteller served as chair of the PC, a role he held in prior summits. The PC held 11 meetings by conference call beginning May 5, 2014 and ending September 15, 2014. It conferred electronically both prior to and following this time period.

The 2014 PC members were:

- **Amy Apon**, Chair of the Computer Science Division of the Clemson University School of Computing, former Director of the Arkansas High Performance Computing Center, and past Chair of the Coalition for Academic Scientific Computation.
- **Anthony (Tony) Baylis**, Assistant Department Manager for the Computing Applications and Research Department in the Computation Directorate at Lawrence Livermore National Laboratory.
- **Michael Corn**, Deputy CIO and CISO for Brandeis University.
- **Barbara Fossum**, NEES deputy center director and former managing director of Purdue University's Cyber Center and Computer Research Institute.
- **Kelly Gaither**, Director of Visualization, Texas Advanced Computing Center.
- **Ardoth Hassler**, Associate Vice President of University Information Services & Executive Director, Office of Assessment and Decision Support at Georgetown University and former Senior Information Technology Advisor in the Office of the Chief Information Officer in the NSF Office of Information and Resource Management, Division of Information Systems.
- **William "Clay" Moody**, Computer Science PhD candidate and an active duty US Army Major stationed as an Army Fellow at Clemson University. Following his PhD studies, he has an appointment to the faculty of the Department of Electrical Engineering and Computer Science at West Point, the United States Military Academy.
- **Rodney Petersen**, interim Executive Director of the Research and Education Community Security Collaborative (previously known as SecuriCORE) and former Managing Director of the EDUCAUSE Washington Office and a Senior Government Relations Officer.
- **Mark Servilla**, Lead Scientist, Network Information System at LTER Network Office (LNO).

4 The Call for Participation and Program

The full agenda and biographies are attached to this report as Appendices A and B[8].

The PC solicited input on challenges and desired summit topics from the Large Facilities via the NSF's Facility Security Working Group (FacSec) and issued a call for participation (CFP) to the community requesting submissions in the form of: (a) white papers one to five pages in length, focused on unmet cybersecurity challenges or lessons learned, (b) one to two-page abstracts for proposed half and full-day trainings, or (c) student applications.[9] Additionally, the PC invited

[8] The full summit program is also available on the CTSC website, http://trustedci.org/2014summit/

[9] http://trustedci.org/cfp2014; see also Appendix C.

specific community leaders as well as experts from outside the community to give presentations and participate in panels directed at aspects of the challenges identified from input from the Large Facilities and in the submitted white papers.

The CFP represented a new direction for program planning in 2014, designed to elicit a greater degree of community participation in developing the agenda, executing the summit, and increasing our ability to identify summit findings that represent the concerns, successes, and aspirations of our community. All submitted white papers are collected in Appendix D. Ultimately, the CFP process proved a success, and drove a great deal of the resultant program, including 4 case study presentations, 3 panel topics, and all the training sessions.

On August 26, we offered a full day of training in response to 2013's overwhelmingly positive feedback and strong attendance. Descriptions of each training session, including slide sets for most are appended as Appendix E.[10]

August 27 and 28 followed a workshop format designed to identify both the key cybersecurity challenges facing Large Facilities and the most effective responses to those challenges. Highlights of the event included keynotes offered by Dr. Phyllis Schneck, Deputy Under Secretary for Cybersecurity, U.S. Department of Homeland Security, and Matthew Rosenquist, Cyber Security Strategist, Intel. In addition to the CFP-driven portions of the program, the plenary workshop saw significant contributions from NSF, as well as colleagues from the broader scientific and cybersecurity communities. On August 27, Program Committee members and community members led 5 "table talk" discussions during lunch, the content of which are summarized in Appendix F, and many attendees came together again on their own time for an informal dinner that evening.

5 Participants

As with prior summits, attendance was by invitation only, with registration fee, and was inclusive of the NSF CI and Large Facility community. Our invitation list was based on the invitation list from the 2013 summit, and was updated to account for changes in the community, suggestions from NSF staff, and speakers to address specific topics of the summit. We also expanded the list to extend invitations to every Large Facility for which we could identify an appropriate contact. In summary, the invitation list included those with direct cybersecurity responsibilities in NSF Large Facilities and CI projects, NSF project principal investigators, and other key stakeholders and risk owners to ensure that NSF cybersecurity evolves to address their needs. Additionally, we invited individuals from outside the NSF

[10] See also, http://trustedci.org/2014trainingsessions

community (*e.g.*, Department of Energy, Internet2, higher education) to avoid being insular, maintain and develop new relationships, and encourage infusion of additional perspectives.

One hundred twenty-two (122) individuals registered for the summit, and 117 attended (including speakers, tutorial presenters, panelists, and the program committee). A listing of the attendees and their affiliations is in Appendix G. Fifty-nine (59) attendees participated in the August 26 training sessions. Sixty-four (64) individuals -- over one half of all registrants -- participated in planning, spoke, provided training, co-authored a CFP submission, and/or led a lunch table talk. Twenty-four (24) attendees work at Large Facilities. Twenty (20) attendees work at the NSF.

The following 69 NSF-funded projects or facilities, including 14 Large Facilities (marked with "◆"), were represented at the summit:

- Advanced CyberInfrastructure for High Performance Data Intensive Computing
- Atacama Large Millimeter Array (ALMA)
- Blue Waters
- Bro Center of Expertise
- CC-NIE or CC*IIE projects (x 7)
- CI-SEEDS: Seeding the Next Generation Cyberinfrastructure Ecosystem
- CoCoA
- Collaborative Research: 100G Connectivity for Data-Intensive Computing at JHU
- COmanage
- Comet
- Cornell Energy Recovery Linac (ERL)
- Cornell High Energy Synchrotron Source (CHESS) ◆
- Cornell Laboratory for Accelerator-based ScienceS and Education (CLASSE)
- Center for Trustworthy Scientific Cyberinfrastructure (CTSC)
- Dark Energy Survey
- Data Observation Network for Earth (DataONE)
- Distributed Web Security for Science Gateways
- Daniel K. Inouye Solar Telescope (DKIST) ◆
- DMR-1332208 (at CHESS)
- EAGER: Report on International Data Exchange Requirements (RIDER)
- EarthCube Initiative Cyber-infrastructure for Geosciences
- Earthcube Building Blocks
- Extreme Science and Engineering Discovery Environment (XSEDE)
- Green Bank Telescope (GBT) (part of NOAO)
- NSF GEO-SciSIP-STS-OCI-INSPIRE 1249607, "Enabling Transformation in the Social

Sciences, Geosciences, and Cyberinfrastructure"

- Gemini Observatory ✦
- GENI-Global Environment for Network Innovation
- HTCondor
- HUBzero
- IceCube South Pole Neutrino Observatory (IceCube) ✦
- International Computer Science Institute (ICSI)
- INSPIRE
- International Ocean Discovery Program ✦
- NSF IRNC
- Laser Interferometer Gravitational-Wave Observatory (LIGO) ✦
- Large Synoptic Survey Telescope (LSST) ✦
- Long Term Ecological Research Network (LTER)
- MRI[11]: Acquisition of High Performance Computing Instrument for Collaborative Data-Enabled Science
- MRI: Acquisition of 100TF Graphics Processor Laboratory for Multiscale/Multiphysics Modeling
- MRI: Development of Data-Scope - A Multi-Petabyte Generic Data Analysis Environment for Science
- National Center for Genome Analysis Support (NCGAS)
- National Center for Supercomputing Applications (NCSA)
- National Ecological Observatory Network (NEON) ✦
- National High Magnetic Field Laboratory (Magnet Lab) ✦
- National Optical Astronomy Observatory (NOAO) ✦
- National Radio Astronomy Observatory (NRAO) ✦
- National Solar Observatory ✦
- National Superconducting Cyclotron Laboratory (NSCL) ✦
- Network for Earthquake Engineering Simulation (NEES) ✦
- Open Science Data Cloud
- Open Science Grid (OSG)
- Pittsburgh Supercomputing Center (PSC)
- San Diego Supercomputer Center (SDSC)
- SI2-SSI: Sustaining Globus Toolkit for the NSF Community (Sustain-GT)
- SI2-SSI: SciDaaS – Scientific data management as a service
- Stampede
- Texas Advanced Computing Center (TACC)
- Thirty Meter Telescope Observatory

[11] *i.e.*, Major Research Instrumentation

- Very Large Array (VLA) (part of NRAO)
- Very Long Baseline Array (VLBA) (part of NRAO)
- Web10G
- Wrangler

Finding 4 from the 2013 summit stated "Future program committees should take on gender, age, and racial/ethnic diversity in the community and summit attendance as a strategic imperative for future summits." The lack of gender, age, and racial/ethnic diversity at that event was objectively obvious and pointed out by several attendees. Moreover, we recognize that diverse participation is both a socially relevant outcome for NSF[12] and a particular challenge in the cybersecurity community in general[13]. Thus, we expressly addressed the topic with the PC, identifying two members to spearhead efforts (Baylis, Hassler), and the group sought to encourage diverse participation via the invitees, speakers, panelists, and PC itself. The CFP expressly gave priority to those students from groups underrepresented in the NSF information security workforce. We note that Baylis has specific experience in this area as chair of the Supercomputing Broader Engagement in 2008 and participated in that committee in 2009. Ultimately, the PC supported the participation of three outstanding student applicants: Jasmine Bowers, Christopher Gullo, and Paul Lordier.

In order to gather baseline data related to this diversity effort, 2014 registrants had the option to provide their ethnicity/race and gender/sex. The aggregated responses to the those items follow. Voluntary responses to these questions show:

Ethnicity / Race	
Asian or Southeast Asian	7
Black or African American	3
Hispanic or Latino	3
Native Alaskan or American Indian	1
Multiracial	0
White or Caucasian	77
Other Ethnicity	0
Other (space provided)	0
Prefer not to answer	7

[12] *See*, NSF GPG, Section II.C.2.d.i

[13] *See, e.g.*, *Agents of Change: Women in the Information Security Profession.* A whitepaper derived from the 2013 (ISC)2 Global Information Security Workforce Study. Available from: https://www.isc2cares.org/uploadedFiles/wwwisc2caresorg/Content/Women-in-the-Information-Security-Profession-GISWS-Subreport.pdf

No Answer Provided	23

Gender / Sex	
Female	17
Male	73
No Answer Provided	32

6 Attendee Evaluations

We sought attendee evaluations of the summit via two SurveyMonkey surveys. One survey gathered feedback on the summit generally; the other requested feedback specific to the August 26 training sessions.

6.1 Attendee Survey

A summary of the general survey results is appended to this report as Appendix H. The responses were generally very positive. We summarize the results of the general survey below.

Forty-four (44) attendees (approximately 38% of all attendees) responded to the general "Attendee Survey." The organizers did not submit responses, but the survey was open to all other participants. We did not request the names of respondents, and have redacted some information from the appended report to further protect the anonymity of respondents.

The quantified and categorical results (*e.g.*, rating scales, yes/no questions) were very favorable. Selections follow:

- To Question #5, "How would you rate your overall experience with the 2014 summit?" 100% of respondents selected "Good" or "Excellent."

- Regarding Question #7, "Was this summit better than what you expected, worse than what you expected, or about what you expected?," the summit at least met the expectations of 100% of respondents, exceeding the expectations of 84% of respondents.

- To Question #8, "How useful to your work was the information discussed at the summit?" 100% of respondents gave ratings of "moderately useful," "very useful," or "extremely useful," with 77% providing the higher two responses.

- To Question #10, “Would you like to attend future summits?” 88.64% responded “Yes,” with the remaining 11.36% responding “Maybe.”

Questions 11 and 12 sought open-ended responses, and were designed to elicit critique and discern highly-valued aspects of the experience. While the generally positive results of the above-referenced questions provide context, these open-ended questions have proved a useful communication tool. Observations follow:

- Question 11 asked, “How can we improve the summit experience in the future?”
    - Of the 23 respondents to this question, 6 suggested more opportunities for interaction among participants and cross-project benchmarking, particularly around sharing practical, usable information (*e.g.*, BoFs, more interactivity between panels and audience members). An example response follows:

      “More sharing by NSF CI projects about what is (and isn’t) working for them, what their top risks/concerns are, what their future plans are. The HUBzero presentation was an excellent example of what we need more of in future summits.”

- Question 12 asked, “Were there any aspects of the summit you found particularly useful or important? If so, please explain.”
    - Of the 26 respondents, 5 praised the panel discussions and 4 highlighted the training sessions as particularly useful or important.
    - Three (3) respondents noted the importance of NSF’s presence and contribution.
    - Three (3) respondents highlighted networking opportunities.

6.2 Tutorial Evaluation

The responses to the tutorial-specific surveys were very positive generally, and included constructive feedback, as well as ideas for future training offerings. For simplicity, we asked attendees to complete one survey with several repeated questions to allow sorting differentiated responses for morning and afternoon sessions. The aggregated ratings in Questions 1 through 10, and 13 through 18 are attached as Appendix I. We summarize a few aggregate responses below:

- To Question 3, “Based on your overall experience with the August 26 training sessions, would you participate in training offered at future summits?” 30 (*i.e.*, 86%) of 35 respondents selected “Yes,” 4 selected “Maybe,” and 1 selected “No.”

- To Questions 7 and 15, “How would you rate your overall experience with the [morning/afternoon] training?” 84% of responses were “Excellent” or “Good.”

- To Questions 9 and 17, “Was this [morning/afternoon] training better than what you expected, worse than what you expected, or about what you expected?” 93% of responses indicated that expectations were met or exceeded. Forty-seven (47%) of responses were “Quite a bit better” or “A great deal better.”

- To Questions 10 and 18, “How useful to your work was this [morning/afternoon] training?” 70% of responses were “Very Useful” or “Extremely Useful.”

The responses for the individual tutorials were filtered and reported back to their respective tutorial leaders, including responses to Questions 11 and 19, “How can we improve this training session in the future?” and Questions 12 and 20, “Were there any aspects of [morning/afternoon] training you found particularly useful or important? Please explain.”

7 Findings and Recommendations

The following Findings and accompanying Discussions are observations regarding the state of cybersecurity practice, challenges, and consensus in our community. They are based on the 2014 summit’s presentations, panels, discussions, and evaluations. For each finding, we provide related Recommendations to the NSF CI and Large Facility community.

*

Finding A: It is up to the NSF CI and Large Facility community to adopt baseline expectations and evaluative metrics for our cybersecurity programs.

> Discussion: Finding 6 of the 2013 summit stated, "The community needs to develop a better understanding of the expectations for their cybersecurity programs and how to meet those expectations." The discussion following this finding includes, "(T)he community is still not certain what the expectations are for a cybersecurity program or how they go about fulfilling those expectations ....(T)here is a subset of the community that expects NSF to provide greater clarity, while others believe we can make progress as a community." The panel on Large Facilities’ Cybersecurity Challenges and Success tackled that finding head on, and representatives of NSF clarified that, as a sponsoring organization, it provides guidance, but is not positioned to prescribe the precise structure of appropriate project cybersecurity programs or practices. As such, the NSF CI and Large Facility community faces both the challenge of determining baseline expectations and best practices, and

the opportunity to tailor these practices to our needs outside the confines of a compliance-oriented regime. Similarly, like the cybersecurity community more generally (as highlighted by Matthew Rosenquist), we face the challenge of having few usable outcome metrics by which to measure the success of a cybersecurity program. As such our best practices and processes form the most usable and reliable metrics by which to evaluate our programs.

Recommendation 1: **The NSF CI and Large Facility community should define its own best practices for cybersecurity rather than anticipating detailed direction from NSF. Clearly setting our own standards will help protect us from compliance directives not as well-suited to our community.**

*

Finding B: Risk-based approaches are the appropriate and increasingly dominant means by which NSF projects and the broader community address information security.

Discussion: Risk-based approaches to information security dominated discussions throughout the summit, including both keynote addresses and particularly during the panel on Large Facilities' Cybersecurity Challenges and Success. The NSF CI and Large Facility community is not bound by a highly prescriptive regulatory regime, and (like the information security community more broadly) is embracing programmatic approaches to information security risk that are more mature than purely technical or entirely *ad hoc* responses. Panelists discussed the utility of risk-based methods for determining when to accept residual risk versus push forward with additional controls, as contemporary risk-based approaches highlight managing risk rather than entirely eliminating it, and embrace identification of key assets, detection, response, and recovery in addition to prevention. These approaches account for the risk not only to information and information systems, but organizational interests such as science mission and reputation (as highlighted in the Threat Profile panel). While participants discussed a variety of case examples involving rare or unique instruments, data, and/or institutional relationships, risk-based processes appear sufficiently generalizable, flexible, and technology neutral, so as to serve as a common point of reference.

***Recommendation 2*. The NSF CI and Large Facility community should implement a risk-based approach to cybersecurity that leverages broader best practices as much as possible, while addressing and balancing the community's particular needs around unique scientific instruments, data, openness, multi-organizational relationships, and project lifespans.**

*

Finding C: Cybersecurity is a "whole-of-organization" endeavor, requiring input and buy-in both vertically (from PI's and directors to staff and users) and horizontally (*e.g.*, scientists, legal, IT, HR) across project organizations, and coordination with cooperating, hosting research institutions.

> Discussion: The roles of principal investigators, directors, and other project leaders vary considerably with respect to cybersecurity among NSF CI projects and Large Facilities. However, the importance of leadership involvement in cybersecurity emerged as a repeated theme of discussion (*e.g.*, in the panel on Large Facilities' Cybersecurity Challenges and Success, the panel on the Threat Profile for NSF Large Facilities and Cyberinfrastructure, and in Matthew Rosenquist's keynote address). The Large Facilities panel discussed that a critical issue for project leadership involves clarifying expectations and processes regarding who "owns" information security risk, including who has authority to accept information security related residual risk. Rosenquist and others highlighted the need for communication and education across departments in order bridge knowledge gaps in creating policy and ensure that security is effectively integrated into the project organization.

Recommendation 3: **The NSF CI and Large Facility community should identify and share best practices for how to successfully integrate security throughout and acrosss project organizations.**

Recommendation 4: **The NSF CI and Large Facility community should develop a common understanding of how risk responsibility and acceptance practices are most efficiently and appropriately distributed among project personnel and stakeholders.**

*

Finding D: Community building and information sharing must play an increasingly central role in supporting NSF projects as they develop and maintain their respective cybersecurity programs and practices.

> Discussion: In her keynote, Dr. Phyllis Schneck emphasized that trust is our #1 tool for shifting the advantage in favor of cybersecurity, from leveraging direct personal relationships to utilizing formal threat intelligence sharing systems. Many summit participants agreed; when asked an open question about how the summit can be improved, a quarter of respondents to the attendee evaluation survey suggested more

opportunities for sharing among themselves. The panel on The Role of Information Sharing in Large Facility Security discussed this range of relationship-leveraging practices as well, including the desire not only for more usable technical threat information, but also increased opportunities for sharing experiences and resources in terms of how to kickstart a cybersecurity program, governance, policy development, day-in-day-out practices, peer audits and reviews, as well as when and how to get outside assistance in handling security incidents. The panel and attendees explored ideas for progress ranging from a survey study of what threat intelligence sources are currently used by projects, to the possibility of standing up a specialized incident response team for NSF CI projects. Taking our broader community into account, NSF's Anita Nikolich highlighted the foundation's broader cybersecurity research and education initiatives, and other attendees and speakers (including Dr. Schneck and Purdue's Saurabh Bagchi) emphasized the potential for the NSF CI and Large Facility community to contribute to transition-to-practice and workforce development, as well as benefit from advances in the science of cybersecurity.

***Recommendation 5*: The NSF CI and Large Facility community should explore ways of collaboratively developing and maintaining cybersecurity programs, such as sharing materials, services, policies, practices, lessons learned, and collaborative/peer reviews.**

***Recommendation 6*: The NSF CI and Large Facility community should continue to find ways of sharing real-time data in order to foster continuity of expertise and gain as much of an advantage as possible in defending ourselves. Existing cross-organizational mechanisms (*e.g.*, REN-ISAC, EDUCAUSE, Internet2) should be evaluated in terms of how they could be leveraged.**

*

Finding E: In addition to the challenges already identified, the summit revealed several open questions (or *known unknowns)* where research is likely necessary in order to understand and communicate the relevant processes, implications, and applicability to our community.

***Recommendation 7*: We recommend the NSF CI and Large Facility community undertake or support a research effort to increase understanding and communicate that knowledge or know-how for each of the following open questions:**

D. **What is the threat profile for our community, and can insights into threat actors and their motivations positively impact the efficiency and effectiveness of our cybersecurity programs and risk management processes?**
    a. Discussion: These questions were the subject of a lively disagreement between a

keynote speaker and an attendee.

E. **When and how does privacy intersect with NSF CI cybersecurity efforts in terms of (i) legal and regulatory requirements; (ii) our community's norms, values, and stakeholder relationships; and (iii) being a barrier to and/or enabler of science?**
    a. Discussion: The discussion during the panel focused on privacy revealed the multifaceted nature of the privacy topic, and a lack of common understanding.
F. **How do we include and meaningfully address software assurance, quality, or supply chain in the context of the project cybersecurity programs, and the summit itself?**
    a. Discussion: An attendee pointed out that the summit devoted little if any discussion to software security.

8 Closing Thoughts from the Organizers

The 2014 summit was very well-received, and we believe the event fulfilled the dual purposes set out in the early planning stages: (a) to support the development of a trusting, collaborative community; and (b) to substantially address that community's core cybersecurity challenges. We again thank the Program Committee and all who responded to the CFP, spoke, provided training, and actively participated, for making the 2014 summit a success.

As organizers, our goal has been to push the 2014 and future summits to maximize their positive impact on cybersecurity for the NSF CI and Large Facility community, and we believe 2014 saw a number of improvements over the 2013 event. With the success of the CFP process, the program was more community-driven, and the program was even more deeply substantive than in 2013. The discussions benefitted a great deal from the presence of, strong participation from, and frank discussions with NSF program officers and personnel. The summit brought together many attendees, projects, and facilities who were not present in 2013. All this community engagement and depth have supported the drafting and vetting of a more detailed set of Findings and Recommendations. For CTSC, the summit was once again a forum for forming new relationships and an opportunity to plan new engagements, as well as a chance to socialize CTSC's *Guide to Developing Cybersecurity Programs for NSF Science and Engineering Projects*.[14] Our attendee surveys showed overwhelmingly positive evaluations of the event, as well as thoughtful critique and new ideas.

One of the most encouraging -- and yet most challenging -- things we observe is a strong desire in the community for more opportunities to share materials, services, practices, and lessons learned. We plan to address that desire in the 2015 summit, as well as consider how we can support these activities between the summits. In 2013, we set up the Trusted CI Forum as an

[14] http://trustedci.org/guide

online set of tools to help support community interaction and continuity from summit to summit. On its own, the Trusted CI Forum has not proven successful in fostering a great deal of community activity. As a result, CTSC is working with the REN-ISAC and other members of the community to determine more precisely what content, format, and fora will best meet the community needs, including increased opportunities for these types of interactions at the summit itself.

Diversity in attendance and addressing 2013 concerns became a strategic item for the PC for 2014. The 2014 summit was a great improvement over 2013 in terms of gender and age inclusiveness, in part due to the PC's focused effort and in part due to attendance by and participation from NSF personnel. We are determined to continue efforts to appropriately encourage diversity / inclusion in future summits, determine appropriate process and outcome metrics for this effort, and leverage the baseline data we collected as factual background for future discussions.

For the 2015 summit we will continue the successful process of program building by convening a program committee and issuing a call for participation. We hope to see even more of the agenda driven by community submissions. The focus of the 2015 summit will be addressing the 2014 Recommendations and documenting Large Facilities community progress. A secondary focus will be maximizing the positive impact on the broader scientific CI ecosystem by considering how Large Facility practices relate to medium-sized projects.

Appendix A
The Agenda

Program Agenda

2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

August 26 - August 28 Westin Arlington Gateway Arlington, Virginia

http://trustedci.org/2014summit/

Updated August 22, 2014

Organizers: James Marsteller, Craig Jackson, Von Welch

Training Day

Tuesday, August 26

http://trustedci.org/2014trainingsessions

7:00am	Registration and Continental Breakfast (Hemingway Pre-Function)
8:00am	Morning and All Day Training Sessions Begin ● Bro Platform Training Workshop ● Developing Cybersecurity Programs for NSF Projects ● Vulnerabilities, Threats, and Secure Coding Practices
9:45am	*Coffee Break*
10:00am	Resume
12:00pm	*Lunch provided*
1:00pm	Afternoon Training Sessions Begin and All Day Training Sessions Resume ● Bro Platform Training Workshop (continued) ● HPC, HIPAA, and FISMA: Meeting the Regulatory Challenge through Effective Risk Management ● Incident Response Training
3:00pm	*Coffee Break*
3:30pm	Resume
5:00pm	Sessions End
Evening:	*Dinner on your own*

continued on next page

Plenary Session

Wednesday, August 27
F. Scott Fitzgerald C

7:00am	Sign-In and Continental Breakfast (Pre-Function C)
8:00am	Welcome and Goals (Jim Marsteller)
8:10am	NSF Address: Irene Qualters, Division Director of Advanced Cyberinfrastructure (ACI)
8:30am	Keynote Address: Dr. Phyllis Schneck, Deputy Under Secretary for Cybersecurity, U.S. Department of Homeland Security
9:30am	*Coffee Break*
10:00am	Panel: Large Facilities' Cybersecurity Challenges & Successes *Panelists*: Steve Barnet (IceCube), Bret Goodrich (DKIST), Cliff Jacobs, Bill Kramer (Blue Waters) *Moderator*: Amy Northcutt (NSF)
11:00am	NSF/ACI Perspective on Cybersecurity (Anita Nikolich)
11:20am	CTSC Observations, Perspective, and Vision (Von Welch)
11:40am	NSF Bro Center for Expertise (Robin Sommer)
12:00pm	Lunch and Table Topics - *Lunch provided*
1:00pm	Case Study: Curbing Abusive Behavior in Science Gateways (Pascal Meunier, HUBzero)
1:30pm	Case Study: Cybersecurity Operations in a Multi-Institutional Academic Setting: The NEES Story (Saurabh Bagchi, Fahad Ali Arshad, Gaspar Modelo-Howard)
2:00pm	Panel: Privacy Concerns at Large Research Facilities *Panelists*: Mike Corn (Brandeis U.), Celeste Matarazzo (LLNL), Nigel Sharp (NSF), Heidi Wachs (Gartner) *Moderators*: Ardoth Hassler (Georgetown U.), Rodney Petersen (EDUCAUSE)
3:00pm	*Coffee Break*
3:30pm	Case Study: XSEDE Leverages Globus Nexus for Identity and Group Management (Steve Tuecke)

4:00pm	Case Study: Managing Security Policies for Federated Cyberinfrastructure (Stephen Schwab and John Wroclawski, USC ISI)
4:30pm	Open Discussion / Summary of the Day's Findings (Von Welch, Craig Jackson, Jim Marsteller)
5:00pm	*Adjourn for the Day*
Evening:	*Dinner on your own.* *Informal Dinner Gathering at World of Beer, 901 N. Glebe Rd., 6:30pm*

Plenary Session (continued)

Thursday, August 28
F. Scott Fitzgerald C

7:00am	Sign-In and Continental Breakfast (Pre-Function C)
7:50am	Welcome Back (Jim Marsteller)
8:00am	Keynote Address: Matthew Rosenquist, Cyber Security Strategist, Intel "Strategic Leadership for Managing Evolving Cybersecurity Risks"
9:00am	Panel: Threat Profile for NSF Large Facilities and Cyberinfrastructure *Panelists*: Amy Butler (GWU), Jeremy Epstein (NSF), David Halstead (NRAO), Matthew Rosenquist (Intel) *Moderator*: David Raymond (Army Cyber Institute, U.S. Military Academy)
10:00am	*Coffee Break*
10:30am	Panel: The Role of Information Sharing in Large Facility Security *Panelists*: Joel Cutcher-Gershenfeld (UIUC), Jim Marsteller (PSC), Kim Milford (REN-ISAC), Abe Singer (LIGO) *Moderator*: Greg Bell (ESNet)
11:30am	Open Discussion / Summary of Summit Findings (Von Welch, Craig Jackson, Jim Marsteller)
12:00pm	Adjourn

Appendix B

Biographies for Speakers, Program Committee, and Organizers

2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

*

Bios for Speakers, Authors, Program Committee Members, Organizers, and Student Awardees

---

in alphabetical order by surname

Jared Allar is a Pittsburgh Supercomputer Center information security analyst. His background covers many aspects of information security including vulnerability discovery, vulnerability coordination, security evaluations of information systems, and incident response. He has done information security work in the fields of health insurance, banking, and academia.

*

Amy Apon is Chair of the Computer Science Division in the School of Computing at Clemson University. She is the current Past Chair of the Coalition for Academic Scientific Computation (CASC), an organization of nearly 80 U.S. academic institutions who are leaders in computational and data-enabled science and engineering. Apon does research in high performance computing clusters and infrastructure for collaborative computing and is leading several initiatives to expand graduate education and research, including the CI SEEDS project, funded by NSF, that is increasing the number of domestic Ph.D. students in areas of data-enabled science at Clemson University. Dr. Apon holds a Ph.D. in Computer Science from Vanderbilt University and Masters and Bachelor's degrees from the University of Missouri - Columbia.

*

Fahad Arshad completed his Ph.D. from Purdue University in the Department of Electrical and Computer Engineering. His research interests focus on developing algorithms for software testing, error detection and failure diagnosis in distributed systems. He is particularly interested in data-driven analysis of computer systems. His work has appeared at top dependability conferences - DSN, ISSRE, ICAC, Middleware and SRDS, and he has been awarded grants to attend DSN, ICAC and ICNP. He has also been an active contributor to security research while working as a cybersecurity engineer at NEEScomm IT, Purdue University. He has recently joined a position as a systems engineer in industry.

*

Security Engineer **Justin Azoff** is responsible for implementing security plans; assisting other NCSA groups in hardening and protecting their systems; and developing, administering and utilizing NCSA's state-of-the-art cybersecurity monitoring infrastructure in support of the Center's objective of providing a highly reliable and functional computing environment. Working with other Security Engineers, Azoff identifies and investigates cybersecurity incidents across NCSA networks and systems and responds to these events, interdicting malicious behavior, mitigating security vulnerabilities, remediating compromised systems and adjusting cybersecurity controls as appropriate to ensure similar malicious behavior is prevented in the future. Azoff has been a Bro user since 2009 and became a Bro developer as part of his security engineer role when he joined NCSA in 2012.

*

Saurabh Bagchi is a Professor in the School of Electrical and Computer Engineering and the Department of Computer Science (by courtesy) at Purdue University in West Lafayette, Indiana. He is an ACM Distinguished Scientist (2013), a Senior Member of IEEE (2007) and of ACM (2009), a Distinguished Speaker for ACM (2012), an IMPACT Faculty Fellow at Purdue (2013-14), and an Assistant Director of the CERIAS security center at Purdue. He is the Cybersecurity Lead for the NSF Center at Purdue called NEEScomm. His work on fault tolerance in distributed systems has been rewarded with recognition of best papers or runner-up awards at several conferences (Sensys 2011, Supercomputing 2012, 2009, SecureComm 2008, etc.) and through the Seed for Success award at Purdue University twice. He is proudest of the 11 PhD students who have graduated from his research group and have gone on to wonderful careers in industry or academia.

*

Steve Barnet has specialized in supporting scientific and academic computing for nearly 20 years. During that time, he has worked in multiple domains including storage, networking, high-throughput computing, and security. He handled his first incident in 1995, a compromised Solaris system providing several important infrastructure services.

Steve is currently works for the IceCube project, a kilometer scale neutrino detector located at the geographic South Pole. He began collaborating with CTSC in 2013 to develop a Cybersecurity plan for the IceCube facility.

*

William Barnett oversees the life sciences research IT practice at Indiana University, both for basic research and for health care research including the IU School of Medicine, where he is an adjunct associate professor in Medical and Molecular Genetics. He is the Co-Director of Translational Informatics at the Indiana Clinical and Translational Sciences Institute (CTSI). Bill is the Director of the National Center for Genome Analysis Support (ncgas.org), which provides bioinformatics and computational support for genomics research. He also oversees the Grid Operations Center for the Open Science Grid. As an Associate Director of the Center for Applied Cybersecurity Research, Bill led the alignment of IU computing and data management systems with HIPAA. He is on the Steering Committee for the Association for American Medical Colleges (AAMC) Group on Information Resources (GIR) and faculty of the AAMC GIR Leadership Institute.

*

Jim Basney is a senior research scientist at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign. Jim leads the CILogon project (www.cilogon.org), which enables federated authentication to cyberinfrastructure. Jim is also the security technical lead for XSEDE (www.xsede.org) Software Development and Integration (SD&I), and Jim is the identity management lead for the Software Assurance Marketplace (SWAMP). Jim maintains the MyProxy credential management software, an “exemplar of success in cyberinfrastructure software sustainability” according to the report from the NSF workshop on CyberInfrastructure Software Sustainability and Reusability (http://hdl.handle.net/2022/6701). Jim is an active participant in The Americas Grid Policy Management Authority and the InCommon Technical Advisory Committee. Jim received his PhD in computer sciences from the University of Wisconsin-Madison where he worked as a graduate research assistant on the Condor project.

*

Tony Baylis of Lawrence Livermore National Laboratory is the Laboratory's Director for the Office of Strategic Diversity and Inclusion Programs. In this position, he is the senior management advocate for diversity and inclusion for the Laboratory. The Office of Strategic Diversity and Inclusion Programs partners with senior management to develop strategies, initiatives, programs, and activities that promote the creation of a diverse and inclusive workforce and work environment. Tony serves as the Laboratory's EEO, AA and Diversity compliance officer as well. In conjunction with these tasks, Tony is responsible for overseeing the laboratory's interactions and successful execution in building, partnering and collaborating with governmental, educational, industrial, community interests and other stakeholders. LLNL has had a long history in working with Minority Serving Institutions, specifically relationships with American Indian Institutions, Hispanic Institutions and Historically Black College and Universities. He represents the Laboratory on the subjects of Diversity and Inclusion, STEM, Outreach Efforts, and Student Programs.

Tony's career represents 26 years of administrative, project, program, technical and organizational management. He has worked in a scientific and technical environment for over 20 years and has worked as an consultant in industry as well. Tony has extensive experience networking with a broad range of academic, industry, government and non-profit organizations that has educated him and helped him in his career. He serves on a number of conference program committees and advisory boards that promote STEM and diversity in science and technical careers. He has been an NSF reviewer and PI/Co-Pi for the Broadening Participation in Computing Program. Tony is also an ACM and ACM SIGGRAPH member, and serves as the Treasurer for ACM SIGGRAPH. He is a graduate of the University of Illinois.

*

Gregory Bell is director of the Scientific Networking Division at Lawrence Berkeley National Laboratory (Berkeley Lab), and director of the Energy Sciences Network (ESnet), the U.S. Department of Energy's high-performance national network - one of the oldest and fastest computer networks in the world. Previously, Bell served as Chief Technology Architect in Berkeley Lab's IT Division, and prior to that he worked as a network engineer at Berkeley Lab. His professional interests include advanced networking technologies, cyber-security models for open science, and data-intensive discovery. Bell earned an AB from Harvard College (English), and a PhD from UC Berkeley, where he wrote an interdisciplinary dissertation on the cultural history of conspiracy belief. Bell has also managed a non-profit agency serving political refugees, and served as an analyst for Amnesty International. He lives in Berkeley with his wife Chalon.

*

Jasmine Bowers is a Lawrence Livermore National Laboratory (LLNL) Cyber Defenders summer scholar and a second year M.S. candidate at North Carolina Agricultural and Technical State University (NCAT). She holds two B.S. degrees in mathematics and computer science from Fort Valley State University (FVSU). She will graduate with an M.S. in computer science in May of 2015. This summer, she worked with two LLNL computer scientists on an iOS mobile application for simplifying and securing password-based authentication. This application will alleviate the burden and security risks associated with requiring a user to keep track of several complex passwords for various websites.

Jasmine is also a research assistant in the NCAT Center for Advanced Studies in Identity Science (CASIS) group, directed by the NCAT computer science department chair. She is studying under the information assurance masters track and her research topic is author identification. At the 2013

Richard Tapia Diversity in Computing Conference, she presented “Android vs. iPhone: What’s Your Personality”, an undergraduate project that analyzed the correlation of users and phone operating system preferences. As an undergraduate, she worked with the Department of Defense as a civilian computer science cooperative education student. At FVSU, she served as treasurer of the ACM chapter, two-term president of the Eta chapter of Delta Sigma Theta Sorority, Inc., personal assistant to the Director of Leadership and Character Development, assistant to the Director of the Cooperative Developmental Energy Program, and teacher assistant. In addition, she served as a mathematics tutor at the local middle school and FVSU tutoring lab.

In her spare time, she provides budget coaching and workshops. She recently presented financial workshops at the FVSU annual iLead Leadership Conference.

*

Amy Butler has 15 years of IT experience, with ten in the creation and deployment of solutions protecting information assets and ensuring confidentiality, availability and integrity of a large enterprise environment. Ms. Butler is currently the AVP, Information Security and Compliance at The George Washington University. Previously, she held positions at The Coca Cola Company, Secore, Inc and Peking University. She is a Lecturer at GW's Graduate School of Business as well as its Graduate School of Computer Science. Ms. Butler holds an MBA from The George Washington University and specializes in the development and implementation of enterprise security strategies.

*

Randal Butler serves as Deputy Director for CTSC and leads CTSC EOT activities. He is director of the NCSA’s Cybersecurity Directorate, Chief Security Officer for NCSA and the Security Officer for the NSF XSEDE project. Previously, he led security efforts for the National Computational Science Alliance and was the NCSA PI of the NSF National Middleware Initiative GRIDS Center.

*

Michael Corn is the Deputy CIO and CISO for Brandeis University. His areas of interest include privacy, identity management, and cloud services. He has been an active speaker and author on security and privacy and has participated in numerous Educause and Internet2 initiatives. He is a member of the Internet2 Netplus Product Advisory Board and until recently was also a member of the Box.com and Splunk Product Advisory Boards, as well as the Kuali Ready Product Board.

Prior to joining Brandeis he was the CISO and Chief Privacy and Security Officer of the University of Illinois at Urbana-Champaign. He is a graduate of the University of Colorado at Boulder and the University of Illinois at Urbana-Champaign.

*

Joel Cutcher-Gershenfeld is a Professor and former Dean in the School of Labor and Employment Relations (LER) at the University of Illinois. He is also a Senior Research Scientist with the National Center for Super Computing Applications (NCSA) and holds a courtesy appointment in Industrial and Enterprise Systems Engineering (IESE) at the University of Illinois. Joel also serves as a visiting Professor in Work and Organizations at the University of Sydney, Australia.

He is an award-winning author who has co-authored or co-edited eleven books, including *Ford-UAW Pivots: Transforming Work and Relationships to Deliver Results* (MIT Press, 2015 forthcoming),

Multinational Human Resource Management and the Law (Edward Elgar, 2013), *Valuable Disconnects in Organizational Learning Systems* (Oxford University Press, 2005), *Lean Enterprise Value* (Palgrave, 2002), *Knowledge-Driven Work* (Oxford University Press, 1998), and *Strategic Negotiations* (Harvard Business School Press, 1994), and over eighty five articles on high performance work systems, transformation in labor-management relations, negotiations and conflict resolution, economic development, and engineering systems. His current research centers on stakeholder alignment in complex systems – a foundation for 21st Century institutions. Along with his co-inventors, he has a patent pending on a new visualization method designed to help see points of alignment and misalignment among stakeholders.

Joel was the 2009 President of the Labor and Employment Relations Association (LERA). Prior to coming to the University of Illinois, Joel served as a Senior Research Scientist and Executive Director of the Engineering Systems Learning Center, with a joint appointment in MIT's Sloan School of Management and MIT's Engineering Systems Division, as well as a Visiting Associate Professor at Babson College, and an Associate Professor at Michigan State University.

Joel has extensive experience leading large-scale systems change initiatives with public and private stakeholders in Australia, Bermuda, Canada, Denmark, England, Iceland, Italy, Jamaica, Mexico, New Zealand, Panama, Poland, Spain, South Africa, and the United States. He holds a Ph.D. in Industrial Relations from MIT and a B.S. in Industrial and Labor Relations from Cornell University.

*

Kyle Chard is a Senior Research Project Professional at the Computation Institute, a joint venture between The University of Chicago and Argonne National Laboratory. He received a PhD degree in Computer Science from Victoria University of Wellington in 2011. His research focuses on applying cloud-based techniques to large scale research data management as part of the Globus project. His research interests also include distributed meta-scheduling, Grid and Cloud computing, economic resource allocation and social computing.

*

Patrick Duda is a member of NCSA's Cybersecurity directorate and is currently assigned to work on CTSC. His responsibilities are to aid in the EOT efforts under the direction of Randy Butler. Most of this work is aimed at developing training programs to disseminate security information to NSF funded CI projects. Prior to joining NCSA Patrick worked with several software development companies. At NCSA he has worked on GRID computing and various other science projects.

*

Jeremy Epstein is lead program director for NSF's Secure and Trustworthy Cyberspace (SaTC) program, NSF's flagship cybersecurity research program. He is on loan to NSF from SRI International, where his research areas including voting system security and software assurance. Jeremy has spent 25 years in the security field as a researcher, product developer, consultant, and program manager. He is associate editor in chief of IEEE Security & Privacy Magazine, and founder of the ACSA Scholarships for Women Studying Information Security (SWSIS) program. He holds an MS from Purdue University in Computer Sciences, and is ABD from George Mason University.

*

Barbara Fossum is the Deputy Director for the George E. Brown, Jr. Network for Earthquake Engineering Simulations (NEES), at Purdue University in Lafayette, Indiana. In this capacity, Barbara directs the day-to-day operation and the development of cyberinfrastructure to support the $105 million NSF distributed network of 14 earthquake engineering research centers. Barbara comes to Purdue from the NSF where she was a Program Manager from 2001 to 2004, for the Information Technology Research initiative within the Office of Cyberinfrastructure Research. While currently devoting her time to Large Facility operations and management, she continues to be engaged in supercomputing activities and scientific visualization.

*

Ian Foster is Director of the Computation Institute, a joint institute of the University of Chicago and Argonne National Laboratory. He is also an Argonne Senior Scientist and Distinguished Fellow and the Arthur Holly Compton Distinguished Service Professor of Computer Science.

Ian received a BSc (Hons I) degree from the University of Canterbury, New Zealand, and a PhD from Imperial College, United Kingdom, both in computer science. His research deals with distributed, parallel, and data-intensive computing technologies, and innovative applications of those technologies to scientific problems in such domains as climate change and biomedicine. Methods and software developed under his leadership underpin many large national and international cyberinfrastructures.

Dr. Foster is a fellow of the American Association for the Advancement of Science, the Association for Computing Machinery, and the British Computer Society. His awards include the Global Information Infrastructure (GII) Next Generation award, the British Computer Society's Lovelace Medal, R&D Magazine's Innovator of the Year, and an honorary doctorate from the University of Canterbury, New Zealand. He was a co-founder of Univa UD, Inc., a company established to deliver grid and cloud computing solutions.

*

Kelly Gaither, Director of Visualization, Texas Advanced Computing Center

*

Bret Goodrich is the Software Manager for the Daniel K. Inouye Solar Telescope (DKIST). He is responsible for the High Level Software and Controls group, including the software development for the telescope, instrument cameras, data handling, observatory control, and architectural framework. In addition, he is responsible for the computing assets of the facility, including their specification, performance, security, and operational procedures. He has worked for over 30 years on telescope software and information technology at Kitt Peak National Observatory, Gemini Observatory, and the National Solar Observatory. He has participated in the design and development of numerous telescope projects and is an active member of the telescope software community.

*

Christopher Gullo is a growing software developer contributing to LLNL's research initiatives in cyber security situational awareness. During the summer of 2014, Chris will expand his computer science skill set working on existing projects with the Lawrence Livermore National Laboratory Cyber Defenders Program that will be influential for years to come.

Chris enjoys spending time outdoors (hiking, biking, running, camping, sports), traveling, flying, following technological innovations, and more. Chris is a rising senior studying Computer Science at Rochester Institute of Technology where he has been honored to lead several class team projects. He participates in Air Force ROTC and plans to both graduate and commission in the United States Air Force in May 2016. His goal is to be a Pilot or Cyberspace Officer in the Air Force.

*

David M. Halstead, Head of IT, CIO, National Radio Astronomy Observatory

Highest Degree: Ph.D., Computational Quantum Surface Chemistry, University of Liverpool, 1990

Experience: 20+ years of experience with HPC systems and high speed storage/network solutions for both research and industry. Extensive knowledge of communications technologies and data intensive systems (genomics, chemistry and astronomy). Operations support for large infrastructure initiatives.

Bio Highlights
1985-92 Doctoral and Post-doctoral research with 10+ peer reviewed journal articles exploring molecular interactions with metal surfaces.
1992 Moved into HPC research at the Scalable Computing Laboratory of Ames Lab, DOE, implementing commodity parallel processing cluster solutions to benefit research in surface science, chemistry, physics and biology. ESNet representative for Lab.
2002 Hired by Celera Genomics to drive the Strategic Platform Initiative; transitioning away from the ~$20M leased computer systems used to sequence the human genome, to scalable HPC systems supporting proteomics and therapeutics research.
2004 Moved to into corporate Applera to manage a team of ~12 Communications Services staff at ~6 locations around the US coordinating special IT projects including multiple $100M+ Mergers & Acquisitions events.
2008 hired as NRAO CIO in support of North American ALMA and NRAO-wide IT and Science Computing services. Now manage 26 staff under Business Office and Data Management & Software at 3 main US locations, supporting 15 sites and interfacing with the Joint ALMA Observatory in Chile.

Community Service
Organizing Committee for Super Computing Conference series: SC94, SC99, SC05, SC10; SC13; SC14. Reviewer for multiple NSF HPC programs/awards. Founding member of new ACM's SIG HPC for Education.

*

Ardoth Hassler is Associate Vice President of University Information Services & Executive Director, Office of Assessment and Decision Support at Georgetown University. Her work focuses on policy, planning and research, including being the PI for an NSF CC-NIE award. She also supports institutional research, business intelligence, data warehousing and reporting. She was on loan to the National Science Foundation 2007-2011 where she served as Senior Information Technology Advisor in the Office of the Chief Information Officer in the NSF Office of Information and Resource Management, Division of Information Systems. Her activities included work related to cybersecurity best practices for large research facilities, working on technology policy for the Foundation and large research facilities, assisting NSF in joining the InCommon Federation and introducing concepts of single-sign-on logon to Research.gov, leading the SSN Be Gone project to remove SSNs from FastLane and other systems where there was no business need, working on NSF's Got Green initiative, etc.

She has prior experience serving on the program committees of the NSF Cybersecurity Summit, EDUCAUSE Annual Conferences, etc.. She has a BS in Math (CS minor) from Oklahoma State University and an MS in Biostatistics from the University of Oklahoma.

*

Elisa Heymann is an Associate Professor in the Computer Architecture and Operating Systems Department at the Autonomous University of Barcelona. She co-directs the MIST software vulnerability assessment project in collaboration with her colleagues at the University of Wisconsin.

She is also in charge of the Grid security group at the UAB, and participates in two major Grid European Projects: EGI-InSPIRE and European Middleware Initiative (EMI). Heymann's research interests include security and resource management for Grid and Cloud environments. Her research is supported by the Spanish government, the European Commission, and NATO.

Heymann received her M.S. and Ph.D. degrees in Computer Science from the Autonomous University of Barcelona (Spain) in 1995 and 2001 respectively.

*

Craig Jackson is Senior Policy Analyst at Indiana University's Center for Applied Cybersecurity Research (CACR), where his research interests include risk management, security, and identity management. He serves as the project manager for the Center for Trustworthy Scientific Cyberinfrastructure (CTSC); is part of the security team for the DHS-funded Software Assurance Marketplace (SWAMP); and is part of the DOE-funded XSIM (Extreme Scale Identity Management) project. He is a graduate of the IU Maurer School of Law (J.D.'10) and IU School of Education (M.S.'04). As a member of the Indiana bar, Mr. Jackson has litigated in federal and state court, primarily representing government and corporate clients in constitutional and tort claims. His research, design, and project management background includes work at IU School of Education's Center for Research on Learning and Technology and Washington University in St. Louis School of Medicine. He is a member of Phi Beta Kappa, and was a Lien Honorary Scholar at Washington University in St. Louis. He is married with 2 kids and 2 dogs. In his free time, he crashes BMX bikes and writes indie movie scores.

*

Dr. **Clifford A. Jacobs** worked for the National Science Foundation (NSF) for 30 years and for 25 years of that time provided oversight to the National Center for Atmospheric Research (NCAR) and its managing organization University Corporation for Atmospheric Research (UCAR). His oversight responsibilities cover a wide range of topics, such as acquisition of supercomputers, the development of world-class climate and weather models, the initiation and maturation of cyberinfrastructure necessary to delivery environmental data observations and products through the Unidata program, coordinated collaborative activities among Federal agencies, participation in the working group to develop NSF clarification of its data policy, the development of requirement for a data management plan, and chaired an internal group of cyberinfrastructure for NSF-sponsored large facilities.

Dr. Jacobs has represented the geosciences in a variety of NSF studies and initiatives related to high performance computing and information technology, observing facilities, and best practices in the operation and management of facilities. In addition, he assisted with the oversight, planning and execution of several complex agency activities, including the operation and management of major

facilities and the EarthCube endeavor. Dr. Jacobs co-chairs an internal Directorate working group on Geoinformatics and data and serves a member of GEO facilities working group. While serving in the Division of Advanced Cyberinfrastructure, he continue his efforts to engage the NSF staff and the community in a dialog about cybersecurity for NSF-sponsored large facilities.

Currently, Dr. Jacobs is consulting through Clifford A. Jacobs Consulting, LLC.

*

William T.C. Kramer is Director and Principle Investigator of the Blue Waters Project and is the Director of the UIUC/NCSA @Scale Program office. Bill is responsible for leading all aspects of the Blue Waters project, a National Science Foundation-funded project at NCSA. Blue Waters the most powerful general purpose computational and data analytics available to open science, system available. It is one of the most powerful resources for the nation's researchers. It is the only public Top-5 systems in the world that chose not to list on the Top-500 list.

Previously Bill was the general manager of the NERSC at Lawrence Berkeley National Laboratory (LBNL) was responsible for all aspects of operations and customer service for NASA's Numerical Aerodynamic Simulator (NAS) supercomputer center. He also served as the CSO in those organizaions. Blue Waters is the 20th supercomputer Kramer deployed and/or manages, deployed and managed large clusters of workstations, five extremely large data repositories, some of the world's most intense networks. He has also been involved with the design, creation and commissioning of six "best of class" HPC facilities.

He holds a BS and MS in computer science from Purdue University, an ME in electrical engineering from the University of Delaware, a PhD in computer science at UC Berkeley.

Kramer's research interests include large-scale system performance evaluation, systems and resource management, job scheduling, fault detection and resiliency, and cyber protection. Kramer has taught classes and tutorials on large scale system management, computer architectures, cyber-protection and visualization.

*

Lee Liming is a Technical Communications Manager at the Computation Institute, a joint venture between The University of Chicago and Argonne National Laboratory. He has spent fourteen years working with scientists from many fields of study to build computing systems capable of supporting their ever-growing data and computing needs. Past collaborations have included civil engineers, space scientists and astronomers, climate scientists, high-energy physicists, energy scientists, cosmologists, social scientists and librarians, neuroscientists, cancer researchers, and, of course, computer scientists. Prior to working at the University of Chicago and Argonne, Lee was a Sr. Product Manager and Principal Engineer at ProQuest Information and Learning and an information technology manager at the University of Michigan. Lee received a B.S.E degree in Computer Engineering at the University of Michigan.

*

Paul Lordier a Senior at California State University Sacramento where he is pursuing a B.S. in Computer Science with a concentration in information assurance / cyber security, and a minor in Geographic Information Systems. His expected graduation date is Spring 2015. Paul is a recipient of the Cybercorps Scholarship for Service, a program sponsored by the National Science Foundation that

funds students pursing cyber security related programs with a goal of placing them in government cyber security jobs. Paul recently completed a summer internship in the Cyber Defenders program at Lawrence Livermore National Laboratory.

*

As the Information Security Officer of the Pittsburgh Supercomputing Center, **James A. Marsteller, Jr.** (CISSP) is responsible for ensuring the availability and integrity of the PSC's high performance computing assets. Jim has over 12 years experience in the information security field and greater than 17 years of professional experience in the field of technology. Prior to working at PSC, he was a program manager for the Carnegie Mellon Research Institute that provided information security consulting services for government agencies and Fortune 500 companies. Jim leads the XSEDE Incident Response team and is XSEDE's security officer. He is a Co-PI for the Center for Trustworthy Scientific Cyberinfrastructure (CTSC). Jim chaired the program committee for the three most recent past summits, 2008, 2009, and 2013.

*

Celeste Matarrazzo is a data science expert with more than 27 years of service to the Lawrence Livermore National Laboratory's (LLNL) Computation Directorate. Celeste is presently the Associate Program Leader for Network Exploitation within the Global Security Principal Associate Directorate and currently the Principal Investigator for a large LLNL funded strategic initiative research project in cyber security situational awareness called Continuous Network Cartography following on from a successful research effort she led from November 2008 through September 2011. She is also the program manager for LLNL's Cyber Defenders Summer Intern Program. Celeste was previously a division leader who provided oversight and technical leadership for computer scientists and technicians addressing global security issues. Celeste also was the project leader for the Advanced Simulation and Computing Program's Scientific Data Management effort. Celeste has a B.S. in Mathematics and Computer Science from Adelphi University and pursued her graduate studies at the University of Wisconsin- Madison.
*

Michael McLennan is a Senior Research Scientist at Purdue University and Director of the HUBzero Platform for Scientific Collaboration. He created the Rappture toolkit as part of that platform. He has more than 20 years of software development experience in both academic and corporate environments, with an emphasis on computer-aided design tools and user interface design.

Dr. McLennan received a Ph.D. in 1990 from Purdue University for his dissertation on dissipative quantum mechanical electron transport in semiconductor heterostructure devices. He became a Tcl enthusiast when he joined Bell Labs in 1992 to work on tools for semiconductor device and process simulation. He is co-author of "Effective Tcl/Tk Programming" (published by Addison-Wesley) and "Tcl/Tk Tools" (published by O'Reilly and Associates). He also developed [incr Tcl], an object-oriented extension of Tcl, which is now used by thousands of developers worldwide, on projects ranging from the TiVo digital video recorder to the Mars Pathfinder.

*

Pascal Meunier is the head of security and operations at HUBzero. He has 15 years of experience working in computer security, starting at Purdue University CERIAS and continuing at HUBzero. He has been an editor for MITRE's CVE since its early days and contributed to related efforts. He created and taught secure programming classes at Purdue and maintains an active CISSP certification.

*

Kim Milford became Executive Director of REN-ISAC in April 2014. As Executive Director, Ms. Milford works with members, partners, sponsors and advisory committees to direct strategic objectives in support of member institutions, providing services and information that allows them to better defend local technical environments while overseeing administration and operations. She joined Indiana University in June 2007 and served in several different roles leading IT, information policy, and university privacy initiatives during her tenure. Most recently, Ms. Milford served as Chief Privacy Officer, coordinating privacy-related efforts, serving on IU's Assurance Council, chairing the Committee of Data Stewards, and directing the work of the University Information Policy Office, which includes IU's IT incident response team. Prior to joining Indiana University, Ms. Milford served as the Information Security Officer at the University of Rochester, where she successfully incorporated security plans and operations into strategic IT initiatives. In that role, she developed and led an information security program that included disaster recovery planning, identity management, incident response and user awareness. As Information Security Manager at UW-Madison's Department of Information Technology from 1998 - 2004, she assisted in the establishment of the university's information security department and co-led in the development of an annual security conference. Ms. Milford has provided presentations on information technology at various national conferences and seminars and participated in the authorship of articles and courseware. Ms. Milford has a B.S. in Accounting from Saint Louis University in St. Louis, Missouri and a J.D. from John Marshall Law School in Chicago, Illinois.

*

Barton Miller is Professor of Computer Sciences at the University of Wisconsin. He is Chief Scientist for the DHS Software Assurance Marketplace research facility. He co-directs the MIST software vulnerability assessment project in collaboration with his colleagues at the Autonomous University of Barcelona. He also leads Paradyn Parallel Performance Tool project, which is investigating performance and instrumentation technologies for parallel and distributed applications and systems. His research interests include systems security, binary
and malicious code analysis and instrumentation extreme scale systems, parallel and distributed program measurement and debugging, and mobile computing. Miller's research is supported by the U.S. Department of Homeland Security, U.S. Department of Energy, National Science Foundation, NATO, and various corporations.

In 1988, Miller founded the field of Fuzz random software testing, which is the foundation of many security and software engineering disciplines. In 1992, Miller (working with his then-student, Prof. Jeffrey Hollingsworth, founded the field of dynamic binary code instrumentation and coined the term "dynamic instrumentation". Dynamic instrumentation forms the basis for his current efforts in malware analysis and instrumentation.

Miller was the chair of the IDA Center for Computing Sciences Program Review Committee, a member of the Los Alamos National Laboratory Computing, Communications and Networking Division Review Committee, and has been on the U.S. Secret Service
Electronic Crimes Task Force (Chicago Area), the Advisory Committee for Tuskegee University's High Performance Computing Program, and the Advisory Board for the International Summer Institute on Parallel Computer Architectures, Languages, and
Algorithms in Prague. Miller is an active participant in the European Union APART performance tools initiative.

Miller received his Ph.D. degree in Computer Science from the University of California, Berkeley in 1984. He is a Fellow of the ACM.

*

Pascal Meunier is the head of security and operations at HUBzero. He has 15 years of experience working in computer security, starting at Purdue University CERIAS and continuing at HUBzero. He has been an editor for MITRE's CVE since its early days and contributed to related efforts. He created and taught secure programming classes at Purdue and maintains an active CISSP certification.

*

Gaspar Modelo-Howard is a Senior Researcher at Narus, a big data analytics for cybersecurity company and wholly owned subsidiary of the Boeing Company. He is also the Director for the ARGUS Information Security and Networking Lab at the Technological University of Panama. Gaspar has worked for over 14 years as a cyber-security consultant and engineer and as a college professor. He has a PhD in Computer Engineering from Purdue University and a MSc in Information Security from Royal Holloway, University of London. His current research interests lie at the intersection between machine learning and system security, particularly in the areas of malware detection, signature generation and web security. Gaspar is a Member of USENIX and ACM, and a Senior Member of IEEE.

*

William "Clay" Moody is a Computer Science Ph.D. Candidate at Clemson University, Clemson, SC. Clay is an active duty U.S. Army Major and will join the faculty of the Electrical Engineering and Computer Science Department at the United States Military Academy at West Point upon the completion of his doctoral studies. His research is focused on designing, modeling, and building applications that introduce the military concept of maneuver allowing parallel and distributed systems to be provisioned, optimized and secured. Clay is a founding member of the United States Cyber Command at Fort Meade, MD and a former Cyber Battle Captain in the Joint Operations Center. He holds a M.S. in Computer Networking from North Carolina State University.

*

Anita Nikolich is Program Director for Cybersecurity in the Division of Advanced Cyberinfrastructure at the National Science Foundation (NSF). Prior to her work at the NSF she served as the Executive Director of Infrastructure at the University of Chicago. Past assignments include Director of Global Data Networking at Aon and Director of Security for Worldcom. She has explored how information technology and secure networking can best support the creation and sharing of scientific knowledge in virtual, mobile and physical contexts. She holds a Master of Science from The University of Pennsylvania and a Bachelor of Arts from the University of Chicago.

*

Amy Northcutt was appointed Chief Information Officer of the National Science Foundation in January 2012. In this capacity, she is responsible for NSF's information technology investments, governance, policy, and planning. Prior to this appointment, Ms. Northcutt served as Deputy General Counsel of the Foundation from 2001 - 2012. Ms. Northcutt holds a J.D., *magna cum laude*, from Boston College Law School, an A.M.R.S. from the University of Chicago; and a B.A. from Smith College.

*

Rodney Petersen is the interim Executive Director of the Research and Education Community Security Collaborative, previously known as SecuriCORE. It is a new joint project between EDUCAUSE, Internet2, and Indiana University to establish a service organization to help improve cybersecurity at colleges and universities. Recently, he was the Managing Director of the EDUCAUSE Washington Office and a Senior Government Relations Officer. He also previously directed the EDUCAUSE Cybersecurity Initiative and was the lead staff liaison for the Higher Education Information Security Council. Prior to joining EDUCAUSE, he served as the Director of IT Policy and Planning in the Office of the Vice President and Chief Information Officer at the University of Maryland. He previously held the position of Campus Compliance Officer in the Office of the President at the University of Maryland where he mediated disputes and handled grievances under the Human Relations Code, including claims of discrimination or harassment that increasingly involved misuse of the Internet. He also completed one year of service as an Instructor in the Academy for Community Service for AmeriCorps National Civilian Community Corps where he taught alternative dispute resolution and facilitated service learning projects. He began his professional career in higher education as the Resident Student Life Director at Michigan State University. He is the co-editor of a book in the EDUCAUSE Leadership Strategy Series entitled "Computer and Network Security in Higher Education". He is also a founding member of the Association of College and University Policy Administrators and the author of "A Primer on Policy Development for Institutions of Higher Education" and "A Framework for IT Policy Development". He writes and speaks regularly on topics related to higher education cyber law and policy. He received his law degree from Wake Forest University. He also received a certificate as an Advanced Graduate Specialist in Education Policy, Planning, and Administration from the University of Maryland.

*

Irene M. Qualters is currently Division Director of Advanced Cyberinfrastructure (ACI) at the National Science Foundation (NSF). ACI is responsible for programs with a total annual budget in FY2013 of over $200 million. These programs support the acquisition, development, and provisioning of state-of-the-art cyberinfrastructure resources, tools, and services essential to the conduct of 21st century science and engineering research and education. ACI is also responsible for the NSF-wide vision, strategy, planning and coordination for research cyberinfrastructure. She joined NSF as a Program Director in December 2009, participating in multidisciplinary, interagency and international activities as well as overseeing several major computational projects within the division's portfolio, including the Blue Waters project at NCSA/UIUC and the Stampede project at TACC/UT at Austin. Irene has a Master's degree in Computer Science. Prior to beginning her NSF responsibilities, she had a distinguished 30-year career in industry, with executive leadership positions for research and development organizations within the technology sector. During her twenty years at Cray Research, in increasingly larger leadership roles, she participated in the development of the first commercially successful vectorizing compiler, the first multiprocessor version of Unix and Cray's landmark massively parallel computer, the T3E. Subsequently, for six years, as Vice President, she led the Research Information Systems for Merck Research Labs (MRL). She is expert in parallel computer system architectures and in a wide variety of software from scientific applications to compilers to file systems and operating systems.

*

Warren Raquel has been the Head of Operational Security and Incident Response and the National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign (UIUC) for the last year where he leads a Security Operations team that provides network security for NCSA and associated projects like Blue Waters and XSEDE. Prior to that he was a Security Analyst for the Office of Privacy and Information Assurance for UIUC where he did Incident Response and Digital Forensics. Warren has been a highly active member of the Higher Education security community for over a decade.

*

LTC **David Raymond** is an Armor Officer in the U.S. Army and is currently serving as Director of Research in the Army Cyber Center at West Point. He holds Bachelor's and Master's Degrees in Computer Science from the United States Military Academy and Duke University, and a Ph.D. in Computer Engineering from Virginia Tech. LTC Raymond has significant operational experience as an Armor officer, to include serving as a tank platoon leader during Operation Desert Storm and as a tank battalion executive officer during Operation Iraqi Freedom. He is a CISSP, Certified Ethical Hacker (CEH) and holds Global Information Assurance (GIAC) Certifications in Incident Handling, Intrusion Detection, Unix/Linux Security Administration, and Penetration Testing. LTC Raymond teaches senior-level computer networking and cyber security courses at West Point and conducts research on information assurance, network security, and online privacy.

*

Matthew Rosenquist joined Intel Corp in 1996 and benefits from 20 years in the field of security. Mr. Rosenquist specializes in security strategy, measuring value, and developing cost effective capabilities which deliver the optimal level of security. Currently, a cybersecurity strategist for the Intel Security Group, he helped in the formation of this global organization which brings together security across hardware, firmware, software and services. Previously, he managed the security playbook for Intel's PC strategy planning group, encompassing all security features landing in the PC. Mr. Rosenquist built and managed Intel's first global 24x7 Security Operations Center, oversaw several internal platform security products and services, deployed the enterprise intrusion detection program, and was the first Incident Commander for Intel's worldwide IT emergency response team. He has conducted hundreds of security investigations leading to arrests and successful prosecutions in defense of corporate assets. He ran security for Intel's multi-billion dollar worldwide mergers and acquisitions activities and justified the security strategy protecting Intel's global manufacturing capability.

Mr. Rosenquist is active in the industry, speaks at conferences, consults with industry partners, and has published acclaimed white papers, blogs, videos and audio-casts on a wide range of information security topics. He is very passionate about security and information technology, his chosen career path, and strives to blend practical risk mitigation practices and information technology capabilities to achieve an optimal level of security.

Mr. Rosenquist is active in the industry, speaks at conferences, consults with industry partners, and has published acclaimed white papers, blogs, videos and audio-casts on a wide range of information security topics. He is very passionate about security and information technology, his chosen career path, and strives to blend practical risk mitigation practices and information technology capabilities to achieve an optimal level of security.

*

Stephen Schwab is a Senior Computer Scientist with the University of Southern California's Information Sciences Institute, where his research draws broadly from the systems, networking, computer architecture, and information security communities. He is a long-time contributor to the DETER Cyber Security testbed project, focusing on modeling of experimental phenomena and testbed architecture. He currently leads the DARPA SAFERLab project, focused on assessing technology for anonymous and non-blockable Internet communication through the definition and testbed realization of forward-looking motivating scenarios. He also leads the DARPA-sponsored Quasar Vetting project, investigating how to detect tampered (malicious) firmware pre-installed on devices within the commercial supply chain. In the larger community, Schwab has held from 2008 to the present time the role of Security Architect for NSF's Global Environment for Network Innovations (GENI) initiative, aimed at deploying national-scale research infrastructure for the networking and distributed systems communities.

*

Dr. **Phyllis Schneck** serves as the Deputy Under Secretary for Cybersecurity for the National Protection and Programs Directorate within the Department of Homeland Security (DHS). Dr. Schneck is the chief cybersecurity official for DHS and supports its mission of strengthening the security and resilience of the nation's critical infrastructure.

Dr. Schneck came to DHS from McAfee, Inc., where she was Chief Technology Officer for Global Public Sector. Dr. Schneck served eight years as chairman of the National Board of Directors of the FBI's InfraGard program and founding president of InfraGard Atlanta.

Before joining McAfee, Dr. Schneck was Vice President of Research Integration for Secure Computing. She also worked as the Vice President of Enterprise Services for eCommSecurity; served as Vice President of Corporate Strategy for SecureWorks, Inc.; and, was Founder and Chief Executive Officer of Avalon Communications.

Dr. Schneck earned her Ph.D. in Computer Science from Georgia Tech.

*

Mark Servilla is Lead Scientist, Network Information System at LTER Network Office (LNO). At LNO, Mr. Servilla's primary responsibility is the implementation of the LTER Network Information System—a system of standards and applications that support the interoperability of distributed LTER research sites, thus enabling synthetic science at the Network level and beyond. To achieve a successful Network Information System, he will rely on his skills as a computer scientist to use the latest computing technologies for maximum effectiveness within the NIS, while utilizing his experience as an earth scientist to better serve the needs and understand the requirements of LTER, associated scientists, and the field of Ecology in general. Prior to his current position at LNO, Mark's most recent role in the private sector at Photon Research Associates (PRA), Inc. was as architect of a web-based application (GeoServer TM) that provided the discovery, management, and exploitation of geospatial data, including Earth observation imagery and GIS vector objects. Mark holds graduate degrees in Earth and Planetary Sciences (Volcanology) and Computer Science, both from the University of New Mexico.

*

Anurag Shankar oversees HIPAA and other regulatory compliance activities at the University Information Technology Services (UITS) at Indiana University (IU). He spearheaded the technical

effort that led to the HIPAA alignment of UITS systems in 2008. He is a computational astrophysicist by training and has a Ph.D. in Astronomy from the University of Illinois at Urbana-Champaign. After postdoctoral work in Astronomy at the University of Arizona and IU, he switched professions to IT in 1995. He started his IT career as a senior Unix systems programmer at Brown University and then moved back to IU in 1997. He has spent the past seventeen years with the Research Technologies division of UITS at IU, playing a variety of roles that include managing Unix support, massive data storage, and the national Teragrid project, and supporting the research mission of the IU School of Medicine. He has been responsible for building several of IU's large data storage environments, for establishing research computing services for IU's Indiana Genomics Initiative and other life sciences efforts, and for co-building an information infrastructure and technology solutions for the Indiana Clinical and Translational Sciences Institute (CTSI).

*

Nigel Sharp is the Program Director for the Large Synoptic Survey Telescope project, in the Division of Astronomical Sciences (AST) in the Directorate for Mathematical and Physical Sciences (MPS) of the National Science Foundation (NSF). He has just worked LSST through the process of federal funding for major projects. He has some additional programmatic responsibilities, too minor to mention. After three degrees in physics, mathematics, and astrophysics at the University of Cambridge (the real one), Nigel moved to Texas and had a varied career in astronomy theory and observation, including instrumentation and telescope management. His service work has included supercomputer access and numerical methods consulting, and systems management, networking and security at an NSF FFRDC. After all that, it made sense to join NSF and continue to work on service to the community from the funding end of things. He has been involved in NSF's cyberinfrastructure initiatives, was part of the working group for interdisciplinary research, and helped to define and to implement NSF's data management plan requirement.

*

Abe Singer, Chief Security Officer, LIGO

*

Robin Sommer is a Senior Researcher at the International Computer Science Institute, Berkeley, and he is also a member of the cyber-security team at the Lawrence Berkeley National Laboratory. Robin Sommer's research focuses on network security and privacy, with a particular emphasis on high-performance network monitoring in operational settings. He is leading the development of the open-source Bro network security monitor, and he is a co-founder of Broala, a recent start-up offering professional Bro services to corporations and government.

*

Susan Sons serves as a Senior Systems Analyst at Indiana University's Center for Applied Cybersecurity Research. Susan comes to CACR and CTSC from a background in abuse management and web application development. She's a founding member of the Internet Civil Engineering Institute (ICEI), and is a co-author of The Definitive Guide to Drupal 7. Her interests include penetration testing, vulnerability management, security-conscious development practices, historical cryptography, and open source security tool sets.

*

Steven Tuecke is Deputy Director of the Computation Institute (CI) at The University of Chicago and Argonne National Laboratory, and co-leads the Globus project (www.globus.org) with Dr. Ian Foster. His focus is on the development of sustainable, cloud-based, software-as-a-service data management solutions to accelerate research. Prior to CI, Steven was co-founder, CEO and CTO of Univa Corporation from 2004-2008, providing open source and proprietary software for the high-performance computing and cloud computing markets. Before that, he spent 14 years at Argonne as research staff. Tuecke graduated summa cum laude with a B.A in mathematics and computer science from St. Olaf College.

*

Heidi L. Wachs is a Research Director on the Gartner for Technical Professionals Identity & Privacy Strategies Team where she publishes, presents, and advises clients on best practices for data privacy, information classification, and identity and access governance.

Prior to joining Gartner, Heidi was the Chief Privacy Officer & Director of IT Policy for Georgetown University. She combined her professional background in technology policy and public relations to develop and implement sound policy to preserve the privacy and integrity of those who use Georgetown University information systems and their data. Before Georgetown, Heidi worked in government relations and advocacy focusing on intellectual property and technology policy issues. Earlier in her career she worked in public relations representing technology-focused clients.

Heidi graduated cum laude with a BA in Journalism from Lehigh University and earned her JD with a concentration in Intellectual Property from the Benjamin N. Cardozo School of Law. She is a Certified Information Privacy Professional, and a member of the District of Columbia Bar and the United States Supreme Court Bar.

*

Von Welch is the director of Indiana University's Center for Applied Cybersecurity Research (CACR) and PI for the Center for Trustworthy Scientific Cyberinfrastructure, a project dedicated to helping NSF science projects with their cybersecurity needs. His expertise lies in applied research and practice of cybersecurity for distributed systems. Other roles include serving as CSO of the Software Assurance Market Place, a DHS-funded facility to foster software assurance and software assurance research, PI on a Department of Energy funded grant focused on identity management for extreme-scale scientific collaboration, and serving the Open Science Grid as an identity management expert. Previously he has worked with a range of high-visibility projects to provide cybersecurity to the broader scientific and engineering community, including TeraGrid, Open Science Grid, Ocean Observatory Infrastructure, and GENI. His work in software and standards includes authoring two IETF RFCS and the contributing to the creation of the well-known CILogon and MyProxy projects.

*

John Wroclawski is director of the University of Southern California's Information Sciences Institute's Internet and Networked Systems division, with responsibility for the strategic direction of this 40-member research organization. The division holds a historic and continuing role in the development of the Internet, and today maintains active programs of research in areas such as Internet protocols and architecture, network and distributed system security, sensing and sensor nets, network measurement and characterization, cyberphysical systems, and the Smart Grid. Wroclawski's immediate technical interests include the architecture, technology and protocols of large, decentralized communication systems such as the Internet, architectural aspects of

cyberphysical systems, and the core principles of self-organizing and self-structuring architectures. His contributions to the development of large-scale cyberinfrastructure for the networking and cybersecurity communities include an ongoing role as chief scientist of the DHS-supported DETER Cybersecurity testbed and service in the planning and early development stages of NSF's GENI project.

Appendix C
Call for Participation

Call for White Papers, Training, and Student Applications

2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

August 26 - 28 ✶ Westin Arlington Gateway ✶ Arlington, VA

http://trustedci.org/2014summit/

Theme: *Large Facility Cybersecurity Challenges and Responses*

It is our great pleasure to announce that the 2014 Cybersecurity Summit will take place Tuesday, August 26th through Thursday, August 28th, at the Westin Arlington Gateway near National Science Foundation Headquarters in Arlington, VA. On August 26th, the Summit will offer a full day of information security training. The second and third days will follow a workshop format designed to identify both the key cybersecurity challenges facing Large Facilities and the most effective responses to those challenges.

Spanning six years from 2004-2009 and reinstated in 2013, the annual NSF Cybersecurity Summit serves as a valuable part of the process of securing the NSF scientific cyberinfrastructure by providing the community a forum for education, sharing experiences, building relationships, and establishing best practices.

The NSF CI ecosystem presents an aggregate of complex, unique cybersecurity needs (e.g., scientific data and instruments, unique computational and storage resources, complex collaborations) as compared to other organizations and sectors. This community has a unique opportunity to develop information security

practices tailored to these needs, as well as break new ground on efficient, effective ways to protect information assets while supporting science. The Summit will bring together leaders in NSF CI and cybersecurity to continue the processes initiated in 2013: Building a trusting, collaborative community, and seriously addressing that community's core cybersecurity challenges.

Call for White Papers

Please submit brief white papers focused on NSF Large Facilities' unmet cybersecurity challenges, lessons learned, and/or significant successes. White papers (and presentations) may be in the form of position papers and/or narratives and may be one to five pages in length.

Criteria: All submitted white papers will be included in the 2014 summit report. The Program Committee will select the most relevant, reasoned, and broadly interesting for presentation during the Summit Plenary Session (Aug 27-28). A limited amount of funding is available to assist with travel for accepted submissions.

Extended submission deadline: July 12, 2014

Submit to: Craig Jackson, scjackso@indiana.edu

Word limit: 400 to 2000 words (~1-5 single spaced pages)

Notification of acceptance: July 16, 2014

Call for Information Security Training

Please submit brief abstracts from individuals or teams willing to present half and full-day training on August 26. Training may be targeted at technical and/or management audiences. Areas of interest include, but are not limited to, cybersecurity planning and programs, risk assessment and management,

regulatory compliance, identity and access management, networks security and monitoring, secure coding and software assurance, physical security in the context of information security, and information security of scientific and emerging technologies. The Program Committee will select the most community-relevant and broadly interesting training sessions for presentation during the first day of the summit (Aug 26).

Extended submission deadline: July 12, 2014

Submit to: Craig Jackson, scjackso@indiana.edu

Word Limit: 600 words

Notification of Acceptance: July 16, 2014

Call for Student Applications

Please submit a one page reference or cover letter and a student resume detailing how the student would benefit from attending the Summit. Recognizing that inclusivity and diverse participation is both a socially relevant outcome for NSF and a particular challenge in the cybersecurity community in general, the Program Committee will consider diversity when selecting successful applications.

Cover letters should address the student's interest in science and/or information security. Up to five successful student applicants will receive invitations to attend the Summit, and the opportunity for reimbursement of travel expenses. All submissions will be reviewed by the Program Committee and organizers.

Extended submission deadline: July 12, 2014

Submit to: Craig Jackson, scjackso@indiana.edu

Word limit for cover letters: 600 words

Notification of Acceptance: July 16, 2014

Program Committee and Organizers

Amy Apon, Chair of the Computer Science Division of the Clemson University School of Computing and former Director of the Arkansas High Performance Computing Center.

Anthony (Tony) Baylis, Director, Office of Strategic Diversity Programs at Lawrence Livermore National Laboratory.

Michael Corn, Deputy CIO and CISO for Brandeis University.

Barb Fossum, NEES deputy center director and former managing director of Purdue University's Cyber Center and Computer Research Institute.

Kelly Gaither, Director of Visualization, Texas Advanced Computing Center.

Ardoth Hassler, Associate Vice President of University Information Services & Executive Director, Office of Assessment and Decision Support at Georgetown University and former Senior Information Technology Advisor in the Office of the CIO in the NSF Office of Information and Resource Management, Division of Information Systems.

Craig Jackson (Organizer), Senior Policy Analyst, Center for Applied Cybersecurity Research, Indiana University, and Project Manager, Center for Trustworthy Scientific Cyberinfrastructure (CTSC)

James Marsteller (Organizer and Program Committee Chair), Information Security Officer, Pittsburgh Supercomputing Center, and Co-PI, Center for Trustworthy

Scientific Cyberinfrastructure (CTSC).

William "Clay" Moody, Computer Science PhD candidate and an active duty US Army Major stationed as an Army Fellow at Clemson University. Following his PhD studies, he has an appointment to the faculty of the Department of Electrical Engineering and Computer Science at West Point, the United States Military Academy.

Rodney Petersen, interim Executive Director of the Research and Education Community Security Collaborative (previously known as SecuriCORE) and former Managing Director of the EDUCAUSE Washington Office and a Senior Government Relations Officer.

Mark Servilla, Lead Scientist, Network Information System at LTER Network Office (LNO).

Von Welch, (Organizer), Deputy Director, Center for Applied Cybersecurity Research, Indiana University, and PI, Center for Trustworthy Scientific Cyberinfrastructure (CTSC).

Appendix D
White Papers Submitted in Response to the CFP

CYBERSECURITY OPERATIONS IN A MULTI-INSTITUTIONAL ACADEMIC SETTING: THE NEES STORY

July 7, 2014

White Paper for the 2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

Saurabh Bagchi[1,2], Fahad Ali Arshad[1,2], Gaspar Modelo-Howard[3,ξ]

(1) NEEScomm, Purdue University
(2) School of Electrical and Computer Engineering, Purdue University
(3) Narus, Inc.
(ξ) Gaspar worked at NEEScomm when he was involved in the work presented in this paper.

1. What is NEES and what role does Cybersecurity Play in it?

Earthquakes and tsunamis can be devastating not only to the infrastructure of a society, but also to families, the community, and people's sense of security. To reduce the impact of these events, fundamentally to save lives, the George E. Brown, Jr. Network for Earthquake Engineering Simulation (NEES) originated as a national, multi-user, research infrastructure to enable research and innovation in earthquake and tsunami loss reduction, create an educated workforce in hazard mitigation, and conduct broader outreach and lifelong learning activities [1].

In the ten years since officially opening its doors in 2004 to outside users, NEES has created a vibrant collaboratory consisting of unique laboratories and cyberinfrastructure with its collaboration platform, NEEShub, representing hundreds of millions of dollars of investment. The NEES collaboratory has served tens of thousands of users from over 210 different nations. In 2009, Purdue University took over from NEES Inc. as the manager of a network of 14 advanced laboratories [2] connected by a cyberinfrastructure. The NEES Community and Communications Center (NEEScomm) was established in West Lafayette, IN. The anticipated end of that project is May 2015, with the expectation of a future solicitation.

Participating universities in NEES include: Cornell University; Lehigh University; Oregon State University; Rensselaer Polytechnic Institute; University at Buffalo, State University of New York; University of California, Berkeley; University of California, Davis; University of California, Los Angeles; University of California, San Diego; University of California, Santa Barbara; University of Illinois, Urbana-Champaign; University of Minnesota; University of Nevada, Reno; and the University of Texas, Austin. Each of these university-based laboratories enabled researchers to explore a different aspect of the complex way that soils and structures behave in response to earthquakes and tsunamis. The laboratories were available not just to researchers at the universities where they are located, but to investigators throughout the United States who were awarded grants through NSF's annual NEES Research (NEESR) Program and other NSF programs.

In July 2010, NEEScomm released the first version of the **NEEShub**, the collaboration platform of NEES researchers [3] and based on HUBzero, a content management system built to support scientific activities [6]. Linking the NEES experimental facilities to each other, to NEEScomm, and to off-site users, this unique system of information technology resources has enabled researchers participating on-site or remotely to collect, view, process, and store data from NEES experiments at the NEEScurated central repository (Project Warehouse), to conduct numerical simulation studies, and to perform hybrid (combined experimental and numerical) testing involving one or more NEES equipment sites.

Role of cybersecurity in NEES

NEES has developed a comprehensive cybersecurity approach that includes best practice cybersecurity policies and mechanisms at NEEScomm and an annual security audit at each of the NEES sites [4]. The goal of our cybersecurity plan is to enable earthquake engineering and science to proceed unimpeded by security outages affecting the NEEShub, either the computational nodes or the data warehouse. This goal calls for careful tango between the needs of doing science quickly and keeping the IT assets secure from attacks. The cast of characters, i.e., the "stakeholders", who are most directly influenced by the cybersecurity practices are the earthquake engineers at the sites, the IT managers at the sites (who are responsible for maintaining the local machines and software on them), and the software development team at NEEScomm. The attack surface is large because we run an "open cyberinfrastructure", i.e., anyone with a valid email address can get access to (some) assets from NEEShub. Also, since the machines at the individual sites are used to access NEEShub, we are also tasked with auditing the security of these machines. The level of staffing is difficult to pinpoint unambiguously because the cybersecurity staff, who are within the broader NEEScomm IT organization, are liberally and routinely helped by people in the broader IT organization. The core cybersecurity staff is composed of a faculty member (Bagchi), a cybersecurity staff engineer (Howard), and a half-time graduate student (Arshad). Additionally, approximately 2 people from the broader group routinely perform cybersecurity activities such as installation of security patches and unlocking accounts.

The result of a well-documented planning process and then the implementation and deployment has been no reportable cybersecurity incident in the 5 years of existence. The cybersecurity activities have resulted in no major complaint from any of the stakeholders. Next, we describe the insights we have obtained by running this kind of a cyberinfrastructure for the past 5 years.

2. Insights in running cybersecurity operations in a multi-institutional university setting

There are several unique challenges that arose in the cybersecurity operations in our environment. The three top ones among them are:

1. Different universities have different cybersecurity policies and our policies had to "play nice" with them, including those at Purdue. For example, the password strength and password change policies differed widely.

   Solution: We sometimes enforced stronger rules and had them apply to those machines that are part of the NEES network. Thus, an "inner shell" of machines and other equipment (routers, PLC controllers, etc.) were created within the university IT resources and different policies applied on them, as they share the network with non-NEES equipment.

2. We were responsible for doing security audits of equipment which we did not own or have root access on. This applied to the IT equipment at the sites (for which we did an annual security scan, plus periodic low intensity scans) as well as some equipment at Purdue (which were administered by the HUBzero team).

   Solution: We negotiated elevated privileges for the purpose of the security scanning. Thus, we installed the tools for scanning to have the higher privileges but not give the root account to us. This highlighted the human element of doing cybersecurity in an environment such as ours. We had to tread carefully on the sensitivities of IT professionals at many organizations and believe we were successful in this and currently have the feeling of pulling together as part of a big team, rather than the very conceivable alternative – an adversarial relationship.

3. How to test for external threats as well as internal threats? The notion of perimeter security is ingrained in security products and the cybersecurity controls in place at NEEScomm as well as at the sites are no exception. The question was should we test for threats that can be launched externally (i.e., from outside the campus network) as well as internally.

   Solution: We decided in favor of doing the testing both ways primarily because with large campuses and some of the users likely to fall prey to security attacks (such as, phishing) it is likely that there will be insider threats. We therefore set up VPN connections for many of the sites' IT infrastructure and launched attacks through there, as well as launch attacks from pure external IP addresses. Philosophically, the cybersecurity controls at NEEScomm (at Purdue) had to be more stringent simply because we are in a position providing service to all the sites and we are developing software for wide use. For example, we employed static scanning of software as well as dynamic discovery of vulnerabilities for NEEScomm IT infrastructure.

3. Cybersecurity Research and Practice: Where do the Twains Meet?

All of us, the co-authors, are cybersecurity researchers and this project has made us keenly aware of the different worlds of research and applications of the research. Here we give a whirlwind tour of where the former aids the latter and where the twains are still far apart. In reading this, you should be cognizant of the fact that academic security research sometime neglects the issues of building robust tools, the user interface for the tool, and applicability of the research prototype to a variety of operational environments. It neglects these issues in favor of intellectual novelty in the work.

We made wide use of some tools that came out of academic research and are still the subject of vigorous academic research, an appreciable amount of which feeds back into the tools. Such tools are Bro (intrusion detection), OpenVAS/Nessus (general scanner for vulnerability assessment of systems/protocols), nmap (hosts and services discovery tool), nikto (scanner for vulnerability assessment

of web servers), Splunk (Log visualization and analysis tool for both security and operations), Tripwire (host-based intrusion detection system), Coverity (static code analysis tool used to test Java software at NEES), Zed Attack Proxy (ZAP) tool (a penetration testing tool for testing PHP code at NEES), fail2ban (a reactive service that bans when an automated bot or a hacker tries to login more than X number of times), iptables (to manage firewall rules on hosts). Some of these tools understandably needed significant effort in customization to our specific environment.

As academic researchers, we shone the mirror on ourselves and realized that academic security research could be benefited by exposure to specific operational challenges that an environment like NEES provided. One is security configuration management, which deals with the fact that security administration needs more help through tools. This is keeping in line with the increasing complexity of today's computer systems, which are more distributed and host an increasing number of applications. A problem faced here, and a possible reason why this area has seen relatively little work is the unavailability of security datasets from real-world systems. A second aspect is the need to concentrate on the false positives problem for signature-based detection systems. The primary reason that administrators improve a given signature over time, manually, is to reduce the false-positive rate. Research efforts should aim to do this more automatically even if it requires sacrificing accuracy considerably. Administrators do not want to waste time or even take actions on false alarms. Administrators get desensitized after using a tool with high number of false alarms, a phenomenon also well known in a hospital setting [5], and this may lead to missing the true alarms.

4. Conclusion

In this article we have described the cybersecurity activities in NEES, an NSF-funded center for earthquake engineering research, with research being done by scientists at 14 sites throughout the country. We have highlighted some of the unique challenges that arise in such a multi-institution setting and our pragmatic efforts at solving them. We concluded the article by considering the interplay between academic security research and practical security controls in a production environment such as ours.

5. References

[1] Purdue University, "George E. Brown, Jr Network for Earthquake Engineering Simulation," At: https://nees.org/
[2] Purdue University, "NEEShub laboratories," At: https://nees.org/sites-mainpage/laboratories
[3] NEEScomm IT, "NEEShub Release 6.0 – March 26, 2014," At: https://nees.org/resources/7666/download/Release_6_Communication_v2.pdf
[4] NEEScomm IT, "NEEShub – cybersecurity," At: https://nees.org/explore/security
[5] Alice Crites (Washington Post), "Too much noise from hospital alarms poses risk for patients," At: http://www.washingtonpost.com/sf/feature/wp/2013/07/07/too-much-noise-from-hospital-alarms-poses-risk-for-patients/

[6] Purdue University, “Hubzero, Platform for Scientific Collaboration,” At: http://hubzero.org

Case Study: XSEDE Leverages Globus Nexus for Identity and Group Management

Lee Liming, Steven Tuecke, Ian Foster, Kyle Chard
The University of Chicago, Computation Institute, Globus project

Abstract: *Research collaborations that share data and computing tools need a way to manage user identities, profiles, and groups. With members at multiple institutions, institutional services are unsuitable for these projects. Developing and maintaining custom solutions is challenging given the plethora of security protocols available and the need for scalable, robust, and highly available implementations. Globus Nexus[1] [2] is a professionally hosted service that offers these capabilities to research teams in a professional, reliable, cost-effective manner. XSEDE (Extreme Science and Engineering Discovery Environment)[3] is a data and computing services federation in the United States. XSEDE supports researchers and educators at U.S.-based institutions, including federal research labs and commercial organizations. In this paper, we present XSEDE's need to upgrade its services for identity and group management and its selection of Globus Nexus to provide this functionality.*

1. Globus Nexus - A platform for identity and group management

Globus Nexus is a professionally hosted service for user and group management tasks, with a particular focus on the needs of scientific communities. It provides features that are important to research applications such as identity provisioning; an "identity hub" that links identities from different systems to a single Globus identity; profile management; user-oriented group management; and branded web interfaces. Globus Nexus implements best practice approaches for each of these features. It implements delegated security protocols such as OAuth 2.0; provides sophisticated workflows for email validation and group membership modification; and implements sophisticated user-defined policies regarding permissible actions.

Globus Nexus is the identity management service for Globus.[4] It offers flexible application programming interfaces (APIs) that make it easy for end users and developers to access its functionality. Globus Nexus provides a Web browser interface, a command-line interface accessible via standard SSH clients, and a REST API.

---

[1] Kyle Chard, Mattias Lidman, Josh Bryan, Tom Howe, Brendan McCollam, Rachana Ananthakrishnan, Steven Tuecke, Ian Foster. "Globus Nexus: Research Identity, Profile, and Group Management as a Service." Submitted to The 10th IEEE International Conference on e-Science, 2014.

[2] Ananthakrishnan, R.; Bryan, J.; Chard, K.; Foster, I.; Howe, T.; Lidman, M.; Tuecke, S. "Globus Nexus: An identity, profile, and group management platform for science gateways and other collaborative science applications." IEEE International Conference on Cluster Computing (CLUSTER) , 23-27 Sept. 2013.

[3] https://www.xsede.org/.

[4] https://www.globus.org/.

Globus Nexus is offered as a hosted platform-as-a-service (PaaS) operated by a non-profit cost center at the University of Chicago on behalf of the research community. Built on commercial cloud services from Amazon (EC2, S3, Elastic load balancing, SMS) and widely used open source solutions (Cassandra, Elastic Search), Globus Nexus is a high-availability, professionally operated, best-of-breed service for research and academic identity management. While many of Globus's functions are available to academic institutions and non-profits at no cost, premium services are offered via a subscription model. This "freemium" approach allows the University of Chicago to maintain and grow the service and offer high-quality user support.

In the three years since deployment, Globus Nexus has been adopted by large research projects and manages more than 16,000 registered users linked to more than 6,500 external identities, and over 800 unique groups with more than 3,400 combined active memberships.

2. Identity management for XSEDE

XSEDE is a data and computing services federation that serves the United States research and academic community. Directly supporting more than 10,000 researchers and their associated teams, XSEDE is a critical element of the national science cyberinfrastructure.

As the XSEDE project has sought to broaden its service to the science and academic community, new ways of using XSEDE have appeared. These new activities include science gateways (in which a small group of developers use XSEDE to construct a custom application for a much larger set of researchers) and campus bridging (where campus IT administrators build "bridges" that make it easier for their local researchers to use XSEDE). These new ways of using XSEDE demand a more flexible set of identity management functions. Table 1 summarizes the identity management needs now known to XSEDE, with those that XSEDE did not initially support highlighted in boldface.

Table 1. XSEDE's identity management needs

User identity functions	Group functions	Authentication and authorization
● Create an identity ● Manage the identity profile ● Add and verify an email address ● **Link to a federated identity** ● Reset password ● **Disable one's own identity** ● Disable and enable an identity	● **Create a group** ● **Manage the group profile** ● **Manage group membership and assign member roles** ● **Invite members to a group** ● **Request membership in a group** ● **Disable and enable a group** ● **Delete a group** ● **List groups in which one is a member** ● **View a group profile**	● Obtain credentials of a particular type **(e.g., OAuth2)** from an XSEDE identity provider ● Use **[OAuth2]** credentials to authenticate with a relying party ● Delegate a credential to a relying party ● Use a credential to make an authorization decision ● Use a delegated credential to access another service on user's behalf ● **Use a group in an authorization policy**

Federated identities and OAuth 2.0

XSEDE identities are critical for understanding and managing user behavior on XSEDE. Of course, most researchers who use XSEDE already have digital identities established with other providers. These providers include academic institutions and departments, other computing centers, and commercial services. Rather than starting a new identity from scratch, a new XSEDE user should be able to link to existing identities from his or her home institution or a previous collaboration. Once the link is established, the XSEDE user can use the linked identity to obtain credentials for use throughout the XSEDE system. For example, when I first use XSEDE, I should essentially be able to say, "I am lliming@uchicago.edu and I can prove it, and I've never used XSEDE before. Sign me up." (If I've previously used XSEDE, I should be able to easily link this University of Chicago identity to my existing XSEDE identity.) I should then be able to use my lliming@uchicago.edu credentials (by successfully authenticating myself to University of Chicago) to sign in to XSEDE and use XSEDE services.

This style of identity linking is becoming more common among academic and commercial service providers. A specific mechanism that appears to be gaining momentum is OAuth 2.0.[5] For example, OAuth 2.0 is supported and used routinely by Google, Microsoft, and Facebook. Bringing XSEDE's users into the OAuth 2.0 community will prepare XSEDE to leverage commercially provided services: a longstanding goal.

User-defined groups

Researchers who use XSEDE increasingly need to share their work (particularly their data) with colleagues. These colleagues are often at different institutions, and they may or may not themselves be XSEDE users. The rules under which this sharing occurs are sometimes complicated, and it is the researchers who know best what rules should apply in a particular situation. Thus, XSEDE users need the ability to define and manage groups of their colleagues and use those groups for establishing sharing rules. Registering with XSEDE in order to become a group member should not be burdensome. (This is another reason why identity linking is important.)

As Table 1 makes clear, before this work began, XSEDE users could not create or manage their own groups. XSEDE provides each project leader (leader of a user team) with a group he/she can manage that controls access to the team's compute allocation, but that group is created and ultimately managed by XSEDE staff and cannot be used for other purposes.

Build vs. buy

XSEDE was faced with two pressing user needs: support for identity federation and support for user-defined groups. This situation presented a classic "build vs. buy" decision: should XSEDE's internal developers add these features to the existing systems, or should XSEDE find a way to leverage off-the-shelf tools instead? Sustainability was a key consideration: managing the ongoing cost of providing and operating services. XSEDE's primary funder, the National Science Foundation (NSF), is charged with encouraging innovation and scientific advancement. It is

[5] http://tools.ietf.org/html/rfc6749

discouraged from funding ongoing maintenance costs for existing systems. Thus, a plan to recover at least some of the ongoing maintenance cost of new services is an important project requirement.

When this decision arose in the past, commercial systems and services--designed for individual users or for single-enterprise use--could not offer the kinds of federation and flexibility needed by a multi-institutional scientific service provider. However, the rise of commercial "cloud" services--designed for shared use across many enterprises--has brought these requirements to the fore in the commercial sector, leading to standards such as OAuth 2.0.

Globus has enthusiastically adopted these commercial services as they have become available. Like XSEDE, Globus takes sustainability very seriously, and has developed a robust strategy for managing and recovering costs. In fact, it is operated as a non-profit cost center at the University of Chicago. Its subscription model for premium services allows it to recover the costs of commercially offered services and customized user support for the academic community.

Given the close similarity of requirements between what XSEDE needs and what Globus Nexus provides, the fact that Globus already leverages commercially offered services where possible, and the fact that Globus is already managing its costs and cost recovery in a way that suits NSF, XSEDE has chosen to acquire its new federated identity and group management features using Globus Nexus. The alternative would have been a costly development cycle and the need to develop a new cost recovery mechanism.

3. Integrating Nexus

XSEDE relies heavily on its XSEDE User Portal (XUP)[6] to provide a consistent, uniform interface for users who interact with XSEDE. XUP provides the user interface for commonly used functions, such as creating a new identity, logging in, and managing the user profile. To maintain this consistency, users are rarely redirected to other websites to perform activities. Instead, XUP will use the Globus APIs to perform operations on the user's behalf. In the few cases where truly new functionality is offered (e.g., using a federated ID to login, creating and managing groups), the user will be seamlessly redirected to a Globus Nexus web interface with a "skin" (user interface) that looks and behaves exactly like XUP, and will then be returned to XUP. In most cases, users will not realize that they visited another website. This integration pattern has already been used successfully by two other research systems to leverage Globus Nexus: Kbase[7] and BIRN[8].

The Globus and XSEDE user communities have developed independently for some time, so a username on Globus may refer to a different person than the same username on XSEDE. Similarly, individuals may already have different usernames on Globus and XSEDE. We resolve this issue by enabling Globus to maintain multiple user namespaces. By adding a namespace to a username (e.g., tuecke@globus.org vs. tuecke@xsede.org), Globus can distinguish the usernames that

[6] https://www.xsede.org/.
[7] https://gologin.kbase.us/
[8] https://access.birncommunity.org/SignIn

originated in Globus from the usernames that originated in XSEDE. We will maintain these separate namespaces indefinitely, using the linking feature to combine identities for people who have both XSEDE and Globus usernames. Thus, Lee Liming can have a single Globus identity linked to usernames lee@globus.org and lliming@xsede.org.

XSEDE will use Globus Nexus's group functionality in several ways, including the following:

- Users will create their own groups, for example to control access to shared data via XSEDE's Globus service or via XSEDE's global federated file system (GFFS).
- In specific instances, XSEDE's staff will create groups and pre-assign users to those groups, optionally giving some initial members the ability to invite additional members to these group while XSEDE staff retains overall administrative control of the group.
- XSEDE plans to maintain three administrative groups: one of all people known to XSEDE, another of all people who have ever had their identities "vetted" (verified in detail) by XSEDE accounts personnel, and a third for all people who currently have allocations to use portions of the XSEDE system.[9]

4. Immediate benefits and a path forward

By leveraging Globus's Nexus platform for identity and group management, XSEDE will be able to deliver critical new pieces of functionality to its users without committing to maintain and enhance its own implementation of those functions over the coming years.

- XSEDE's User Portal can provide an "alternate login" feature that allows a user to login using familiar credentials from their home institution or other research collaborations.
- XSEDE services can count on the availability of OAuth 2 tokens (increasingly common in the "cloud service" commercial space) for identifying and authenticating users.
- XSEDE users can form their own groups and use those groups to control access to shared resources.

XSEDE will use this functionality--including new features subsequently made available by Globus--while contributing only a portion of the overall maintenance cost. The full cost is spread across the rest of the Globus user community, at universities, research laboratories, and other funding agencies. In return, other Globus users gain the benefits of easier access to XSEDE services and access to Globus features that were added to support the XSEDE integration, such as multiple user namespaces.

Adding OAuth 2.0 support to its identity system opens XSEDE, its services, and its user community to a wider world of cloud services, both commercial and academic. It could well be that using the same authentication framework as Google, Microsoft, and Facebook will, in time, lead to further alliances for XSEDE, its users, and its partners.

[9] Vetting is a precondition for being permitted to use some portions of the system. Users are typically given a limited period of time to use these portions of the system.

Curbing Abusive Behavior of Science Gateways

Pascal Meunier, Michael McLennan
Purdue University, West Lafayette, Indiana

The abuse and misuse of shared system resources is not new, but it is often surprising to communities building science gateways, who naively expect their user base to be ethical and well-behaved. When abuse occurs, it is difficult to combat, because it takes different forms and evolves as gateways add new, engaging features. The more feature-rich the facility, the more there can be ways to abuse it. Our experience comes from creating and hosting more than 40 science gateways based on the HUBzero platform. These gateways span a variety of disciplines, from nanotechnology to climate modeling to healthcare to engineering education, to name a few. These sites have substantial user communities with tens of thousands, or even hundreds of thousands, of users every year. Because of that, HUBzero has been the target of old email SPAM, but also HTTP proxy abuse, link SPAM, uploaded advertisements of many formats and types, uploaded rootkits, viruses (e.g., hidden in uploaded images), and other (mostly unsuccessful) exploits.

The cost of defending against abuse is a constant drain on human and computing resources. Monitoring and patrolling manually has issues of scale, whereas automatic methods may reduce the usability of the features or produce many false positives. The trade-off between usability, security, and management costs is especially thorny for the hubs we support, because we want to have a low barrier to entry to the public and offer rich functionality on a low budget. The following is an overview of what didn't work, what almost works, what we want to try next, and what we wish we could do.

What didn't work:

• *Modsecurity*

This Apache module is a web application firewall, mused to block recognizable attacks. Recognition relies on many complex, multi line regular expressions. It works for standard web sites, but is mismatched for complex applications like HUBzero. HUBzero has custom communications (for VNC) riding on top of HTTP, as well as a RESTful Application Programming Interface (API), and many components. We were unable to modify the regular expressions used by modsecurity while having assurance that they were still effective at blocking all the attacks they should. In most cases all, we could do, given the time we had, was to entirely disable them. The regular expressions are so complex as to be effectively black boxes to us.

• *Manually managed IP blocklists*

Some abusers kept creating fake accounts for the purpose of abusing science gateways. We attempted to block the IP addresses and networks they used, but apparently they had access to many proxies or other hosts, and the technique was ineffective due to the human cost. This remained true even when the capability was granted to many different administrators and other

interested parties that noticed the abusive accounts. Likewise, manually blocking the IP addresses or networks of various kinds of attackers was ineffective.

What almost works:

• *Spamhaus blocklists*

To decrease the amount of submitted and uploaded SPAM, we used an Apache module that queries the Spamhaus service. This service maintains blacklists of IP addresses used to send SPAM, as well as IPs indicating infected computers. The Apache module allowed people to view the site, but not submit material. This was modestly effective, and in addition made many users realize that their computer was infected. However, we found that organizations forcing their internal users through web proxies would often get their proxies blacklisted. This included universities and hospitals. We whitelist organizations on request, but on some science gateways, the benefit was not deemed worth the inconvenience to legitimate users.

• *Dshield blacklist*

We downloaded and deployed the Dshield blacklist of IP addresses, which is updated regularly based on observed attacks. The list is small and updates are easily automated. Surprisingly, the signature for the list sometimes failed verification, so we skipped some updates. We're not entirely sure of its effectiveness.

• *Automatically managed IP blocks*

We have found that temporarily blocking IP addresses based on undesirable activity was very effective against denial of service attacks, brute force password guessing attacks, automated vulnerability scanning and other recon activities.

• *Anti-virus scanning*

We have found that using ClamAV to scan all uploads was very effective. Nevertheless, we have observed a few false positives and false negatives (confirmed with other anti-virus products through the Virustotal service).

What we still need to try:

• Filter all content through spam detection software, such as SpamAssassin.

• Account vetting and aging to limit what abusers can do with fake accounts. Delays before capabilities are gained makes brand new accounts less valuable and discourages abusers.

What we wish we had or could do:

• An easier to use and customize web application firewall.

• Blocking single IP addresses will be futile with IPv6, due to the very large number of IP addresses available, which makes each IP address effectively disposable. We need software to be able to correlate and aggregate undesirable behavior from ranges of IP addresses, and block appropriately sized networks.

• Use federated lists of bad actors, identified from various factors such as email addresses, IP addresses, and any other useful information, to limit the capabilities of abusers.

• Organizations that cared or monitored whether any of their IP addresses were listed by Spamhaus as used for SPAM or as showing signs of malicious activity.

Network security is an ongoing battle between hosting providers and attackers, with ordinary users caught on the battlefield between them. Ideally, the ordinary users could go about their business unaware of the war raging around them. But occasionally, their experience is interrupted by SPAM, by forced password resets, and by capabilities taken away when attackers abuse them. Still, the ongoing war forces us to constantly improve our support, benefitting all users in the long run.

Cyber Threats: From APT to Non-malicious Insiders

LTC David Raymond, Ph.D.
Army Cyber Institute, West Point, NY, 10996

This unclassified talk will analyze cyber security threats at the strategic, operational, and tactical levels and focus on steps that units should take to minimize the risk of successful attacks on their organizations. While it is helpful to understand the motivations and methods of external actors, it is perhaps more useful to be aware of the vulnerabilities introduced by non-malicious insiders. A combination of inexperienced administrators and untrained or unaware users can make a network extremely vulnerable to even unsophisticated external threat actors.

At the strategic level, the three primary threats are organized cyber crime syndicates, nation-state actors, and hactivists [1] [2]. Organized cyber criminals find ways to monetize their online activities by stealing user account credentials or amassing botnets to rent to spammers. Nation-state actors, some referred to as Advanced Persistent Threat (APT), generally have two goals. First, to gain access to intellectual property to reduce their own research and development costs, and second, to conduct reconnaissance and identify vulnerabilities for future exploitation. Hacktivist groups such as Anonymous are often loosely organized and are motivated to make a political statement. Their techniques range from coordinated DDoS attacks to more high-tech hacking and theft of user information to embarrass the target.

The threat landscape is different at the operational and tactical levels where much of it is internal, and mostly non-malicious. Despite the significant press coverage of some recent high-profile malicious insiders, they are rare compared to non-malicious ones. A 2013 study by Symantec concluded that 64% of system breaches world-wide were caused by system glitches or human factors. Only 35% of breaches are caused by malicious or criminal attack [3].

Personnel shortages and insufficient training often leave units without well-trained mid-grade non-commissioned officers to supervise network device installation and security configuration. Combine this with long and stressful work hours common in deployed environments and the likelihood of error increases. Basic techniques for minimizing attack surface, such as patching and updating, changing default login credentials, and limiting access to critical data, are often not carefully followed. Additionally, these environments can make even careful users vulnerable to social engineering. Furthermore, the close physical proximity of secure and non-secure systems, particularly in deployed environments, makes accidental spillage of classified data commonplace. Extreme care must be taken to put compensating controls in place to reduce the likelihood of compromised classified networks and data.

A third type of insider threat is the circumventer, who knowingly but non-maliciously bypasses security controls in the name of mission accomplishment. An example is a deployed operator who can't get sufficient bandwidth for his unit from trusted NIPR connections so he routes some of the unit's NIPR computers through an untrusted host-nation ISP.

When lack of operator skill results in reduced confidence in preventive and detective controls, compensating controls, such as increased monitoring of vulnerable systems, must be put in place. This talk will delve into the above threats and, time permitting, will provide actionable strategies that organizations should take to mitigate risk in their networks.

---

[1] Verizon RISK Team. (2013) 2013 Data Breach Investigations Report. Technical Report. [Online]. http://www.verizonenterprise.com

[2] Mandiant, Inc. (2013) Mandiant APT1, Exposing One of China's Cyber Espionage Units. Technical Report. [Online]. http://intelreport.mandiant.com

[3] Ponemon Institute. (2013, May) 2013 Cost of Data Breach Study: Global Analysis. White Paper. [Online]. http://www4.symantec.com

Presenter Bio:

LTC David Raymond is an Armor Officer in the U.S. Army and is currently serving as Director of Research in the Army Cyber Center at West Point. He holds Bachelor's and Master's Degrees in Computer Science from the United States Military Academy and Duke University, and a Ph.D. in Computer Engineering from Virginia Tech. LTC Raymond has significant operational experience as an Armor officer, to include serving as a tank platoon leader during Operation Desert Storm and as a tank battalion executive officer during Operation Iraqi Freedom. He is a CISSP, Certified Ethical Hacker (CEH) and holds Global Information Assurance (GIAC) Certifications in Incident Handling, Intrusion Detection, Unix/Linux Security Administration, and Penetration Testing. LTC Raymond teaches senior-level computer networking and cyber security courses at West Point and conducts research on information assurance, network security, and online privacy.

Contact:

LTC David Raymond
Department of Electrical Engineering and Computer Science
Bldg 601
West Point, NY 10996

Email: david.raymond@usma.edu
Office phone: (845) 938-3071

Outline of Talk:

I. High-level threat overview. http://www.verizonenterprise.com
1. Cyber criminals
   - Organized cyber crime
   - Other cyber criminals
2. Nation-state actors and the APT
   - Characteristics of APT
   - Suspected APT actors
      - China (Mandiant Report: http://intelreport.mandiant.com)
      - Russia www.afpc.org/files/november2012.pdf
      - Iran (as a potential emerging APT threat)
3. Hacktivism. Actors and strategies
4. Cyber Terrorism – (Not quite on par with the above 3, but worth mention.) http://www.ctc.usma.edu/posts/defining-cyberterrorism-capturing-a-broad-range-of-activities-in-cyberspace

II. Attacker techniques (briefly)
1. Social Engineering
   - Phishing/spear-phishing
   - Watering holes
2. External attack techniques
   - Network scanning/host enumeration/port enumeration
   - Exploiting vulnerabilities
   - Web application hacking

III. Insider threats
1. Malicious insiders
   - Definition, examples, and difficulty in identifying
2. Accidental insiders
   - Definition and examples
3. “Circumventors”
   - Definition and examples

IV. Overview of effective defense strategies (time permitting)
1. Critical Security Controls (SANS 20 Critical Controls) http://www.sans.org/critical-security-controls/
   - Introduction to controls with pointers to more information
2. Australian Top 4 Mitigation Strategies http://www.dsd.gov.au/infosec/top35mitigationstrategies.htm
3. Risk management and prioritization
   - Risk analysis
   - Return on investment

V. Concluding remarks

Managing Security Policies for Federated Cyberinfrastructure

White Paper submitted to the
2014 NSF Cybersecurity Summit for
Large Facilities and Cyberinfrastructure

Stephen Schwab and John Wroclawski
USC Information Sciences Institute
{schwab, jtw}@isi.edu
July 13, 2014

Significant current and future investment is represented by the collection of large-scale facilities and cyber infrastructures acquired and operated for the benefit of the scientific research community. Unlike commercial infrastructure, these facilities and resources are often unique; significantly specialized to support aspects of the scientific discovery, analysis or computational enterprise; and essential enablers leading to critical breakthroughs and advances.

We observe, however, that the traditional view of such facilities as independent, disconnected, standalone entities serving a single user community is increasingly obsolete. Increasingly, research communities form, and success relies on, large and flexible collaborations, drawing upon unique mixes of expertise, technical capability, and large-scale cyber infrastructure. These collaborations are fluid in nature, often relying on decentralized models of leadership and cooperation at various levels of formality.

In turn, this growth in the scale, dynamism, and multi-disciplinary breadth of human research collaborations drives the technical need to harness significant ensembles of cyber infrastructure, to meet new requirements of both capability and scale. This need increasingly forces the research community toward a new, federated model for architecting, deploying, and managing such infrastructure resources.

This cyber infrastructure model is notable for being fundamentally decentralized, broadly distributed, with individual elements operated by a diverse set of organizations, and subject to complex and competing usage demands. Further, such an infrastructure is typically never 'finished' - sustainability requires that evolvability and fitness for future as well as present purpose be key design considerations. Allowing for future enhancements and augmented capabilities (including those not yet conceived) may well complicate the system, yet is essential because scientific endeavors must leverage the latest in state-of-the-art information technology to keep pace with ever growing data sets and functional demands.

Thus, we observe that forces at two distinct levels – one human, and one technical – drive us towards a view of research cyber infrastructure that is increasingly federated: intrinsically collaborative, large in scale, technically and administratively

decentralized in nature, and capable of being composed dynamically to meet the needs of increasingly sophisticated researcher requirements.

Finally, and critically, this entire eco-system must be approachable and 'user-friendly' if it is to succeed. It must be designed to interact primarily with non-specialist researchers and administrators, rather than IT experts, while being open and accessible to its intended user community, typically from the Internet at large.

These requirements pose significant challenges to scientific infrastructure cybersecurity. The creation and and deployment of security mechanisms and policies that enable confident use of, and organizational control over, all elements of this diverse eco-system *without creating disincentives to collaboration and the creative research process* is one of the central risks facing stakeholders responsible for the funding, construction, management and operation of large facilities and cyber infrastructure.

Catalyzing Federation with Attribute-Based Access Control

We describe in this white paper concrete work that addresses one aspect of this challenge: the creation, management, and implementation of rich, 'audience-appropriate' authorization and access control policy management mechanisms suitable for federated scientific cyber infrastructure in a decentralized, collaborative environment.

We argue that such next-generation policy mechanisms offer a powerful tool for securing national-scale research cyber infrastructures such as those funded by NSF and other U.S. Government agencies in a manner that provides effective security while fostering wide-spread use and catalyzes collaboration.

These mechanisms can effectively, reliably, and *transparently* manage access to facilities, raw data sources, and partially analyzed datasets, while facilitating collaborative use and devolving decision making to researchers spread throughout the collaborating organizations. Pragmatically, such mechanisms complement and build on existing distributed identity management mechanisms such as Shibboleth, which are widely deployed in our target user community.

Our observation builds on a long history of research within the security community directed towards the creation of trust management systems utilizing rigorous, logic-based frameworks that are semantically well defined and amenable to formal reasoning. Such frameworks provide the structure necessary to support federated, decentralized operation while preserving local control, decision-making and prerogatives. They further provide strong assurance of correct operation, together with clearly defined rationale for decisions made, expressible in audit logs for accountability.

To realize these benefits, our focus has centered on the development of Attribute-Based Access Control (ABAC), a specific example of a logic-based trust system, and its deployment in two significant cyber infrastructures. The current ABAC system supports authorization policy expression and enforcement mechanisms that provide:

- Formally grounded policy definition and interpretation. ABAC is based upon rigorous underlying theory and logical formalisms and semantics. Logical underpinnings are leveraged deeply within the system, while users are only exposed to authorization concepts appropriate to their role and domain of expertise;
- Capability to define common vocabulary across communities and organizations. Common, well understood vocabulary may be rapidly adopted for entities, resources, and privileges within common use cases, while preserving the extensibility required to support diverse specialized policies for specialized sub-communities;
- Auditability of requests, authorizations, and policy changes. ABAC decisions result in tangible proofs of authorization derived from distributed policies, or explicit indications of what policies or insufficient privileges resulted in a request being denied;
- Library implementations suitable for incorporation into a range of cyber infrastructures. ABAC software provides a compact library implementation and language bindings for several of the standard programming languages used throughout the Networks, Grid, Cloud and Cyber infrastructure communities.

Together, these capabilities provide a strong foundation for the implementation of strong, secure access and use management capabilities within large-scale, federated cyber infrastructures, while simultaneously facilitating the key objectives of flexible collaboration and local control.

Development Status and Lessons Learned

As described above, ABAC's capabilities are realized in a concrete implementation, colloquially known as *libabac*, that is suitable for incorporation into a range of cyber infrastructures.

The ABAC software distribution provides the basis for meeting the current needs of a cyber infrastructure eco-system relying on federated deployment and decentralized operations. It also lays the groundwork to support rapid evolution of national-scale cyber infrastructure, anticipating new communities of users, new domains of inquiry and concomitant patterns of use, new classes of large-scale facilities, and new patterns of access and connectivity. ABAC has been under development and use for several years, within significant cyber research communities. We briefly describe two major deployments of the *libabac* software.

Our first use case is within the DETER Cybersecurity Testbed[1] (DETERLab) to enforce distributed authorization in support of wide-area, large-scale federated experiments. DETERLab is an internationally available infrastructure operated in support of experimental cybersecurity research, originally sponsored by NSF and

[1] http://deter-project.org

now sponsored by the US Department of Homeland Security. Our requirements and objectives for the DETER project provided initial motivation for the development of *libabac.* The DETER Federation model realizes *federated experimental scenarios,* realizing distinct sub-components of an overall experimental scenario as locally embedded experiments within multiple, distributed facilities, each contributing resources to the scenario. Each facility enforces authorization policies on a per-experiment basis using ABAC. Resources managed by ABAC within the DETER federation range from basic computing and network facilities to dedicated, unique elements with unique use policies, such as hardware-based modeling engines for electric power systems.

Our second use case includes the prototyping and planned production adoption of ABAC within NSF's Global Environment for Network Innovation[2] (GENI) project. The GENI project represents a major, multi-year investment by NISF CISE to create a nationwide suite of infrastructure supporting "at scale" research in networking, distributed systems, and novel networked applications. Within GENI, ABAC will enable a diverse set of testbed control frameworks and resource aggregate managers, operating at institutions throughout the US, to coordinate policies governing computing and communication resource use by network researchers performing experiments embedded across the GENI infrastructure. In this capacity, it will replace an *ad hoc* first generation authorization management infrastructure that has demonstrated significant limitations in current use.

In particular, a (surprisingly simple, yet surprisingly strong) motivating use case for ABAC within the GENI community is the requirement to allow researchers of varying skills and experience to designate hosted tools and portals to act on their behalf *without* obligating these researchers to upload secret keys or passwords to 3rd party servers. ABAC's logic-based approach supports this requirement by allowing individual researchers and/or their supporting organizations to easily express formal statements regarding delegation of authority to 'speak for' the researcher or organization.

This intuitively simple capability, easily expressed within an ABAC framework, substantially enhances broader security goals by limiting risk of disclosure of passwords or private credentials. Beyond this initial example, ABAC further strengthens GENI security by implementing a manageable and understandable least-privilege capability, providing for the first time a principled mechanism within the GENI architecture that can delegate specific limited privileges to a GENI entity without granting the delegated entity the full power of a user's credentials.

To date, our experience with ABAC concepts in deployed systems has demonstrated the generality and value of integrating trust management and authorization policy grounded in formal logic. Looking forward, our key goal is to enhance the intuitive approachability and usability of such a logic-based system.

---

[2] http://geni.net

To accomplish this, it is crucial to recognize that every research community and domain of interest has its own preferred concepts and language with which to describe principals, resources, and authorization policy. It is further important to understand that the management of scientific instruments and cyber infrastructure resources, even those of extreme complexity or value, must be approachable by the relevant researchers themselves or administrators acting on their behalf, rather than requiring IT experts.

Consequently, a core focus of our current work is the creation of interfaces and UI tools that capture and reflect concepts and policies intuitively, in the natural terms used by the research and education community. A further objective is the development of approaches that allow members of a specific sub-community to tailor and shape their own environment without expert assistance. As these usability-oriented capabilities develop, our goal is to offer ABAC, at both the conceptual and implementation level, to the larger cyber infrastructure community as a fundamental building block that addresses key security requirements of large-scale federated systems, meeting both current and future needs of the broad research community whose objectives motivate the ongoing cyber infrastructure agenda.

Appendix E
Training Descriptions

Training Sessions | Aug 26 | 2014 NSF Cybersecurity Summit

Tuesday, August 26 will feature a full day of training, available to all registrants. All but the Bro session are half-day offerings. Seating may fill for some or all sessions, and pre-event registration for individual sessions is required to reserve a seat. Please register by August 19 to guarantee seating, and help us make final preparations. Direct inquiries to Craig Jackson (scjackso@indiana.edu).

Concurrent Morning Sessions

Bro Platform Training Workshop (Full Day)

Instructors: Robin Sommer & Justin Azoff (Bro Center for Expertise)

Bro is a powerful network analysis framework used for security monitoring and network traffic analysis. The user community includes major universities, research labs, supercomputing centers, and government and corporate organizations. In order to gain the most utility out of Bro we encourage users to attend training workshops and participate in the greater online community.

The Bro development team will deliver a full day workshop focusing on such topics as installation and administration, examining logs, learning out-of-the-box and custom Bro scripts, and the Bro Intelligence Framework.

The morning session will focus on explaining what is Bro, how it is used, and out-of-the-box features. The afternoon session will focus more on hands-on exercises and programming in the Bro scripting language.

Required materials: A laptop with an ssh client and VirtualBox installed

Developing Cybersecurity Programs for NSF Projects

Instructors: Jim Marsteller, Susan Sons, Craig Jackson, Jared Allar (CTSC)

Slides (PDF)

Audience: Principal Investigators, Security Professionals, Center and Operational Managers, NSF Program Officers

Team members of the Center for Trustworthy Scientific Cyberinfrastructure (CTSC) will present an interactive

half day session on developing cybersecurity programs for NSF science and engineering projects. The session will be based on a cybersecurity planning guide (see, **trustedci.org/guide**) developed over the past six months with input from the Daniel K. Inouye Solar Telescope (DKIST) project and other members of the CI community.

The purpose of this session is to offer a streamlined approach to developing comprehensive cybersecurity programs for NSF funded projects. The guide has been developed to address the information security requirements outlined in the NSF cooperative agreements. This session will include an instructional review of the cybersecurity planning guide and supporting templates, which can be used to jumpstart program and policy development. Some of the topics that will be covered include:

- Building or Improving a Cybersecurity Program
- Unique and Critical Science Requirements, Constraints, and Security Controls
- Information Security Policies and Procedures
- The Role of Project Leadership
- 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, and more.
- Program Assessment and Evaluation

While this session will be instructional in nature, it is also intended to be an interactive session to seek constructive feedback from attendees to further improve the guide. There will be significant opportunities for discussion and Q&A.

Vulnerabilities, Threats, and Secure Coding Practices

Instructors: Barton P. Miller & Elisa Heymann

Slides (PDF)

Security is crucial to the software that we develop and use. With the growth of both Grid and Cloud services, security is becoming even more critical. This tutorial is relevant to anyone wanting to learn about minimizing security flaws in the software they develop. We share our experiences gained from performing vulnerability assessments of critical middleware. You will learn skills critical for software developers and analysts concerned

with security.

This tutorial starts by presenting basic concepts related to threats, weaknesses and vulnerabilities. We will also show you how to think like an attacker. The rest of the tutorial presents coding practices that lead to vulnerabilities, with examples of how they commonly arise, techniques to prevent them, and exercises to reinforce you skills in avoiding them. Examples come from a wide variety of languages, including Java, C, C++, C#, Perl, Python, and Ruby, and come from real code belonging to Cloud and Grid systems we have assessed. This tutorial is an outgrowth of our experiences in performing vulnerability assessment of critical middleware, including Google Chrome, Wireshark, Condor, SDSC Storage Resource Broker, NCSA MyProxy, INFN VOMS Admin and Core, and many others.

Concurrent Afternoon Sessions

Bro Platform Training Workshop (continued)

See full description above.

HPC, HIPAA, and FISMA: Meeting the Regulatory Challenge through Effective Risk Management

Instructors: Bill Barnett & Anurag Shankar (Indiana University)

Slides (PowerPoint)

With biomedical research emerging as a formidable computing challenge needing support, high performance computing (HPC) is now face to face with regulatory compliance. New language in government grants and contracts is or will soon be requiring compliance with federal cybersecurity standards for protecting research data, whether or not biomedical. This half-day training session will familiarize the participants with relevant regulations, how they apply to HPC, the challenges they present, and offer a standards-based risk management approach to tackling them.

Topics covered will include:

- *HIPAA and FISMA Demystified*. History and introduction to the regulations, what they mean for HPC shops, what they do not.
- *The NIST Risk Management Framework*. Managing information security risk (NIST 800-39), conducting risk assessments (NIST 800-30), security and privacy controls (NIST 800-53), and assessing the controls (NIST 800-53A).

- *Leveraging the Framework*. Scoping, planning, implementing risk assessments, risk mitigation through selected security controls, documentation, ongoing risk management, reviews, and training, implementation at IU as example.

Incident Response Training

Instructors: Warren Raquel, Randy Butler, & Patrick Duda (NCSA)

Slides (PowerPoint part 1, PowerPoint part 2)

Computer incident response is a required capability for any project or activity that is running internet connected services. This tutorial will provide basic information on setting up an incident response program so that the students can prepare their project team or organization for handling an incident investigation. The initial focus of the tutorial will be on identifying the processes, policies, information, and monitoring services that will be required to effectively respond to a security incident. This first section will additionally discuss investigation and analysis tools that might be useful for your investigations. The second part of the tutorial will identify a collection of questions that the incident response team can use to guide them through both the investigation and the mitigation process. The final section will highlight several actual security incidents. Each of these incidents will be discussed in detail starting with how the incident was discovered and then continue through the investigation and mitigation process. The participant should leave the session with an understanding of the basic steps needed to create an incident response program and what to do when an incident occurs.

Appendix F
Table Topic Summaries

Table Topic Summaries

2014 NSF Cybersecurity Summit Large Facilities and Cyberinfrastructure

On August 27, Program Committee members and community members led 5 table talks during lunch. We invited brief summaries of the discussions; these are included below.

Incident Response and Heartbleed

Discussion Leaders: Rodney Petersen and Kim Milford

The table discussion began with a clarification of the topic. First, the Heartbleed incident was intended as an example and not the entire focus of discussion. Second, there was another table discussion called Federated Incident Response which should probably be recast as how to respond to incident when federated credentials or resources are at risk. The discussion then turned to individuals around the table sharing how they responded at a local level (i.e., campus, research organization, etc.) to Heartbleed and what policies and procedures that they have in place to guide their practice. Some, but not all, organizations reported having a formal incident response plan and CERT or CSIRT that convenes for major incidents. A major point of discussion was the lack of a clear definition or trigger for what qualifies as a major incident. Most individuals reported that they pay attention to national information sharing organizations and sources for the determination of the severity of the incident. They also reported sharing information, as needed, up the chain of command so that there were no surprises if and when an incident moved from minor to major. The group agreed that it would be helpful if there was some agreed upon or national classification system (e.g., similar to prior homeland security color codes) they could apply to their local situation or a national event. There was considerable discussion about the similarities with physical events (e.g., active shooters, natural disasters, etc.) and the need for better integration with emergency management planning and response; however, there was also recognition that the physical threats and emergencies were easier to identify and classify.

We also discussed the need for exercising incident situations, both locally and regionally or nationally. The REN-ISAC was generally recognized as a good source of information and an entity with the potential to fully develop a CSIRT capability on behalf of the Research and Education Community. There are some operational limitations, including communications (secure video and voice) and coordination with other entities such as the Higher Education Information Security Council and other research consortiums. The REN-ISAC is viewed as a trusted source for information, although the reliance on email as the primary form of communication or information sharing makes it very difficult to consume information quickly or consult it later as a resource. There needs to be more consolidation or synthesis of the information to make it more useful and to make the most efficient use of member time. There was also discussion of the need to extend access to the information from the REN-ISAC to others (e.g., system administrators, network administrators,etc.) within the REN-ISAC membership community. Some members reported the current information sharing model is too limiting and inhibits the ability of some organizations to take appropriate action. Trust remains a critical ingredient to having an effective incident response. Trust is required at the local organizational level and across organizations which is why events like the NSF Cybersecurity Summit are critical for building relationships and establishing security points of contact.

Federated Incident Response:

Discussion Leaders: Tom Barton and Craig Jackson

- There is almost certainly a desire for information sharing around compromised credentials between the related parties.
- Challenges and open questions include:
    - Presently there is no standard procedure, structure, or protocol. When this information is shared, it is entirely ad hoc.
    - Who to contact at federated org and whether/how much to trust them, and how to broker those trust relationships. A question was raised re: whether a trusted channel would be necessary.
    - Another challenge is time/effort. An IdP may be sitting on a number of exploited accounts at any given time. Sifting through these and communicating with the relevant partners is too much.
    - Identifying the technical means to relay info, and standard of practice to qualify an org to participate in the technical means.
    - Thus, automated and/or routine processes would be desirable (or perhaps necessary) to support federated IR. However, even with an automated process, data must be entered. Do we understand what is the minimum effective communication?
- Trust is likely a key ingredient of any solution. Possible to cobble this together from peer-to-peer relationships, or leverage something like REN-ISAC?
- R&E Federations should probably become trust brokers for this, in addition to other orgs unknown at this time.

Enhancing Research Support By Enabling a Secure Cyberinfrastructure

Discussion Leaders: Ardoth Hassler and Kevin Thompson

Funding from the NSF CC*IIE program, and its predecessor CC-NIE program, is serving to enhance research by improving researchers' connectivity to the campus backbone, campus backbone connectivity to Internet2, and/or implementing a campus Science DMZ to enable collaboration among researchers, access to resources, including those in the Cloud. (See, http://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf14521)

Some issues identified/topics discussed include:

- The approach to implementing/running a Science DMZ is different from the "central IT" approach to running a "commodity" campus network.
- Issues smaller institutions face include with lack of appropriately-experienced staff to deploy advanced networking, a Science DMZ, for example, and more focus on administrative and student needs than an overall understanding of researchers' needs to obtain compute cycles and move big data.
- Throwing equipment at a campus network doesn't help researchers solve the problems that arise when research is competing with undergraduate network traffic.
- An identified need for polices for a Science DMZ, e.g. that it is separate from administrative and general student use. [NB: some policy information is available at http://fasterdata.es.net/science-dmz/]

- A need to articulate best practices for a Science DMZ in a HIPAA context.
- The community is working with NIST on guidance re Science DMZs and distributed firewalls.

AWS and the Impact on High-End Research

Discussion Leaders: Jim Marsteller and Barb Fossum

Miron Livny, PI of the Open Science Grid (OSG) has been pursuing the use of incorporating AWS along with the Condor software. He shared his experience in leveraging AWS for the OSG project. It is extremely expensive to move around petabytes of data per day so they have more interest in using AWS for data storage to help in moving the data around rather than the computation resources.

Some bullet points from the discussion:

- What kinds of negotiations needed to be made with AWS to implement these services? It's very important to have an emergency back-out process and the paperwork should be in place to find a way to smoothly transition if such a response is needed. Service Level Agreements are also important. AWS reliability has not been able to attain the availability that has traditionally been demanded of high-end data centers.

- What does Amazon get from offering this service? Good PR, student workforce, and in the future they hope to develop a business plan to earn money from this endeavor. OSG will try to leverage these services for the future. There are so many factors that come into play - the ability to scale, the security isn't transparent to the HPC user, how do you request information on login etc to determine the health of the system?
- How do you deploy to 50K cutters - done by software that will deploy by the 1K. There are timelines for deploying large amounts of resources but can e done through the open science grid. Using AWS as hot sites for back-up of data during transfer of data between jobs and resources.
- Is there data on the failover? There is a layer that does the job submission to some resource and it doesn't matter what the compute resource is. There does need to be knowledge of the resource so that the software doing submission can determine the way to talk to the resource- i.e. google, amazon, etc.
- HPC and AWS:HPC is vulnerable because they use MPI and if one node goes down the entire job fails. If google or amazon used MPI their success would be nil. When you get HPC nodes they guarantee a certain latency. AWS offered HPC services are much more expensive than the storage services.

Identity and Access Management: State of Practice and Future Directions

Discussion Leaders: Mark Servilla & Jim Basney

- Federated identity management is considered important, but not a hard requirement
- No experience with InCommon
- Perception that federated identity management is still a "hard" problem to solve
- Globus Nexus is a recommended cloud-based solution for identity and group management

Improving Diversity

Discussion Leaders: Tony Baylis and Amy Apon

No summary of general interest to report. Attendance at this table talk was low.

Appendix G
Listing of Attendees and Organizations

List of Attendees

2014 NSF Cybersecurity Summit for Large Facilities and Cyberinfrastructure

First Name	Last Name	Institution/Organization Provided
Andrew K	Adams	Pittsburgh Supercomputing Center
Joshua	Alexander	University of Oklahoma / OU Supercomputing Center for Education & Research (OSCER)
Jamie	Allan	NSF
Jared	Allar	Pittsburgh Supercomputing Center
Amy	Apon	Clemson U.
Winston	Armstrong	San Diego Supercomputer Center at UCSD
Justin	Azoff	NCSA
Saurabh	Bagchi	Purdue University
Steve	Barnet	UW-Madison - IceCube
Bill	Barnett	Pervasive Technology Institute / Indiana University
Tom	Barton	U. of Chicago
Jim	Basney	UIUC
Tony	Baylis	Lawrence Livermore National Laboratory / Rochester Institute of Technology / Air Force ROTC
Steve	Beaty	NCAR
Steve	Beher	National Superconducting Cyclotron Laboratory
Greg	Bell	ESnet / Lawrence Berkeley National Laboratory
Emily	Bell	Rep. Susan Brooks
Jordan	Berg	National Science Foundation
Steve	Berukoff	NSO / DKIST
Kate	Bonvechio	Subcommittee on Emergency Preparedness, Response, and Communications / House Committee on Homeland Security
Devin	Bougie	Cornell Laboratory for Accelerator-based ScienceS and Education -- CLASSE
Jasmine	Bowers	North Carolina A&T State University
Amy	Butler	George Washington University
Randy	Butler	NCSA / University of Illinois
Neil	Canfield	NSF
Jaime	Combariza	Johns Hopkins University
Leslee	Cooper	Indiana University / Center for Applied Cybersecurity Research
Michael	Corn	Brandeis University
Joel	Cutcher-Gershenfeld	School of Labor and Employment Relations (LER), and National Center for Super Computing Applications (NCSA), University of Illinois, Urbana-Champaign
Jeannette	Dopheide	National Center for Supercomputing Applications
Patrick	Duda	UIUC
Don	DuRousseau	The George Washington University
Walter	Dykas	DOE Office of Science
Rudolf	Eigenmann	NSF
Jeremy	Epstein	National Science Foundation
Shane	Filus	Pittsburgh Supercomputing Center
Terry	Fleury	University of Illinois / NCSA
Barb	Fossum	NEES, Purdue University
Phil	Gates	International Ocean Discovery Program
Jill	Gemmill	Clemson University
Kim	Gillies	Thirty Meter Telescope Observatory
Bret	Goodrich	National Solar Observatory
Steve	Grandi	National Optical Astronomy Observatory
Robert	Grossman	University of Chicago
Christopher	Gullo	Lawrence Livermore National Laboratory / Rochester Institute of Technology / Air Force ROTC
David M.	Halstead	National Radio Astronomy Observatory

Ardoth	Hassler	Georgetown University
Victor	Hazlewood	University of Tennessee
Randy	Heiland	CACR, Indiana University
Elisa	Heymann	University of Wisconsin-Madison / Universitat Autonoma de Barcelona
Bob	Houtman	National Science Foundation
Craig	Jackson	IU
Cliff	Jacobs	Clifford A. Jacobs Consulting, LLC
Peter	Jensen	National High Magnetic Field Laboratory
Ken	Klingenstein	University of Colorado
Scott	Koranda	UW-Milwaukee
Bill	Kramer	NCSA / University of Illinois
Miron	Livny	Morgridge Institute for Research
Paul	Lordier	California State University Sacramento
James	Marsteller	CMU
Celeste	Matarazzo	Lawrence Livermore National Laboratory
Natalie	Matson	Committee on Homeland Security
Charles	McElroy	Information Systems / Case Western Reserve University
Michael	McLennan	Purdue University
Kishor	Mehta	National Science Foundation
Nathaniel	Mendoza	Texas Advanced Computing Center
Pascal	Meunier	HUBzero, Purdue University
Kim	Milford	REN-ISAC
Barton	Miller	University of Wisconsin
Tim	Minick	Gemini Observatory
W. Clay	Moody	Clemson U.
Christopher	Morrison	Association of Universities for Research in Astronomy
Nick	Multari	PNNL
Jose	Munoz	NSF
Pat	Murphy	National Radio Astronomy Observatory
Andrew	Neff	NEON, Inc.
Anita	Nikolich	NSF
Amy	Northcutt	NSF
Joy	Pauschke	NSF
Rodney	Petersen	EDUCAUSE
Donald	Petravick	NCSA
Francesco	Pontiggia	Brandeis University
Sarah	Portwood	Indiana University / Center for Applied Cybersecurity Research
Irene	Qualters	NSF
Warren	Raquel	NCSA
David	Raymond	Army Cyber Institute
Ryan	Richmond	AURA
Thomas	Rieker	NSF / DMR Facilities
Brian	Rohler	NEEScomm / Purdue University
Matthew	Rosenquist	Intel
Jim	Rosser	Texas A&M University
Doris	Schiöberg	ICSI
Ryan	Schmitz	National Ecological Observatory Network
Phyllis	Schneck	US Dept of Homeland Security
Steve	Schwab	USC Information Sciences Institute
Mark	Servilla	Long Term Ecological Research Network/University of New Mexico
Anurag	Shankar	Indiana University
Nigel	Sharp	National Science Foundation
Abe	Singer	LIGO
Robin	Sommer	ICSI/LBNL
Susan	Sons	Indiana University, CACR / CTSC
Kristin	Spencer	NSF / BFA /DACS
Brian	Stengel	Technology Services (CSSD), Univeristy of Pittsburgh

Denise	Sumikawa	ESnet
Guebre X.	Tessema	NSF
Vic	Thomas	BBN Technologies
Kevin	Thompson	NSF
Steve	Tuecke	U. of Chicago
Heidi	Wachs	Gartner
Ralph	Wachter	NSF
Amy	Walton	NSF
Jerry	Wanetick	Scripps Institution of Oceanography (SIO); UC San Diego
Von	Welch	Indiana University
Carol	Wilkinson	NSF Large Facilities Office
Jim	Williams	Internet2
Paul	Wisniewski	University of Wisconsin
Shijie	Yang	Cornell High Energy Synchrotron Source

Appendix H
Attendee Survey summary report

Q1 Which options best describe your job or position? Check all that apply.

Answered: 44 Skipped: 0

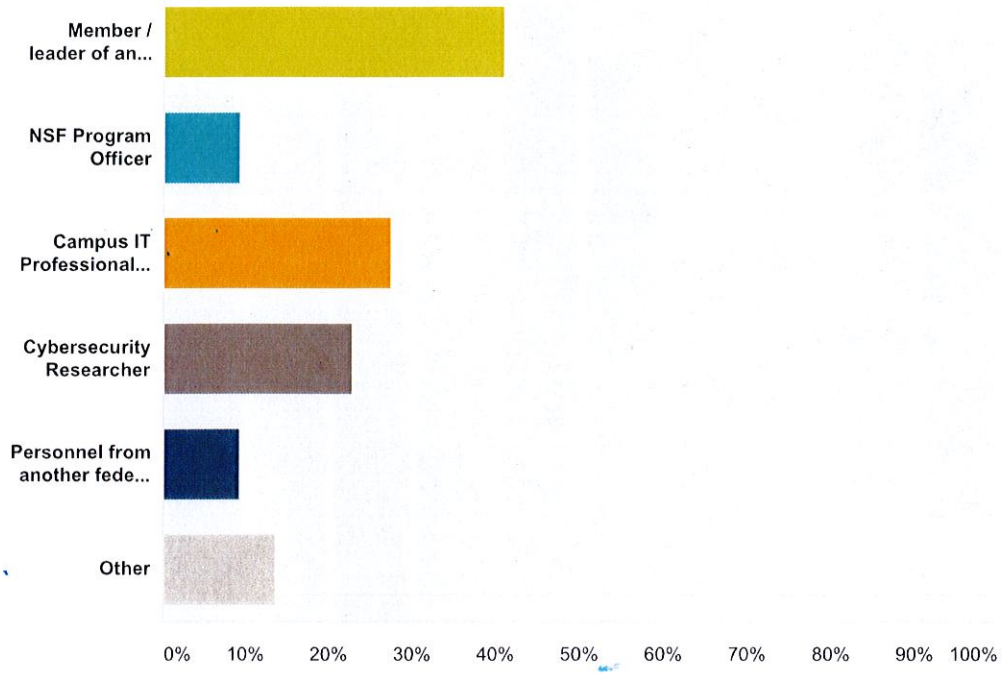

Answer Choices	Responses	
Member / leader of an NSF project	40.91%	18
NSF Program Officer	9.09%	4
Campus IT Professional / CIO	27.27%	12
Cybersecurity Researcher	22.73%	10
Personnel from another federal program (NSA, DOE/ESNet, etc.)	9.09%	4
Other	13.64%	6

Total Respondents: 44

#	If applicable, please state your NSF Project and/or affiliated NSF Directorate. Other comments or clarifications are welcome.	Date
1	National Optical Astronomy Observatory	
2	GENI	
3	National Solar Observatory/AST/MREFC	
4	Mathematical & Physical Sciences (MPS)	
5	Cornell Laboratory for Accelerator-based ScienceS and Education Cornell High Energy Synchrotron Source	
6	Gemini Observatory	
7	NRAO, AST	

8	GEO Sciences
9	EarthCube - Geosciences Directorate
10	Our project has a small NSF grant at this time.
11	National Center for Genome Analysis Support, Bio Directorate.
12	IODP
13	Student, CSUS Undergraduate Computer Science, Concentration in Information Assurance and Cyber Security
14	CTSC
15	CISE - CCF, CNS ENG - CMMI
16	GEO/OCE/Ocean Drilling Program
17	IceCube Project
18	Contracting/Agreements Officer, NSF BFA/DACS
19	CISE

Where do you work primarily?

Answer Choices	Responses	
State/Province:	100.00%	43
Country:	100.00%	43

#	State/Province:	Date
1	Tucson, AZ	
2	tennessee	
3	MN	
4	MI	
5	VA, MD, DC	
6	Oklahoma	
7	Colorado	
8	Washington	
9	VA	
10	New York	
11	SC	
12	Indiana	
13	Hawaii	
14	Illinois	
15	CA	
16	California	
17	VA	
18	CA	
19	New Mexico	
20	Virginia	
21	VA	
22	OH	
23	NY	
24	washington, dc	
25	CA	
26	Indiana	
27	Texas	
28	North Carolina	
29	CA	
30	Illinois	

31	D.C.	
32	Indiana	
33	VA	
34	Pennsylvania	
35	WI	
36	Virginia	
37	New York	
38	Virginia	
39	Massachusetts	
40	Wisconsin	
41	Illinois	
42	Washington	
43	New York	
#	**Country:**	**Date**
1	USA	
2	usa	
3	USA	
4	USA	
5	USA	
6	USA	
7	USA	
8	USA	
9	USA	
10	US	
11	USA	
12	USA	
13	USA	
14	USA	
15	USA	
16	USA	
17	US	
18	US	
19	USA	
20	USA	
21	USA	
22	USA	
23	USA	
24	usa	

25	USA
26	USA
27	USA
28	US
29	USA
30	USA
31	USA
32	USA
33	US
34	USA
35	USA
36	US
37	US
38	US
39	USA
40	USA
41	USA
42	DC
43	USA

Q3 How would you characterize your job in relationship to cybersecurity? Please check all that apply.

Answered: 44 Skipped: 0

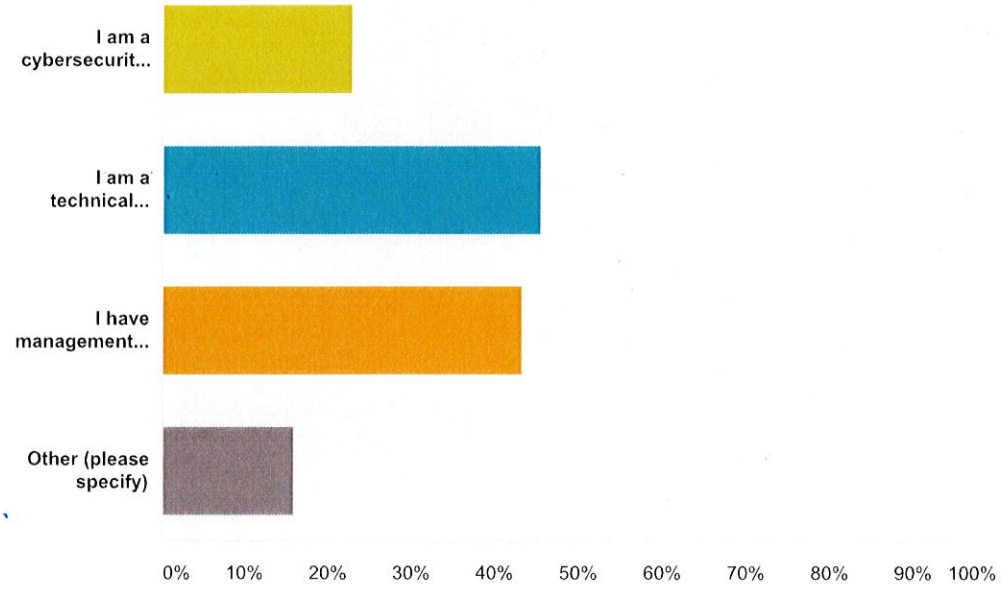

Answer Choices	Responses	
I am a cybersecurity professional	22.73%	10
I am a technical professional who has knowledge of cybersecurity	45.45%	20
I have management responsibility for cybersecurity	43.18%	19
Other (please specify)	15.91%	7

Total Respondents: 44

#	Other (please specify)	Date
1	I direct a national ISAC	
2	consultant helping the community organize	
3	Researcher	
4	Student	
5	Cybersecurity researcher	
6	include requirement in award formation & monitor compliance with cyber terms	
7	Student	

Q4 What sessions of the summit did you attend? Check all that apply.

Answered: 44 Skipped: 0

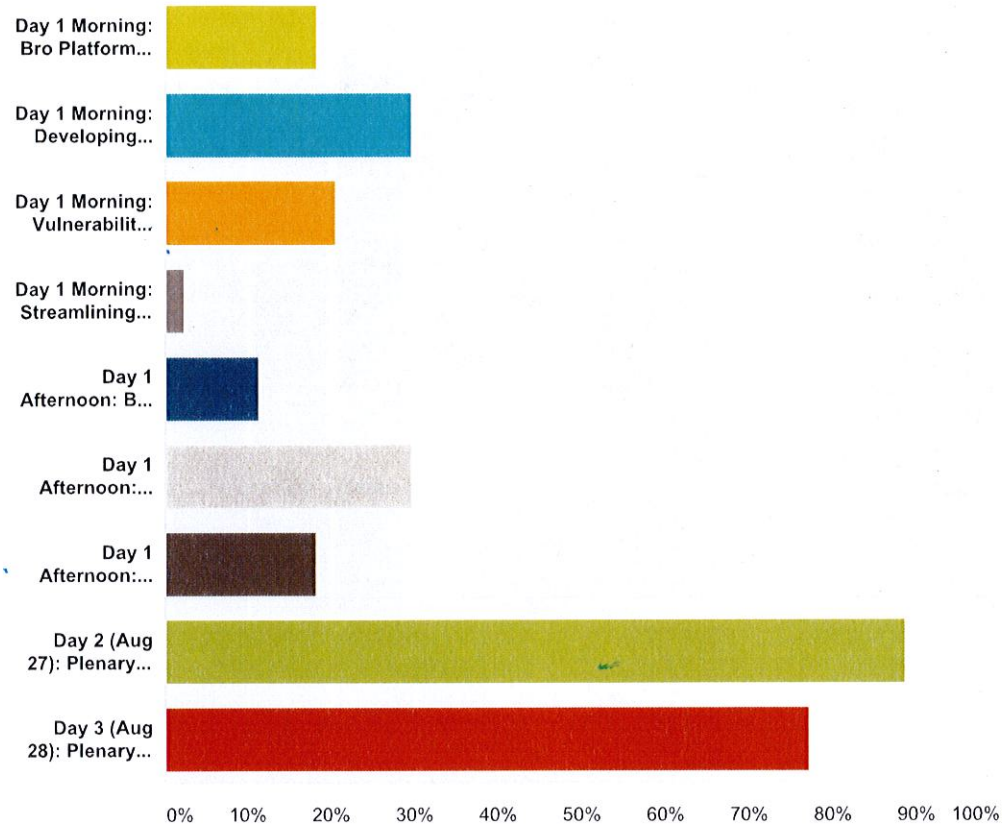

Answer Choices	Responses	
Day 1 Morning: Bro Platform Training Workshop	18.18%	8
Day 1 Morning: Developing Cybersecurity Programs for NSF Projects	29.55%	13
Day 1 Morning: Vulnerabilities, Threats, and Secure Coding Practices	20.45%	9
Day 1 Morning: Streamlining Collaboration with InCommon and Identity Federations	2.27%	1
Day 1 Afternoon: Bro Platform Training Workshop (continued)	11.36%	5
Day 1 Afternoon: Incident Response Training	29.55%	13
Day 1 Afternoon: HPC, HIPAA, and FISMA: Meeting the Regulatory Challenge through Effective Risk Management	18.18%	8
Day 2 (Aug 27): Plenary Session	88.64%	39
Day 3 (Aug 28): Plenary Session	77.27%	34

Total Respondents: 44

Q5 How would you rate your overall experience with the 2014 summit?

Answered: 44 Skipped: 0

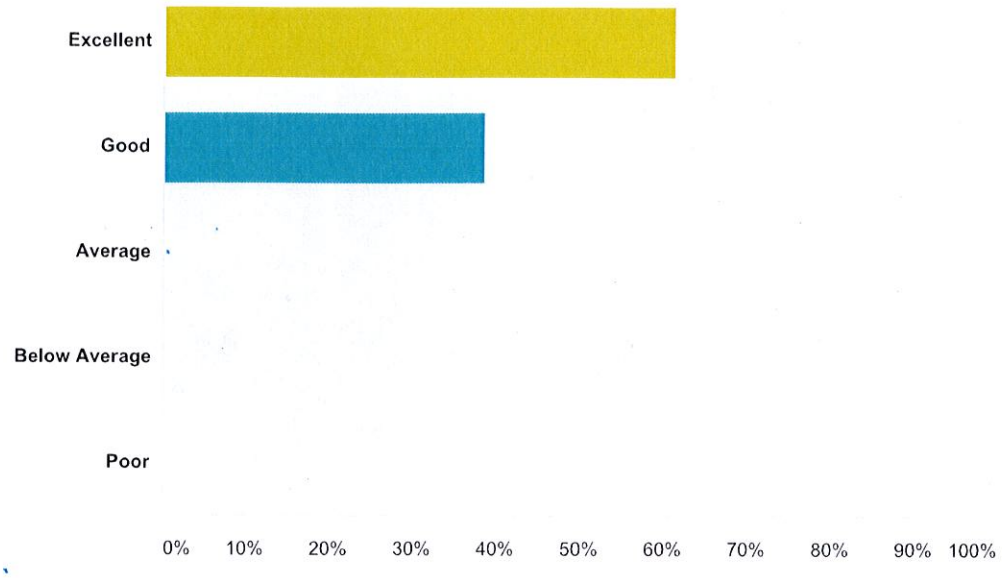

Answer Choices	Responses	
Excellent	61.36%	27
Good	38.64%	17
Average	0.00%	0
Below Average	0.00%	0
Poor	0.00%	0
Total		44

Q6 Please rate your experience with the 2014 summit in these areas:

Answered: 44 Skipped: 0

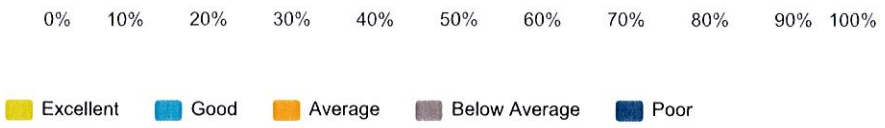

	Excellent	Good	Average	Below Average	Poor	Total Respondents
Topics Addressed	65.91%	31.82%	2.27%	0.00%	0.00%	
	29	14	1	0	0	44
Quality of Presentations	43.18%	56.82%	0.00%	0.00%	0.00%	
	19	25	0	0	0	44
Logistics & Organization	77.27%	22.73%	0.00%	0.00%	0.00%	
	34	10	0	0	0	44
Venue	70.45%	22.73%	6.82%	0.00%	0.00%	
	31	10	3	0	0	44

Q7 Was this summit better than what you expected, worse than what you expected, or about what you expected?

Answered: 44 Skipped: 0

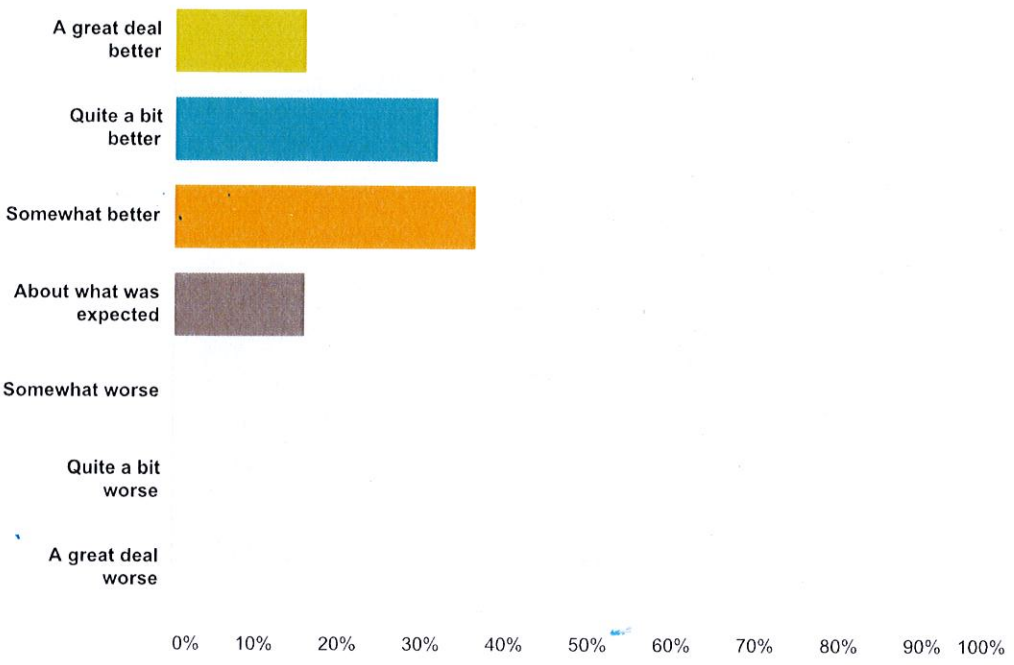

Answer Choices	Responses	
A great deal better	15.91%	7
Quite a bit better	31.82%	14
Somewhat better	36.36%	16
About what was expected	15.91%	7
Somewhat worse	0.00%	0
Quite a bit worse	0.00%	0
A great deal worse	0.00%	0
Total		44

Q8 How useful to your work was the information discussed at the summit?

Answered: 44 Skipped: 0

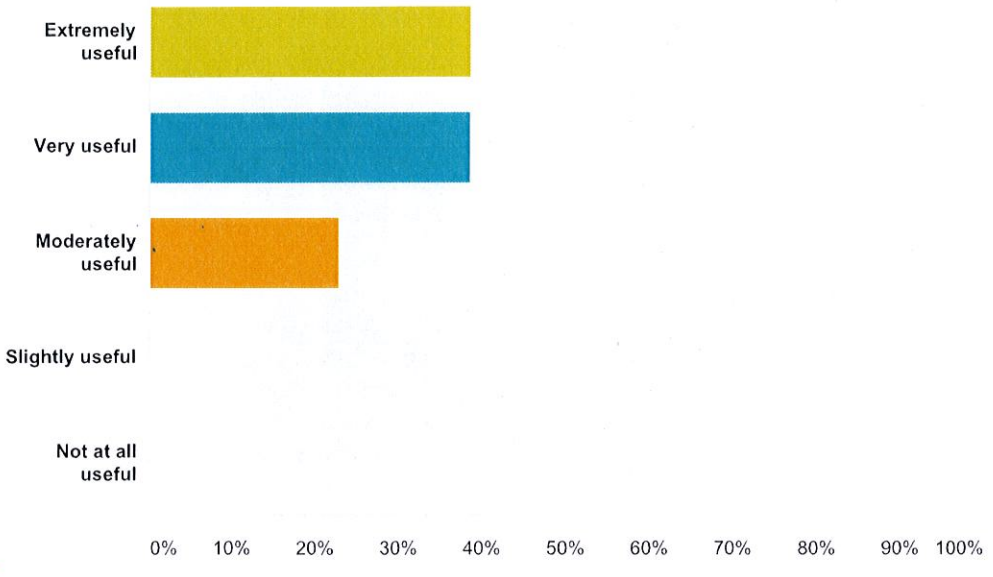

Answer Choices	Responses	
Extremely useful	38.64%	17
Very useful	38.64%	17
Moderately useful	22.73%	10
Slightly useful	0.00%	0
Not at all useful	0.00%	0
Total		**44**

Q9 How would you describe the balance between structured presentations and informal networking opportunities?

Answered: 44 Skipped: 0

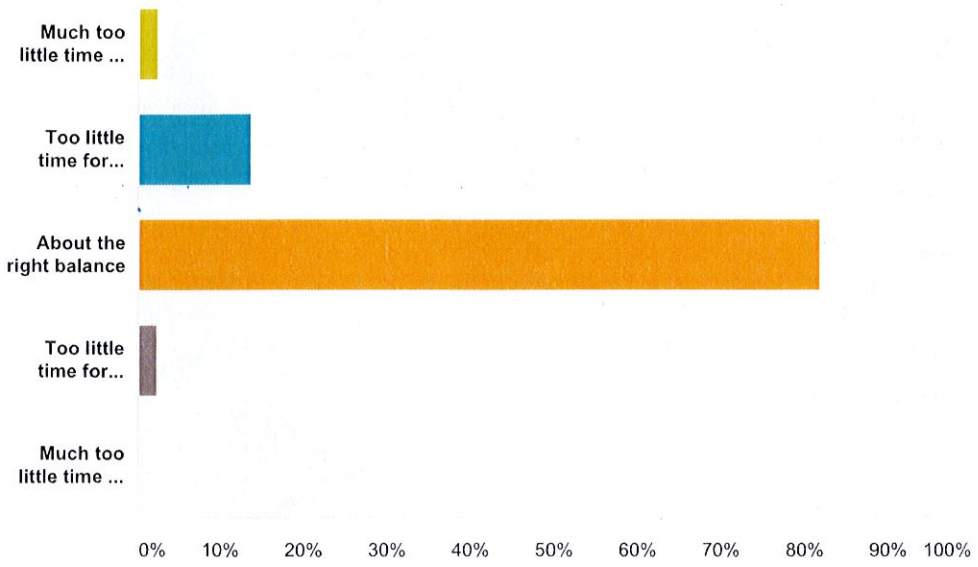

Answer Choices	Responses	
Much too little time for informal networking	2.27%	1
Too little time for informal networking	13.64%	6
About the right balance	81.82%	36
Too little time for structured presentations	2.27%	1
Much too little time for structured presentations	0.00%	0
Total		44

Q10 Would you like to attend future summits?

Answered: 44 Skipped: 0

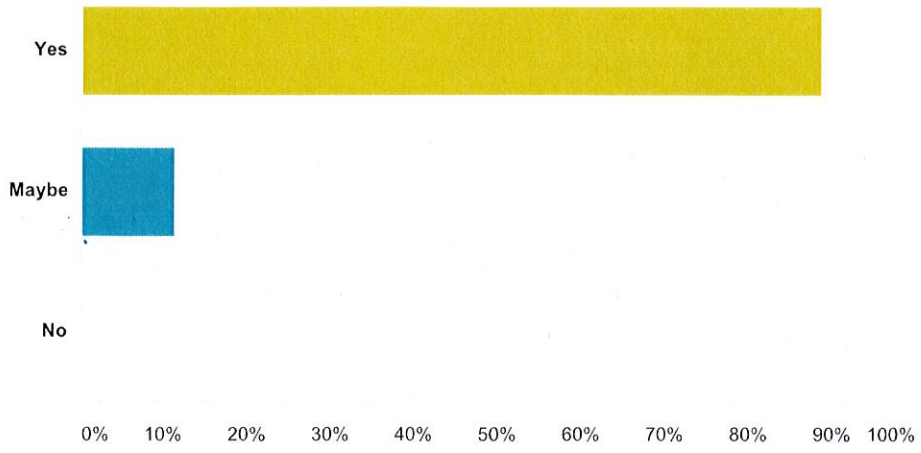

Answer Choices	Responses	
Yes	88.64%	39
Maybe	11.36%	5
No	0.00%	0
Total		44

Q15 How can we improve the summit experience in the future?

#	Responses	Date
1	Ensure that the temperature in the conference rooms is at a comfortable level.	
2	How about creating small break-out sessions that are similar to Birds of a Feather (BoF) sessions? People of similar concerns/ideas/situations meet for 30-60 minutes to discuss their current plans/situations.	
3	I attended the "NSF Facilities" workshop, and thought it could be improved with additional hands-on work. The information content - understanding NIST SPs, FISMA, etc. - was useful for sure, and I think including reference materials is a great use of time and of the webpages.	
4	I would like to see some presentations on what research topics are being pursued today and what we can expect to see in the future.	
5	Retain the full-day training aspect.	
6	While the day 2 lunch discussion groups was a great concept the actual discussions suffered due to numerous conversations (and related noise) happening in the same room. It was noted by myself and others that some of the conference attendees/panel members were repeatedly "hijacking" the microphone and even offering advice and commentary that conflicted with the presentation or presenter. I've been to numerous IT conferences over the course of the past 20-years, and I understand that such a situation can be difficult to mitigate while still extending invitations on a broad scale. Not sure what to suggest here...	
7	The current event is about the right mix of presentations, training and networking. Perhaps including more students would beneficial. Maybe give the students an opportunity to present posters	
8	Encourage NSF participation at the large facilities program level.	
9	Please allow for more time for participants to meet the main presenters. For example, it would have been nice to confer with the Homeland Security representative.	
10	More better case.	
11	As a newcomer, I was hoping to gain more information on the actual threats our projects/institutions are subjected to. I was looking for raw data, statistics, what kinds of things are really happening. This was hinted at in a couple of presentations. I think some kind of "status" of threats/attacks would be very useful.	
12	Do not invite someone like the speaker from Intel. Too much cheerleading, too little information.	
13	More sharing by NSF CI projects about what is (and isn't) working for them, what their top risks/concerns are, what their future plans are. The HUBzero presentation was an excellent example of what we need more of in future summits.	
14	Consider how to make it more engaging for the audience. Fewer presentations and panels and more active audience participation, possibly to include breakouts or brainstorming about problems and solutions.	
15	Have the panel discussions be made more interactive. It seemed that 80% of the panel time was spent on the panelists talking. Have some hands on exercises for the audience to get more involved and also learn more take-aways from the summit.	
16	Thought the start was too early for those living in the DC area and those coming from the West Coast. The training sessions the first day could have been all day, going 9-3 with an icebreaker at 5-7 to promote the most interaction.	
17	Publish the presentations from the NSF folks. Publish the list of attendees and make available at the conference. Show groupings by state, orgs, projects, etc.. in some manner that allows us to report on who was at the conference.	
18	It might be helpful/useful to have a session that take several different cybersecurity programs, and deconstruct them to understand the decision points and the way different projects decide upon which practices will best suit their mission.	

19	Panels were too big, some panelists rambled a bit more than one would expect and got off-topic at times. Not enough time for questions, as a result. Moderators perhaps need to have a little coaching on time management prior to session; I noted attempts by another panelist to shut down excessive talker at one session were somewhat heavy-handed and not conducive to openness/sharing. Panelists who expressed opinions, including unconventional ideas, based on their experience, helped illustrate some of the difficult cybersecurity management issues. It seemed that many participants took away useful ideas and the large facilities could benefit as a result.
20	Better snacks at the break and breakfast but other than that an awesome summit with great speakers and topics. Very useful!
21	I would consider less keynote presentations from people outside of the community. One would be good to bring in another perspective but too many is a distraction, and frankly they did not contribute much.
22	Broad spectrum. Need to ensure we get PIs and not just "techies" attending. More facility (NEON, NEES, etc) would be good.
23	Something interactive / mobile, potentially.

Were there any aspects of the summit you found particularly useful or important? If so, please explain.

#	Responses	Date
1	I very much enjoyed the panels. The topics were good, the panelists were knowledgeable and they were well moderated. The talks were a mixed bag---some were very good and some were mediocre. I would like to see a discussion on how the various cybe-rinfrastructures can learn from one another/help one another. Peer reviews of one other was mentioned, which I think is a good idea. Perhaps some discussion on if this is indeed a good idea, is it feasible, how would be the objectives of such a review, what the process would be, how it might be funded, etc.	
2	The 'Developing Cybersecurity Programs for NSF Projects' training was what I got the most value from for this year's summit. It went through all of the steps needed for an organization of my size and provided very useful templates.	
3	Learning that HIPAA regulations aren't all that bad was very informative.	
4	It is necessary at this stage to keep the NSF involved, so presence of several program officers, past and present, should be maintained. To arrive at "industry practices" it would be useful to have BoF sessions or similar to discuss common approaches; the lunch on Day 2 could have been useful in this regard but needed tighter-knit, more focused conversations.	
5	The networking opportunities were great. The structure of the summit was good but there could have been a little more time for the networking.	
6	The unscripted discussions, especially where there was an interesting disagreement between, say, panelists and audience, was useful. There is no one way to address cyber security and a comparison of different approaches can only help -- it sorts out which approach may be more useful for a particular organization.	
7	The high level industry expert presentations (DHS, Intel) were particularly moving and applicable to cybersecurity initiatives within my organization. It was most awesome that all of the digital content from the three days was made available for download.	
8	The panel discussions were engaging and useful.	
9	Bring together a mix of specialists and those with responsibilities for creating and maintain cybersecurity at facilities was useful and prompted dialog and useful exchanges.	
10	I really enjoyed the information provided by Dr. Phyllis Schneck; although not directly related to NSF funded projects, the information and effort provided by the Federal Government is critical collateral projects such as those funded by NSF. The workshop day is an ideal way to refresh and pick up new and relevant information. Great job!	
11	General awareness of the shared risk and opportunities for normalizing the policy to practice to response cycle even for programs with a short life-cycle. 50% of the battle is won by just holding the annual event to allow for networking! Keep up the great work.	
12	Sybersecurity policy for facility	
13	I thought the presentation by the Intel guy was good at giving an overview of how to address cyber security. Information was transformed. I found the talk by the under secretary from the DHS to be largely rambling.	
14	Comments by NSF program officers were very helpful.	
15	The most valuable session for me was, "Developing Cyber Security Programs For NSF Projects."	
16	Panels seemed to be pretty effective. Need to work on the mic system though.	
17	The discussion around REN-ISAC information sharing restrictions was particularly important and in my opinion addressing this barrier to information sharing should be our #1 priority as a community.	

18	Hearing from the community about actual science projects and their cybersecurity implications was very helpful. Hearing from NSF (briefly) about program directions was useful. Hearing from other parts of the government (DHS) about priorities was helpful. Wrap-Up Discussions at end of first day and second morning were useful.
19	Great job overall. Keep doing this.
20	Rubbing shoulders and hearing from top-notch folks with lots of experience, sometimes with differing ideas
21	Lunch table discussions were great.
22	In the session Thursday morning, the moderator noted that availability appears to be high priority for research centers, and the panel responded with what, in their experience, was important and to which user groups. This turned the focus from types of security issues to research priorities and risks to those priorities, showing how the cyber managers determine where to place their resources. I would have liked to hear more about that, and how their plans dovetail (or not) with the facility's overall strategy and performance.
23	Panels were great. Much better than just hearing talks. I'd like to hear more about specific challenges, not so much on success stories. But if there are success stories, what obstacles did they have to overcome in terms of admin and technical aspects.
24	The talk about HubZero that explained what did NOT work was quite interesting and useful since it validated some of my own experiences.
25	Rosenquist didn't understand his audience. His talk was way too basic.
26	Overall, I found the discussions informative and important. The training was very insightful, and appreciated that aspect as well. The greatest aspect, however, was the ability to network and meet others in similar fields.

Appendix I

Tutorial Evaluation survey summary report

Q1 Which options best describe your job or position? Check all that apply.

Answered: 35 Skipped: 0

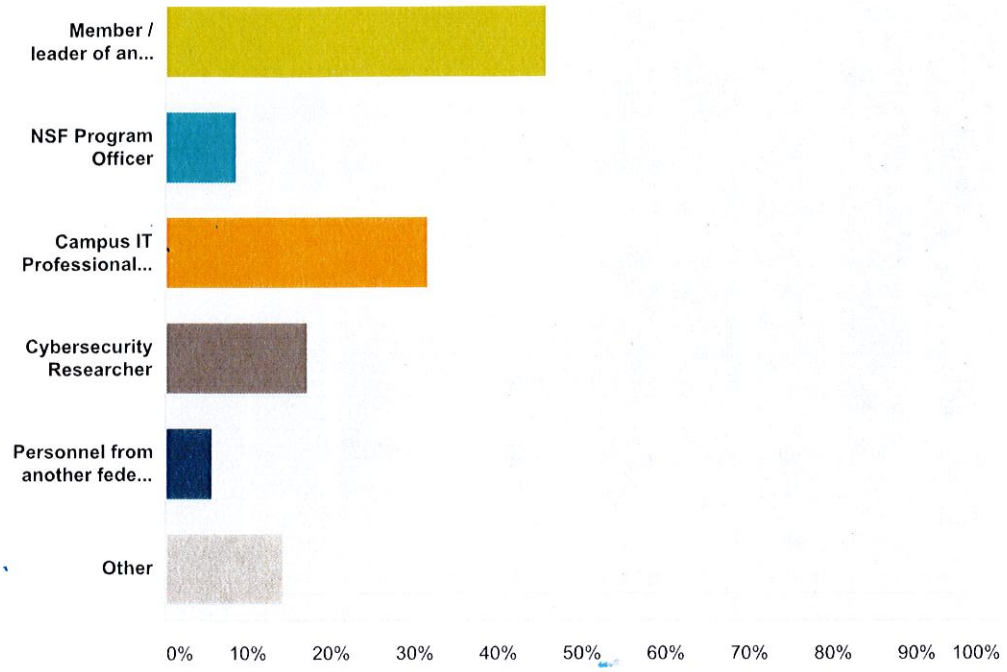

Answer Choices	Responses	
Member / leader of an NSF project	45.71%	16
NSF Program Officer	8.57%	3
Campus IT Professional / CIO	31.43%	11
Cybersecurity Researcher	17.14%	6
Personnel from another federal program (NSA, DOE/ESNet, etc.)	5.71%	2
Other	14.29%	5

Total Respondents: 35

#	If applicable, please state your NSF Project and/or affiliated NSF Directorate. Other comments or clarifications are welcome.	Date
1	NSF CC-NIE grant "OneOklahoma Friction Free Network" NSF grant # ACI-1341028	
2	EarthCube - Geosciences Directorate	
3	GEO/OCE/Ocean Drilling Program	
4	CISE - CNS, CCF ENG - CMMI	
5	ACI	
6	IceCube	
7	I'm the IT Manager for AURA, who has CAs with the NSF to manage Gemini, NOAO, NSO and LSST.	

8	LTER Network - DEB
9	FFRDC with ~500 staff funded out of NSF MPS
10	National Optical Astronomy Observatory.
11	Engineering
12	Cornell Laboratory for Accelerator-based ScienceS and Education Cornell High-Energy Synchrotron Source
13	ENG/CMMI
14	International Ocean Discovery Program
15	CTSC

Q2 How would you characterize your job in relationship to cybersecurity? Please check all that apply.

Answered: 35 Skipped: 0

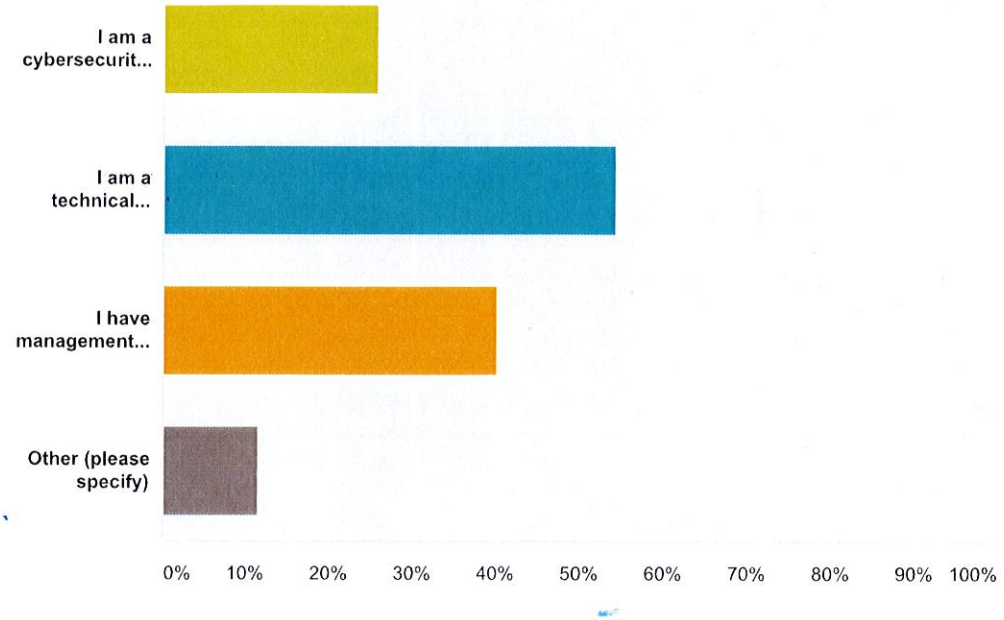

Answer Choices	Responses	
I am a cybersecurity professional	25.71%	9
I am a technical professional who has knowledge of cybersecurity	54.29%	19
I have management responsibility for cybersecurity	40.00%	14
Other (please specify)	11.43%	4

Total Respondents: 35

#	Other (please specify)	Date
1	Cybersecurity researcher	
2	Student	
3	Potential review of cybersecurity	
4	Interested in cybersecurity.. Some knowledge of ICS.	

Q3 Based on your overall experience with the August 26 training sessions, would you participate in training offered at future summits?

Answered: 35 Skipped: 0

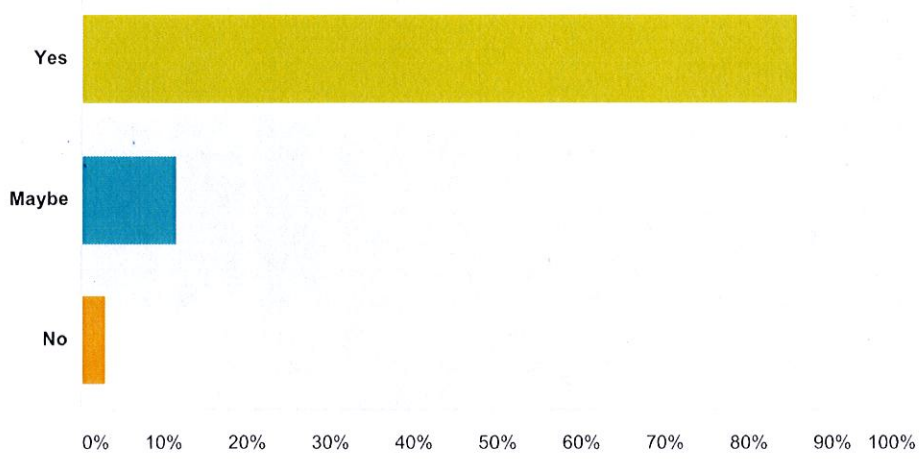

Answer Choices	Responses	
Yes	85.71%	30
Maybe	11.43%	4
No	2.86%	1
Total		35

What training topics would you like to see covered at future summits?

#	Responses	Date
1	ISO & FISMA; The good, the bad, and the ugly.	
2	A session dedicated to cybersecurity at the network level including an example of the physical/logical topology and the various hardware and software network devices that aid in securing the network at each layer.	
3	More Bro Training at different levels of expertise	
4	Password security	
5	Archive and its relationship to Cybersecurity	
6	Keep the secure programming practice tutorial, but make it a full day one	
7	Security minded system architecture - can you build in security?	
8	Reverse engineering Advanced secure coding	
9	I would be interested in ITAR training as it relates to Cybersecurity.	
10	Additional topics on secure coding practices and planning for cyber-security	
11	More technical/operational topics. Really liked the idea of full day of Bro. Maybe expand this to other tools/platforms other institutions use, or are used by the R&E environments. BoFs where organizations share their "special sauce".	
12	Focus on implementation and customization of IDS.	
13	Perhaps a briefing / walkthrough of a conceptually "secure" network. The idea would be that small teams and/or businesses could be more informed on simple and more cost-effective ways to maintain cybersecurity.	
14	More Bro. Network-specific. Science DMZ, etc.	
15	Practical Cybersecurity for "smallish" FFRDC's embedded in a much larger campus network How to "safely" run a data archive	
16	1. security toolkit for researchers 2. risk assessment and management	
17	Cybersecurity Tools	
18	A continuation of the secure coding practices tutorial.	
19	Tools for managing security	
20	1. Cybersecurity strategies in software defined datacenter environments. 2. Security policy deep dive.	
21	Additional security platforms, technologies, etc.	
22	Cybersecurity for ICS. More participatory training.	
23	Policy	
24	Case studies or reviews of Cybersecurity Programs and implementations.	

Q5 Which morning session did you attend?

Answered: 35 Skipped: 0

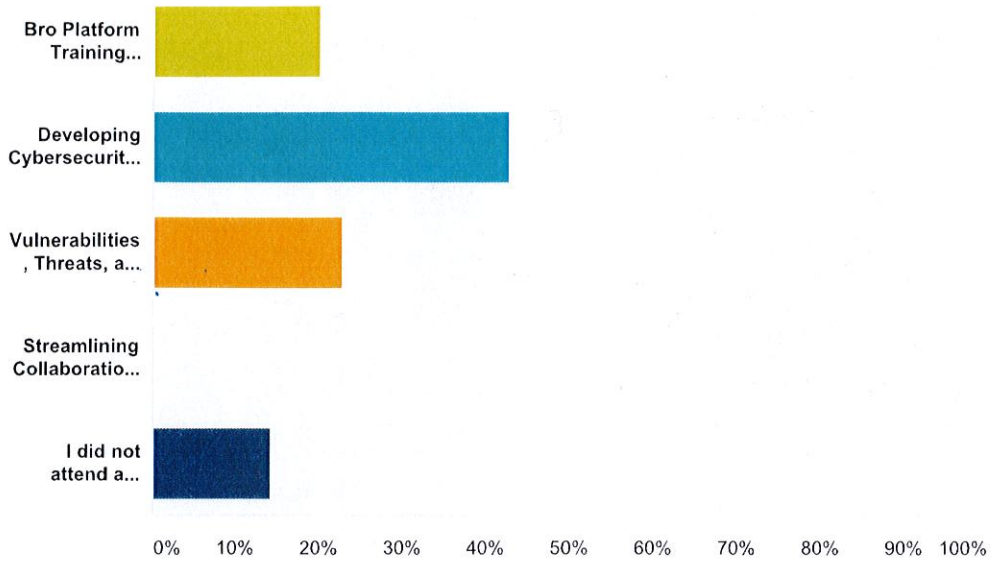

Answer Choices	Responses	
Bro Platform Training Workshop	20.00%	7
Developing Cybersecurity Programs for NSF Projects	42.86%	15
Vulnerabilities, Threats, and Secure Coding Practices	22.86%	8
Streamlining Collaboration with InCommon and Identity Federations: Focus on Supporting International Collaboration	0.00%	0
I did not attend a morning session	14.29%	5
Total		**35**

Q6 How would you rate your level of pre-training familiarity with the topics covered by this morning training session?

Answered: 32 Skipped: 3

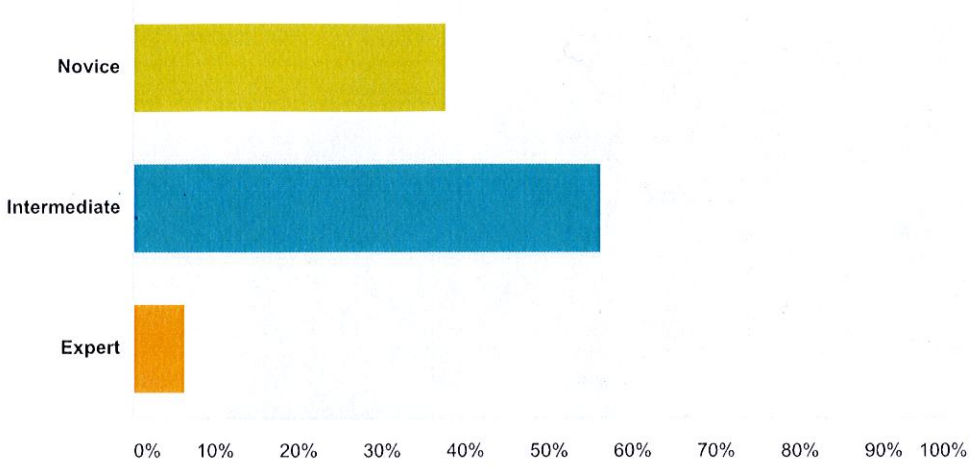

Answer Choices	Responses	
Novice	37.50%	12
Intermediate	56.25%	18
Expert	6.25%	2
Total		32

Q7 How would you rate your overall experience with the morning training?

Answered: 32 Skipped: 3

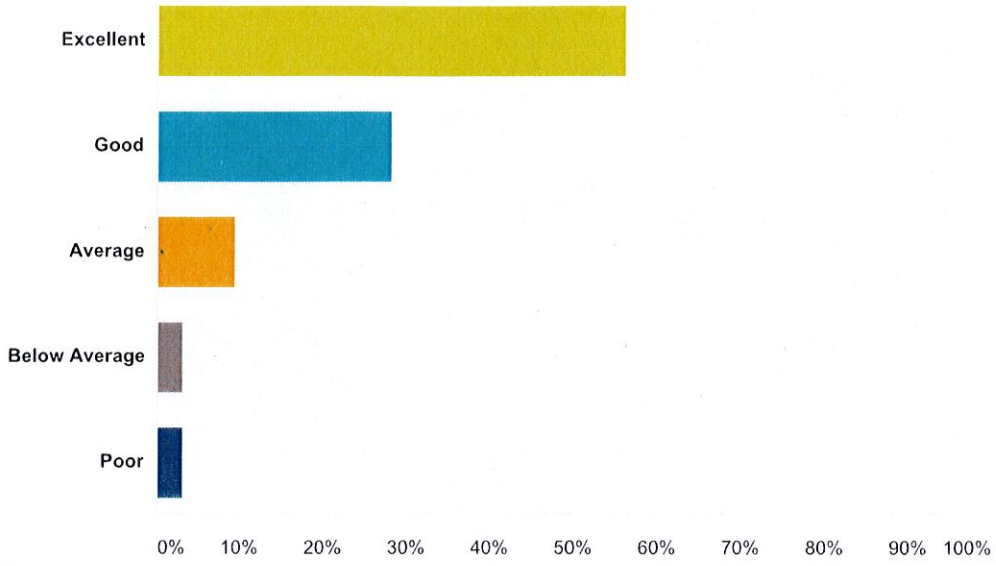

Answer Choices	Responses	
Excellent	56.25%	18
Good	28.13%	9
Average	9.38%	3
Below Average	3.13%	1
Poor	3.13%	1
Total		32

Q8 Please rate your experience with the morning training in these areas:

Answered: 32 Skipped: 3

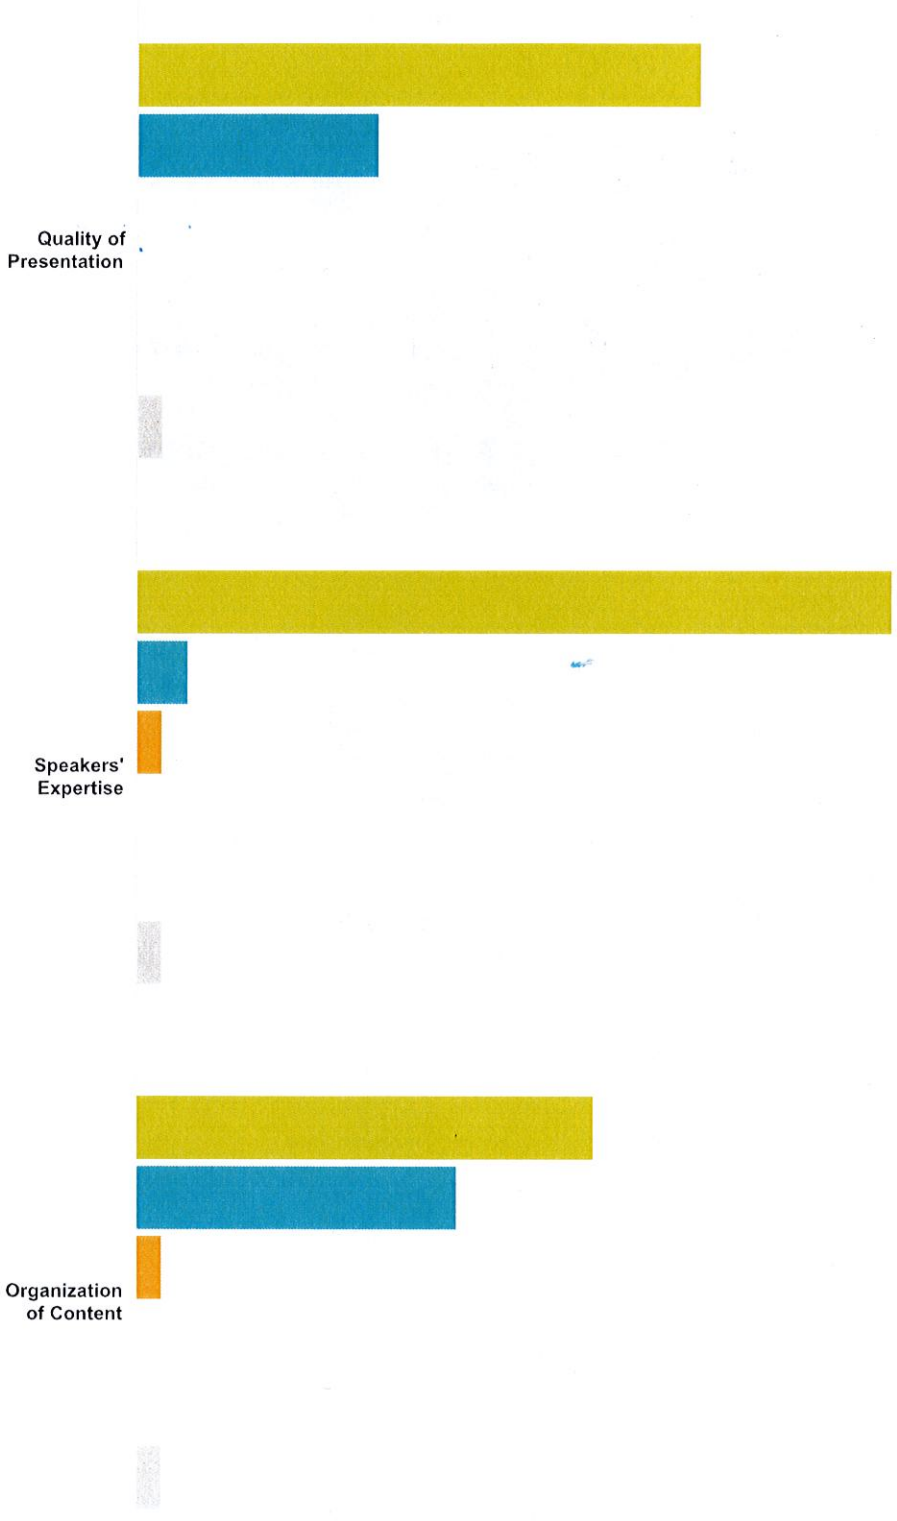

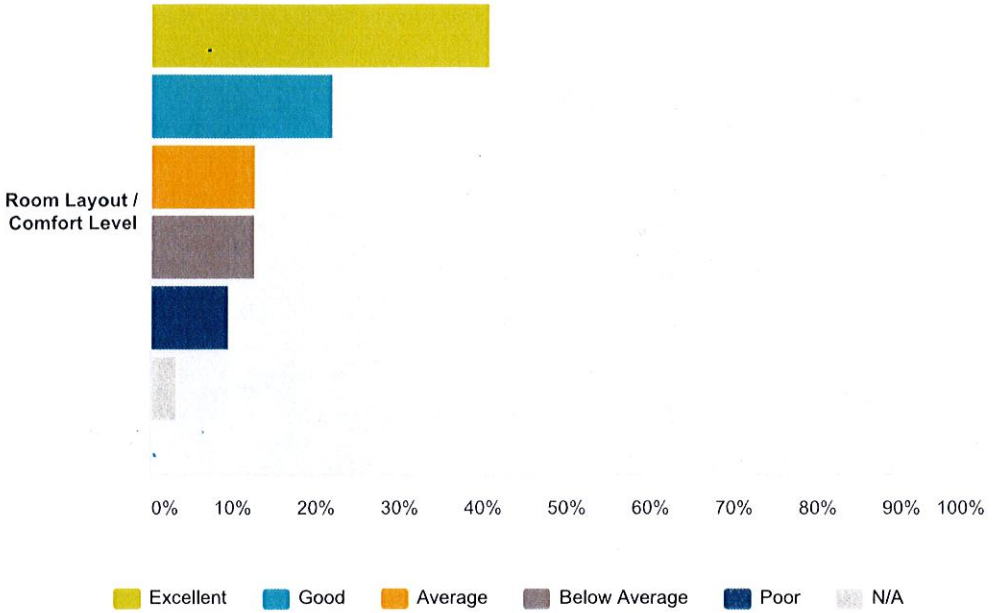

	Excellent	Good	Average	Below Average	Poor	N/A	Total Respondents
Quality of Presentation	67.74%	29.03%	0.00%	0.00%	0.00%	3.23%	
	21	9	0	0	0	1	31
Speakers' Expertise	90.63%	6.25%	3.13%	0.00%	0.00%	3.13%	
	29	2	1	0	0	1	32

Q9 Was this morning training better than what you expected, worse than what you expected, or about what you expected?

Answered: 32 Skipped: 3

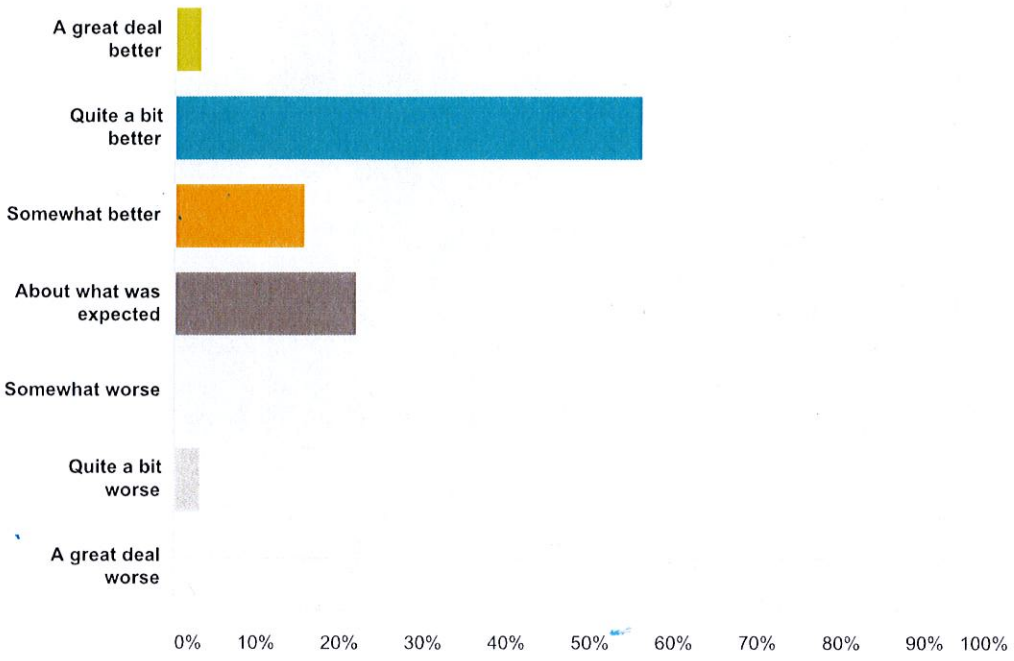

Answer Choices	Responses	
A great deal better	3.13%	1
Quite a bit better	56.25%	18
Somewhat better	15.63%	5
About what was expected	21.88%	7
Somewhat worse	0.00%	0
Quite a bit worse	3.13%	1
A great deal worse	0.00%	0
Total		32

Q10 How useful to your work was this morning training?

Answered: 32 Skipped: 3

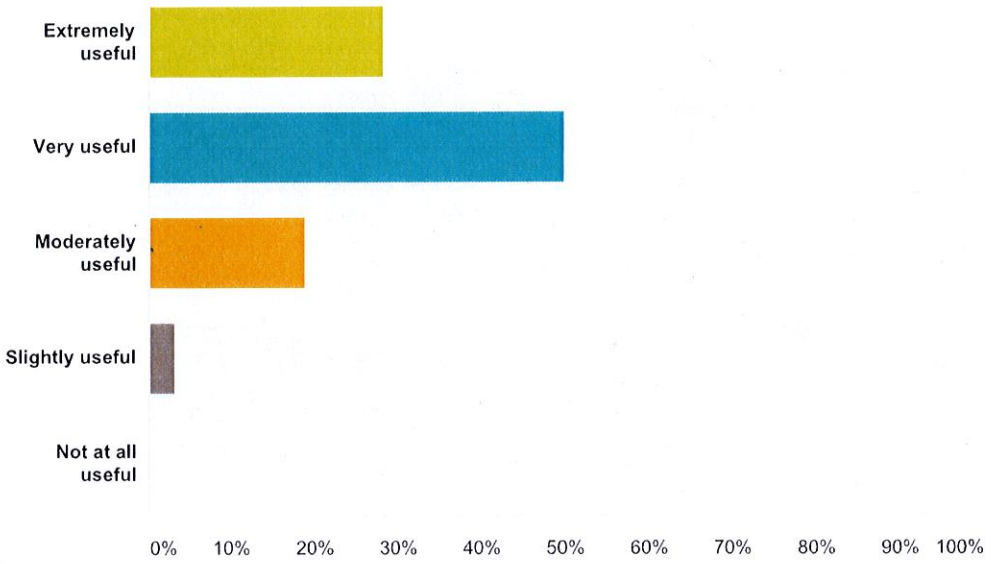

Answer Choices	Responses	
Extremely useful	28.13%	9
Very useful	50.00%	16
Moderately useful	18.75%	6
Slightly useful	3.13%	1
Not at all useful	0.00%	0
Total		32

Responses to Question 11 (How can we improve this training session in the future?) and Question 12 (Were there any aspects of the morning training you found particularly useful or important? Please explain) are open-ended responses directed at specific training sessions. They have been provided to the respective training teams, and are removed from this appendix.

Q13 Which afternoon session did you attend?

Answered: 35 Skipped: 0

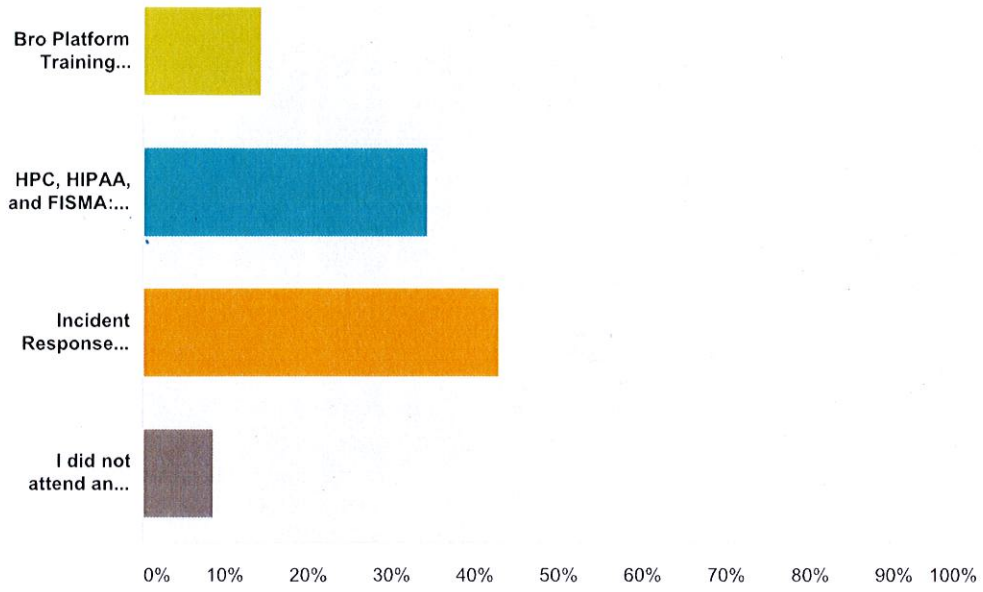

Answer Choices	Responses	
Bro Platform Training Workshop	14.29%	5
HPC, HIPAA, and FISMA: Meeting the Regulatory Challenge through Effective Risk Management	34.29%	12
Incident Response Training	42.86%	15
I did not attend an afternoon session	8.57%	3
Total		**35**

Q14 How would you rate your level of pre-training familiarity with the topics covered by this afternoon training session?

Answered: 32 Skipped: 3

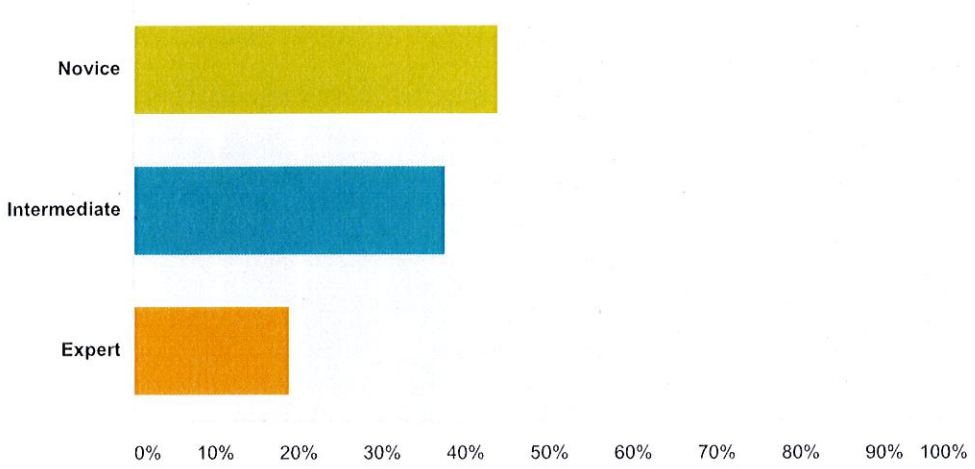

Answer Choices	Responses	
Novice	43.75%	14
Intermediate	37.50%	12
Expert	18.75%	6
Total		32

Q15 How would you rate your overall experience with the afternoon training?

Answered: 30 Skipped: 5

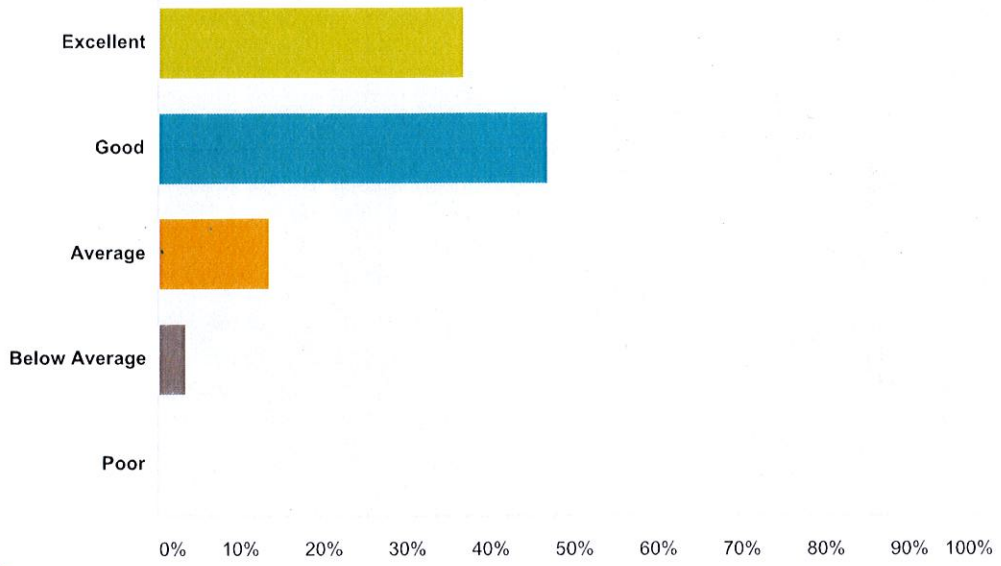

Answer Choices	Responses	
Excellent	36.67%	11
Good	46.67%	14
Average	13.33%	4
Below Average	3.33%	1
Poor	0.00%	0
Total		30

Q16 Please rate your experience with the afternoon training in these areas:

Answered: 31 Skipped: 4

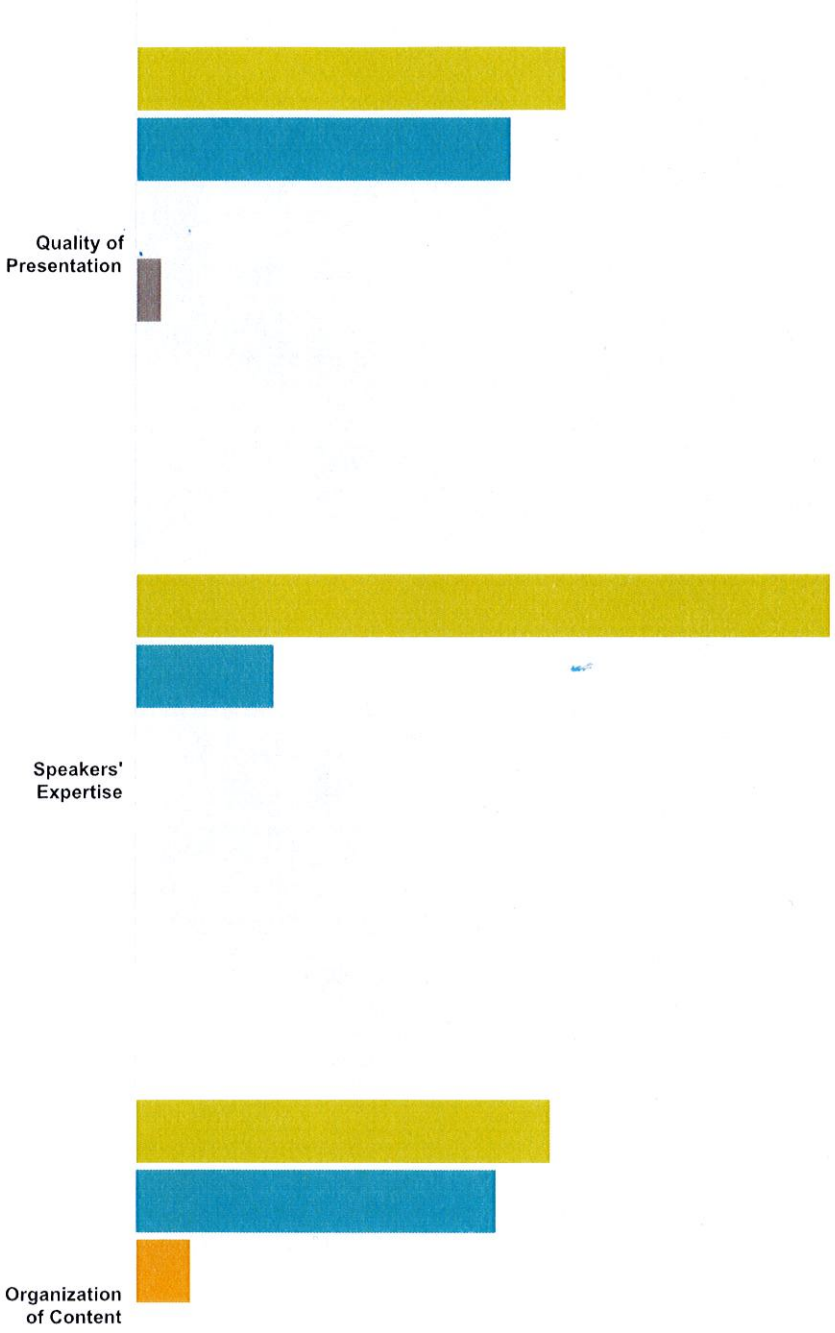

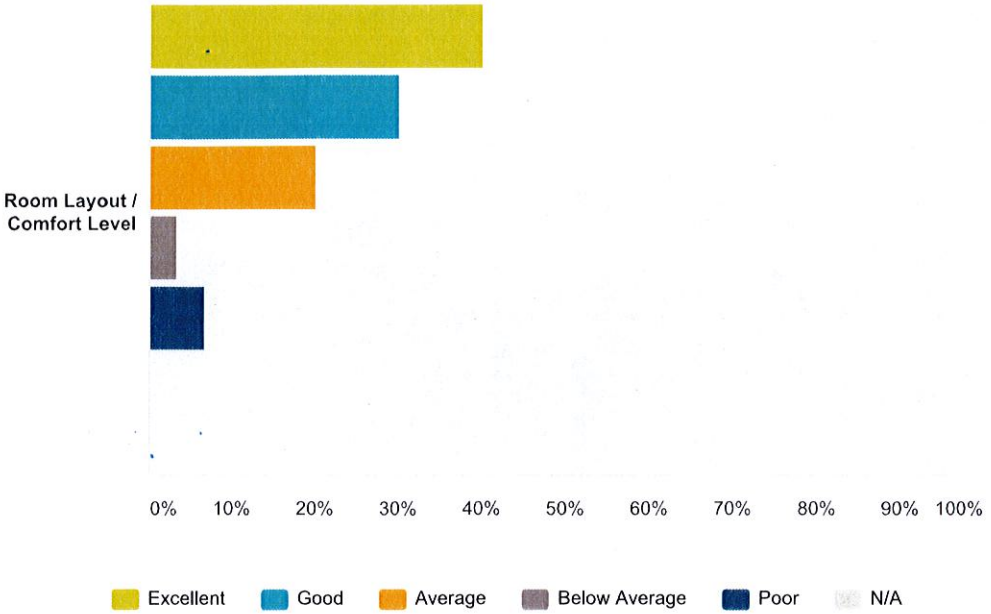

	Excellent	Good	Average	Below Average	Poor	N/A	Total Respondents
Quality of Presentation	51.61%	45.16%	0.00%	3.23%	0.00%	0.00%	
	16	14	0	1	0	0	31
Speakers' Expertise	83.33%	16.67%	0.00%	0.00%	0.00%	0.00%	
	25	5	0	0	0	0	30

Q17 Was this afternoon training session better than what you expected, worse than what you expected, or about what you expected?

Answered: 31 Skipped: 4

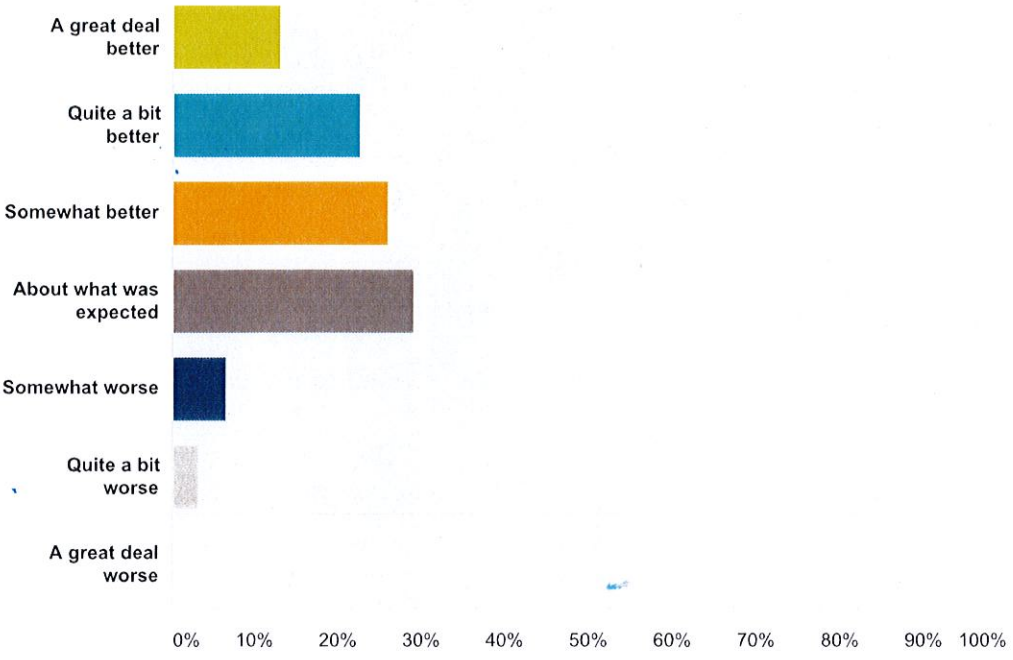

Answer Choices	Responses	
A great deal better	12.90%	4
Quite a bit better	22.58%	7
Somewhat better	25.81%	8
About what was expected	29.03%	9
Somewhat worse	6.45%	2
Quite a bit worse	3.23%	1
A great deal worse	0.00%	0
Total		**31**

Q18 How useful to your work was this afternoon training?

Answered: 31 Skipped: 4

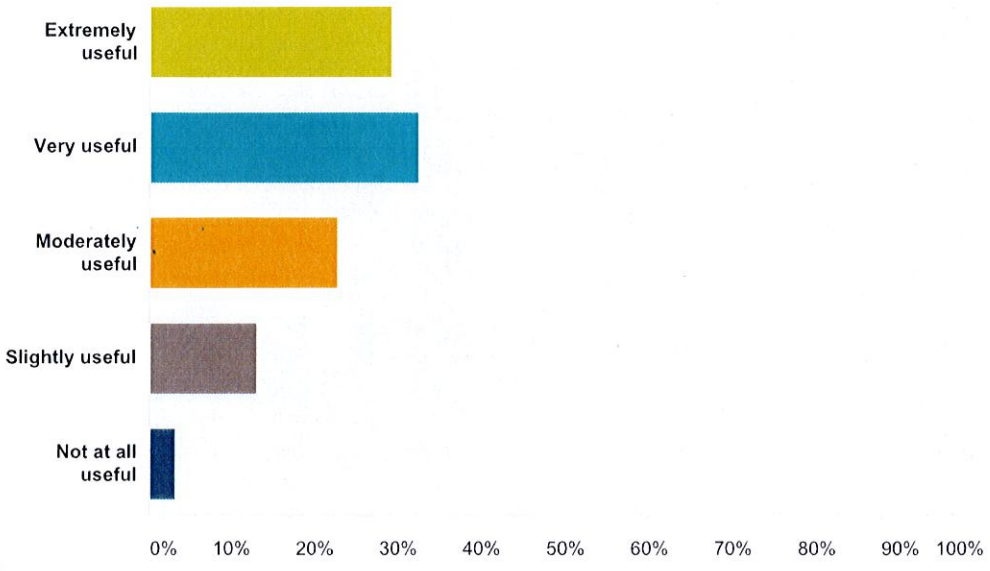

Answer Choices	Responses	
Extremely useful	29.03%	9
Very useful	32.26%	10
Moderately useful	22.58%	7
Slightly useful	12.90%	4
Not at all useful	3.23%	1
Total		**31**

Responses to Question 19 (How can we improve this training session in the future?) and Question 20 (Were there any aspects of the afternoon training you found particularly useful or important? Please explain) are open-ended responses directed at specific training sessions. They have been provided to the respective training teams, and are removed from this appendix.