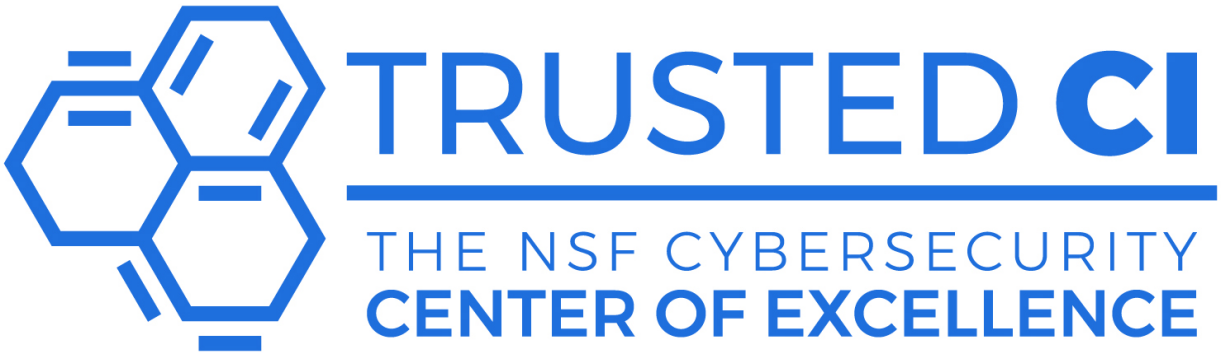

SCiMMA / Trusted CI Engagement

Final Report

December 1, 2020
Distribution: Public

Andrew Adams[1], Shane Filus

[1] Engagement Lead, akadams@psc.edu

About Trusted CI

The mission of Trusted CI is to provide the NSF community with a coherent understanding of cybersecurity, its importance to computational science, and what is needed to achieve and maintain an appropriate cybersecurity program.

This document is a product of Trusted CI. Trusted CI is supported by the National Science Foundation under Grant #1920430. For more information about Trusted CI, please visit: http://trustedci.org/. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Using & Citing this Work

Table of Contents

1 Background

The Scalable Cyberinfrastructure for Multi-Messenger Astrophysics (SCiMMA) project is a planned collaboration between data scientists, computer scientists, astronomers, astro-particle physicists, and gravitational wave physicists[2]. Leveraging NSF investments in astronomical and multi-messenger facilities, and in advanced cyberinfrastructure, SCiMMA intends to prototype a publish-subscribe system based on KAFKA to distribute alerts from gravitational wave, neutrino and electromagnetic observatories to authorized subscribers (initially, public alerts so that all subscribers are authorized, but eventually proprietary alerts). The system will additionally rely on supporting infrastructure, including: machine learning algorithms to analyze and classify alerts; and event databases for richer data mining. The pub/sub prototype will be hosted on cloud resources, including a commercial cloud (e.g., AWS). Upon award completion, SCiMMA will request funding for a sustained distributed institute that will expand the scope and depth of the prototyped system.

To this end, a group from SCiMMA solicited information security guidance from Trusted CI on and-or with various components of their prototype cyberinfrastructure. For example, they sought help in developing an IT security program, identifying appropriate security control sets/catalogs, and performing a risk assessment with a corresponding residual risk registry.

The SCiMMA team consisted of Warren Anderson (lead), Donald Petravick, Margaret Johnson, and Vladislav Ekimtcov, while the Trusted CI personnel were Andrew Adams (lead) and Shane Filus. All contributed .1 of an FTE to the engagement.

The engagement ran from July 1, 2020 to December 31, 2020.

[2] https://scimma.org

2 Engagement Process

After an initial meeting and post-discussions, Trusted CI and the SCiMMA team refined and prioritized SCiMMA's needs by identifying a subset of tasks. These tasks were then used to outline the goals of the engagement, specifically:

1. Perform a security review of SCiMMA's cyberinfrastructure using the Trusted CI Security Program Evaluation worksheet[3] in order to assess the target level of cybersecurity needed.
2. Using information documented in step 1., develop the start of a security program leveraging a master information security policies and procedures document, and
3. Document assets to be used by the security program in step 2.

Additionally, three extra tasks were chosen to pursue, if time allowed, including:

1. Identify an initial set of security controls
2. Identify and develop missing, necessary, security policies
3. Perform a risk assessment and develop a residual risk registry for management

With the goals agreed upon, the engagement started with the SCiMMA team completing the Trusted CI Security Program Evaluation spreadsheet. This was a productive exercise in that it got the team discussing the cybersecurity concerns broached in the evaluation.

From there, the SCiMMA team deemed that having data to show to stakeholders that captured the cyberinfrastructure risk, and thus, the need for security resources was of high priority to the team. So the engagement decided to tackle the 'extra' task of populating the asset-based risk assessment spreadsheet. However, it soon became apparent that SCiMMA had (i) an excessive number of assets, and (ii) their cyberinfrastructure was still in flux. Thus, we shifted our efforts then to the goal of documenting assets, and of those, only critical assets, e.g., admin credentials, source code, DLP backups, etc.

Focusing on just critical assets, however, also proved challenging as the design of the system had yet to harden. Hence, the SCiMMA team attempted to identify what assets would be needed if, e.g., the cloud service was changed from AWS to something else. That is, they had to think about components and critical assets in a more abstract manner.

[3] https://docs.google.com/document/d/1gEMUZLQ6O-RA0yjlv9MYIF_90C3YpUm1LE0ttt9e8O0/

In parallel to this, the SCiMMA team, after attending 'The Framework' workshop at the NSF Cybersecurity Summit[4], sought to revamp their nascent Master Information Security Practices & Procedures document to adopt much of the ideas promoted during the workshop, including leveraging the CIS Controls Tracking tool[5] as their base-line control set. Thus, in conjunction with working on the asset inventory, quality effort was also spent in understanding what controls comprised (at least) implementation group 1 from the CIS control set, and how they would be applied to SCiMMA.

3 Observations & Recommendations

SCiMMA's desire to both identify a control set/catalog for their cyberinfrastructure and then try to understand the residual risk that would still be present after implementing the control set immediately sets them apart from most science gateway-like services. Similarly, they understand the need for a cybersecurity budget and dedicated personnel; both of which are key components of a sound security program.

Other takeaways the SCiMMA team leveraged from the engagement include: mapping CIS 7.1 controls to AWS cloud security controls (which use CIS 6.1); AWS security primers; John Gilligan's "threat sophistication X mission criticality" asset matrix[6]; and an introduction to Trusted CI's Asset-Specific Access and Privilege Specification (ASAPS) documents[7].

Finally, the engagement, for better or worse, highlighted the amount of work in design, documentation, and implementation that will be needed to achieve a strong security posture in a system as complex as SCiMMA; the SCiMMA team is now aware of the resource burden and needs to make certain that this is conveyed to leadership/stakeholders.

Our primary recommendation, once the design of SCiMMA solidifies, is to (i) fully complete documenting all critical assets, (ii) generate ASAPS for all necessary assets identified in (i), (iii) apply the CIS control tracker to those assets, and (iv) identify what residual risk still needs to be conveyed to stakeholders.

A secondary goal would be to complete SCiMMA's security program leveraging the Trusted CI Framework once it is published in early 2021.

[4] https://www.trustedci.org/2020-nsf-summit
[5] https://docs.google.com/spreadsheets/d/1XozhP8QY9mdm1nQyY5YOC26SexgDCqaakBYS4KVcEBg/
[6] https://www.afcea.org/committees/cyber/documents/cybereconfinal.pdf
[7] https://docs.google.com/document/d/1AQCDERMg9YKZk54RqfaXTNYr1RZ1Ms6J73ln7Xdzex4/

4 Future Engagement Possibilities

Like many science gateways that rely on many services (e.g., cloud services, CMS, IAM, home-grown source code, containers), once the system is near production, a future engagement with Trusted CI could: review the security program in place; focus on one particular component, e.g., IAM; and-or leverage a code review. All of these seem beneficial to SCiMMA.