
Stream: Independent Submission
RFC: [8953](#)
Category: Informational
Published: December 2020
ISSN: 2070-1721
Author: K. Moriarty
Center for Internet Security

RFC 8953

Coordinating Attack Response at Internet Scale 2 (CARIS2) Workshop Report

Abstract

The Coordinating Attack Response at Internet Scale (CARIS) 2 workshop, sponsored by the Internet Society, took place on 28 February and 1 March 2019 in Cambridge, Massachusetts, USA. Participants spanned regional, national, international, and enterprise Computer Security Incident Response Teams (CSIRTs), operators, service providers, network and security operators, transport operators and researchers, incident response researchers, vendors, and participants from standards communities. This workshop continued the work started at the first CARIS workshop, with a focus on scaling incident prevention and detection as the Internet industry moves to a stronger and a more ubiquitous deployment of session encryption.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8953>.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction
2. Accepted Papers
3. CARIS2 Goals
4. Workshop Collaboration
 - 4.1. Breakout 1 Results: Standardization and Adoption
 - 4.1.1. Wide Adoption
 - 4.1.2. Limited Adoption
 - 4.2. Breakout 2 Results: Preventative Protocols and Scaling Defense
 - 4.3. Breakout 3 Results: Incident Response Coordination
 - 4.4. Breakout 4 Results: Monitoring and Measurement
 - 4.4.1. IP Address Reputation
 - 4.4.2. Server Name Authentication Reputation C (SNARC)
 - 4.4.3. Logging
 - 4.4.4. Fingerprinting
 - 4.5. Taxonomy and Gaps Session
5. Next Steps
6. Summary
7. Security Considerations
8. IANA Considerations
9. References
 - 9.1. Informative References

Acknowledgements

Author's Address

1. Introduction

The Coordinating Attack Response at Internet Scale (CARIS) 2 workshop [[CARISEvent](#)], sponsored by the Internet Society, took place on 28 February and 1 March 2019 in Cambridge, Massachusetts, USA. Participants spanned regional, national, international, and enterprise Computer Security Incident Response Teams (CSIRTs), operators, service providers, network and security operators, transport operators and researchers, incident response researchers, vendors, and participants from standards communities. This workshop continued the work started at the first CARIS workshop [[RFC8073](#)], with a focus on scaling incident prevention and detection as the Internet industry moves to a stronger and a more ubiquitous deployment of session encryption. Considering the related initiative to form a research group (Stopping Malware and Researching Threats [[SMART](#)]) in the Internet Research Task Force (IRTF), the focus on prevention included consideration of research opportunities to improve protocols and determine if there are ways to improve attack detection during the protocol design phase that could later influence protocol development in the IETF. This is one way to think about scaling response, through prevention and allowing for new methods to evolve for detection in a post-encrypted world. Although the proposed SMART Research Group has not yet progressed, the work to better scale incident response continues through the projects proposed at CARIS2 as well as in future CARIS workshops.

2. Accepted Papers

Researchers from around the world submitted position and research papers summarizing key aspects of their work to help form the shared content of the workshop. The accepted papers may be found at [[CARISEvent](#)] and include:

- Visualizing Security Automation: Takeshi Takahashi, NICT, Japan
- Automating Severity Determination: Hideaki Kanehara, NICT, Japan
- OASIS's OpenC2: Draper and DoD
- Automated IoT Security: Oscar Garcia-Morchon and Thorsten Dahm
- Taxonomies and Gaps: Kirsty P., UK NCSC
- FIRST: Thomas Schreck, Siemens
- NetSecWarriors: Tim April, Akamai
- Measured Approaches to IPv6 Address Anonymization and Identity Association: Dave Plonka and Arthur Berger, Akamai

The program committee worked to fill in the agenda with meaningful and complementary sessions to round out the theme and encourage collaboration to advance research toward the goals of the workshop. These sessions included:

- Manufacturer Usage Description (MUD) [[RFC8520](#)]: Eliot Lear, Cisco
- TF-CSIRT: Mirjam Kühne, RIPE NCC
- M2M Sharing Revolution: Scott Pinkerton, DoE ANL
- Comparing OpenC2 with existing efforts, e.g., [I2NSF](#) [[I2NSF](#)]: Chris Inacio
- Alternate Sharing and Mitigation Models: Kathleen Moriarty, Dell EMC

The presentations provided interesting background to familiarize workshop attendees with current research work, challenges that must be addressed for forward progress, and opportunities to collaborate in the desire to better scale attack response and prevention.

3. CARIS2 Goals

The goal of each CARIS workshop has been to focus on the challenge of improving the overall security posture. The approach has been to identify intrinsic or built-in protection capabilities for improved defense, automation, and scaling attack response through collaboration and improved architectural patterns. It has been assumed that additional training will likely not address the lack of information security professionals to fill the job gap. Currently, there is approximately a [three-million-person deficit](#) [[deficit](#)] for security professionals worldwide, and that is only expected to grow. In preparing for the workshop, the chair and program committee considered that this gap cannot be filled through training but requires measures to reduce the number of information security professionals needed through new architectures and research toward attack prevention. CARIS2 was specifically focused on the industry shift toward the increased use of stronger session encryption ([TLS 1.3](#) [[RFC8446](#)], [QUIC](#) [[QUIC](#)], [tcpcrypt](#) [[RFC8548](#)], etc.) and how prevention and detection can advance in this new paradigm. As such, the goals for this workshop included:

- Scale attack response, including ways to improve prevention, as the Internet shifts to use of stronger and more ubiquitous encryption.
 - Determine research opportunities
 - Consider methods to improve protocols and provide guidance toward goal. For instance, are there ways to build detection of threats into protocols, since they cannot be monitored on the wire in the future?
- Identify promising research ideas to seed a research agenda to input to the proposed IRTF SMART Research Group.

4. Workshop Collaboration

Both CARIS workshops brought together a set of individuals who had not previously collaborated toward the goals of scaling attack response. This is important as the participants span various areas of Internet technology work, conduct research, provide a global perspective, have access to varying data sets and infrastructure, and are influential in their area of expertise. The specific goals, contributions, and participants of the CARIS2 workshop were all considered in the design of the breakout sessions to both identify and advance research through collaboration. The breakout sessions varied in format to keep attendees engaged and collaborating; some involved the full set of attendees while others utilized groups.

The workshop focused on identifying potential areas for collaboration and advancing research.

1. Standardization and Adoption: identify widely adopted and pervasive standard protocols and data formats as well as those that failed.
2. Preventative Protocols and Scaling Defense: identify protocols to address automation at scale.
3. Incident Response Coordination: brainstorm what potential areas of research or future workshops could be held to improve on the scalability of incident response.
4. Monitoring and Measurement: brainstorm methods to perform monitoring and measurement with the heightened need and requirement to address privacy.
5. Taxonomy and Gaps: brainstorm a way forward for the proposed SMART Research Group.

4.1. Breakout 1 Results: Standardization and Adoption

This breakout session considered points raised in the preceding talks on hurdles for automating security controls, detection, and response; the teams presenting noted several challenges they still face today. The breakout session worked toward identifying standard protocols and data formats that succeeded in achieving adoption as well as several that failed or only achieved limited adoption. The results from the evaluation were interesting and could aid in achieving greater adoption when new work areas are developed. The following subsections detail the results.

4.1.1. Wide Adoption

The Transport Layer Security (TLS) protocol has replaced the Secure Sockets Layer (SSL) protocol.

Observations: There was a clear need for session encryption at the transport layer to protect application data. E-commerce was a driving force at the time with a downside to those who did not adopt. Other positive attributes that aided adoption were modular design, clean interfaces, and being first to market.

The Simple Network Management Protocol (SNMP) enables configuration management of devices with extension points for private configuration and management settings. SNMP is widely adopted and is only now, after decades, being replaced by a newer alternative, YANG (a data modeling language) that facilitates configuration management via the Network Configuration Protocol (NETCONF) or RESTCONF. SNMP facilitated an answer to a needed problem set: configuration, telemetry, and network management. Its development considered the connection between the user, vendor, and developers. Challenges did surface for adoption from SNMPv1.1 to 1.2, as there was no compelling reason for adoption. SNMPv3 gained adoption due to its resilience to attacks by providing protection through improved authentication and encryption.

IP Flow Information Export (IPFIX) was identified as achieving wide adoption for several reasons. The low cost of entry, wide vendor support, diverse user base, and wide set of use cases spanning multiple technology areas were some of the key drivers cited.

X.509 was explored for its success in gaining adoption. The solution being abstract from crypto, open, customizable, and extensible were some of the reasons cited for its successful adoption. The team deemed it a good solution to a good problem and observed that government adoption aided its success.

4.1.2. Limited Adoption

Next, each team evaluated solutions that have not enjoyed wide adoption.

Although Structured Threat Information eXpression (STIX) and the Incident Object Description Exchange Format (IODEF) are somewhat similar in their goals, the standards were selected for evaluation by two separate groups with some common findings.

STIX has had limited adoption by the financial sector but no single, definitive end user. The standard is still in development with the US government as the primary developer in partnership with OASIS. There is interest in using STIX to manage content, but users don't really care about what technology is used for the exchange. The initial goals may not wind up matching the end result for STIX, as managing content may be the primary use case.

IODEF was specified by National Research and Education Networks (NRENs) and Computer Security Incident Response Teams (CSIRTs) and formalized in the IETF [[RFC7970](#)]. The user is the security operations center (SOC). While there are several implementations, it is not widely adopted. In terms of exchange, users are more interested in indicators than full event information, and this applies to STIX as well. Sharing and trust are additional hurdles as many are not willing to disclose information.

DNS-Based Authentication of Named Entities (DANE) has DNSSEC as a dependency, which is a hurdle toward adoption (too many dependencies). It has a roll-your-own adoption model, which is risky. While there are some large pockets of adoption, there is still much work to do to gain widespread adoption. A regulatory requirement gave rise to partial adoption in Germany, which naturally resulted in production of documentation written in German -- possibly giving rise to further adoption in German-speaking countries. There has also been progress made in the Netherlands through the creation of a website: <[internet.nl](#)>. The website allows you to test your

website for a number of standards (IPv6, DNSSEC, DANE, etc.). <internet.nl> is a collaboration of industry organizations, companies, and the government in the Netherlands and is available for worldwide use.

IP version 6 (IPv6) has struggled, and the expense of running a dual stack was one of the highest concerns on the list discussed in the workshop breakout. The end user for IPv6 is everyone, and the breakout team considered it too ambiguous. Too many new requirements have been added over its 20-year life. The scope of necessary adoption is large with many peripheral devices. Government requirements for support have helped somewhat with improved interoperability and adoption, but features like NAT being added to IPv4 slowed adoption. With no new features being added to IPv4 and lessons learned, there's still a possibility for success.

4.2. Breakout 2 Results: Preventative Protocols and Scaling Defense

This breakout session followed the sessions on MUD, Protocol for Automated Vulnerability Assessment (PAVA), and Protocol for Automatic Security Configuration (PASC), which have themes of automation at scale. MUD was designed for Internet of Things (IoT), and as such, scaling was a major consideration. The PAVA and PASC work builds off of MUD and maintains some of the same themes. This breakout session was focused on groups brainstorming preventative measures and enabling vendors to deploy mitigations.

One group dove a bit deeper into MUD and layer 2 (L2) discovery. MUD changes sets of filtering control management to the vendor or intermediary MUD vendors for a predictable platform that scales well. While the overall value of MUD is clear, the use of MUD and what traffic is expected for a particular device should be considered sensitive information, as it could be used to exploit a device. MUD has an option of using L2 discovery to share MUD files. L2 discovery, like the Dynamic Host Configuration Protocol (DHCP), is not encrypted from the local client to the DHCP server at this point in time (there is some interest to correct this, but it hasn't received enough support yet). As a result, it is possible to leak information and reveal data about the devices for which the MUD files would be applied. This could multicast out information such as network characteristics, firmware versions, manufacturers, etc. There was some discussion on the use of 802.11 to improve connections [IEEE802.11]. Several participants from this group plan to research this further and identify options to prevent information leakage while achieving the stated goals of MUD.

The next group discussed a proposal one of the participants had already begun developing, namely privacy for rendezvous service. The basic idea was to encrypt Server Name Indication (SNI) using DNS to obtain public keys. The suffix on server IPv6 would be unique to a TLS session (information missing). The discussion on this proposal was fruitful, as the full set of attendees engaged, with special interest from the incident responders to be involved in early review cycles. Incident responders are very interested to understand how protocols will change and to assess the overall impact of changes on privacy and security operations. Even if there are no changes to the protocol proposals stemming from this review, the group discussion landed on this being a valuable exchange to understand early the impacts of changes for incident detection and mitigation, to devise new strategies, and to provide assessments on the impact of protocol changes on security in the round.

The third group reported back on trust exchanges relying heavily on relationships between individuals. They were concerned with scaling the trust model and finding ways to do that better. The group dove deeper into this topic.

The fourth group discussed useful data for incident responders. This built on the first breakout session ([Section 4.1](#)). The group determined that indicators of compromise (IoCs) are what most organizations and groups are able to successfully exchange. Ideally, these would be fixed and programmable. They discussed developing a richer format for sharing event threats. When reporting back to the group, a successful solution used in the EU was mentioned: the [Malware Information Sharing Platform \(MISP\)](#) [MISP]. This will be considered in the review of existing efforts to determine if anything new is needed.

4.3. Breakout 3 Results: Incident Response Coordination

Incident response coordination currently does not scale. This breakout session focused on brainstorming incident response and coordination, looking specifically at what works well for teams today, what is holding them back, and what risks loom ahead. Output from this session could be used to generate research and to dive deeper in a dedicated workshop on these topics.

Supporting:

- Trust between individuals in incident response teams
- Volume of strong signals and automated discovery
- Need to protect network as a forcing function
- Law and legal catalyst, motivator to stay on top
- Current efforts supported by profit and company interests, but those may shift
- Fear initially results in activity or in terms of the diagram used, a burst of wind, but eventually leads to complacency

What creates drag:

- Lack of clear Key Performance Indicators (KPIs)
- Too many standards
- Potential for regional borders to impact data flows
- Ease of use for end users
- Speed to market without security considerations
- Legal framework slow to adapt
- Disconnect in actual/perceived risk
- Regulatory requirements preventing data sharing
- Lack of clarity in shared information
- Behind the problem/reactionary
- Lack of resources/participation
- Monoculture narrows focus

Looming problems:

- Dynamic threat landscape
- Liability
- Vocabulary collision
- Lack of target/adversary clarity
- Bifurcation of Internet
- Government regulation
- Confusion around metrics
- Sensitivity of intelligence (trust)
- Lack of skilled analysts
- Lack of "fraud loss" data sharing
- Stakeholder/leader confusion
- Unknown impact of emerging technologies
- Overcentralization of the Internet
- New technologies and protocols
- Changes in application-layer configurations (e.g., browser resolvers)

4.4. Breakout 4 Results: Monitoring and Measurement

The fourth breakout session followed Dave Plonka's talk on IPv6 aggregation to provide privacy for IPv6 sessions. Essentially, IPv6 provides additional capabilities for monitoring sessions end to end. Dave and his coauthor, Arthur Berger, primarily focus on measurement research but found a way to aggregate sessions to assist with maintaining user privacy. If you can devise methods to perform management and measurement, or even perform security functions, while accommodating methods to protect privacy, a stronger result is likely. This also precludes the need for additional privacy improvement work to defeat measurement objectives.

This breakout session was focused on devising methods to perform monitoring and measurement, coupled with advancing privacy considerations. The full group listed out options for protocols to explore and ranked them, with the four highest then explored by the breakout groups. Groups agreed to work further on the proposed ideas.

4.4.1. IP Address Reputation

There is a need to understand address assignment and configuration for hosts and services, especially with IPv6 [PlonkaBergerCARIS2] in (1) sharing IP-address-related information to inform attack response efforts while still protecting the privacy of victims and possible attackers and (2) mitigating abuse by altering the treatment, e.g., dropping or rate-limiting, of packets. Currently, there is no database that analysts and researchers can consult to, for instance, determine the lifetimes of IPv6 addresses or the prefix length at which the address is expected to be stable over time. The researchers propose either introducing a new database (compare PeeringDB) or extending existing databases (e.g., the regional Internet registries (RIRs)) to contain such information and allowing arbitrary queries. The prefix information would either be provided by networks that are willing or based on measurement algorithms that reverse-

engineer reasonable values based on Internet measurements [PlonkaBergerKIP]. In the former case, the incentive of networks to provide such information is to ensure that privacy of their users is respected and to limit collateral damage caused by access control lists affecting more of that network's addresses than necessary, e.g., in the face of abuse. This is an early idea; Dave Plonka is the lead contact for those interested in helping to develop this further.

4.4.2. Server Name Authentication Reputation C (SNARC)

SNARC is a mechanism to assign value to trust indicators, used to make decisions about good or bad actors. The mechanism would be able to distinguish between client and server connections and would be human readable. In addition, it builds on zero trust networking and avoids consolidation, thus supporting legitimate new players. SNARC has a similar theme to the IP reputation/BGP ranking idea mentioned above. SNARC is not currently defined by an RFC; however, such an RFC would help customers and design teams on existing solutions. The group plans to research visual aspects and underlying principles as they begin work on this idea. They plan to begin work in several stages, researching "trust" indicators, "trust" value calculations, and research actions to apply to "trust". The overarching goal is to address blind trust, one of the challenges identified with information/incident exchanges. Trent Adams is the lead contact for those interested in working with this team.

4.4.3. Logging

The group presented the possibility of injecting logging capabilities at compile time for applications, resulting in a more consistent set of logs, covering an agreed set of conditions. Using a log-injecting compiler would increase logging for those applications and improve the uniformity of logged activity. Increasing logging capabilities at the endpoint is necessary as the shift toward increased use of encrypted transport continues. Nalini Elkins is the lead contact for those interested in developing this further.

4.4.4. Fingerprinting

Fingerprinting has been used for numerous applications on the Web, including security, and will become of increasing importance with the deployment of stronger encryption. Fingerprinting provides a method to identify traffic without using decryption. The group discussed privacy considerations and balancing how you achieve the security benefits (identifying malicious traffic, information leakage, threat indicators, etc.). They are interested in deriving methods to validate the authenticity without identifying the source of traffic. They are also concerned with scaling issues. William Weinstein is the lead contact for those interested in working with this team.

4.5. Taxonomy and Gaps Session

At the start of the second day of the workshop, Kirsty Paine and Mirjam Kühne prepared (and Kirsty led) a workshop-style session to discuss taxonomies used in incident response, attacks, and threat detection, comparing solutions and identifying gaps. The primary objective was to

determine a path forward by selecting the language to be used in the proposed SMART Research Group. Several taxonomies were presented for review and discussion. The topic remains open, but the following key points were highlighted by participants:

- A single taxonomy might not be the way to go, because which taxonomy you use depends on what problem you are trying to solve, e.g., attribution of the attack, mitigation steps, technical features, or organizational impact measurements.
- A tool to map between taxonomies should be automated, as there are requirements within groups or nations to use specific taxonomies.
- The level of detail needed for reporting to management and for the analyst investigating the incident can be very different. At the workshop, one attendee mentioned that, for management reporting, they only use 8 categories to lighten the load on analysts, whereas some of the taxonomies contain 52 categories.
- How you plan to use the taxonomy matters and may vary between use cases. Take, for instance, sharing data with external entities versus internal only. The taxonomy selected depends on what you plan to do with it. Some stated a need for attribute-based dynamic anthologies as opposed to rigid taxonomies used by others. A rigid taxonomy did not work for many from feedback in the session.
- [\[RFC4949\]](#) was briefly discussed as a possibility; however, there is a clear need to update terminology in this publication around this space in particular. This is likely to be raised in the Security Area Advisory Group (SAAG) during the open mic session, hopefully with proposed new definitions to demonstrate the issue and evolution of terms over time.
- Within a taxonomy, prioritization matters to understand the impact of threats or an attack. How do you map that between differing taxonomies? What is the problem to be solved, and what tooling is required?
- Attack attribution had varying degrees of interest. Some felt the public sector cared more about attribution, not about individuals. They were interested in possible motivations behind an attack and determining if there were other likely victims based on these motivations. Understanding if the source was an individual actor, organized crime, or a nation state mattered.

The result of this discussion was not to narrow down to one taxonomy but to think about mappings between taxonomies and the use cases for exchanging or sharing information, eventually giving rise to a common method to discuss threats and attacks. Researchers need a common vocabulary, not necessarily a common taxonomy.

5. Next Steps

The next steps from the CARIS2 workshop are twofold:

1. The research initiatives spawned from the second CARIS workshop require further exploration and development. Fostering this development and creating communities around each proposed project is the first step, with reports back out to the SMART mailing list.
2. The second initiative will be planning for the next CARIS workshop.

6. Summary

When wrapping up the workshop, we reviewed the list of agreed projects to get a feel for actual interest as a follow up. Through the course of the two-day workshop, a larger set of potential research items had been generated, and this gave participants a chance to reassess commitments to better have them match expected outcomes. The highest ranking projects in terms of interest to drive the ideas forward included the following:

- Traffic fingerprinting
- SNARC
- Attack coordination solutions and automated security
- Cryptographic rendezvous
- L2 discovery

7. Security Considerations

There are no security considerations, as this is an informational workshop summary report.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Informative References

- [CARISEvent]** Internet Society, "CARIS2: Coordinating Attack Response at Internet Scale", February 2019, <<https://www.internetsociety.org/events/caris2>>.
- [deficit]** Morgan, S., "Cybersecurity Talent Crunch To Create 3.5 Million Unfilled Jobs Globally By 2021", October 2019, <<https://cybersecurityventures.com/jobs/>>.
- [I2NSF]** IETF, "Interface to Network Security Functions (i2nsf)", <<https://datatracker.ietf.org/wg/i2nsf/about>>.
- [IEEE802.11]** IEEE, "IEEE 802.11 WIRELESS LOCAL AREA NETWORKS", <<https://www.ieee802.org/11/>>.
- [MISP]** MISP, "Malware Information Sharing Platform", <<https://www.misp-project.org/>>.
- [PlonkaBergerCARIS2]** Plonka, D. and A. Berger, "Measured Approaches to IPv6 Address Anonymization and Identity Association", CARIS2 Paper Submission, March 2019, <<https://www.internetsociety.org/events/caris2>>.

- [PlonkaBergerKIP]** Plonka, D. and A. Berger, "kIP: a Measured Approach to IPv6 Address Anonymization", July 2017, <<https://arxiv.org/abs/1707.03900>>.
- [QUIC]** Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", Work in Progress, Internet-Draft, draft-ietf-quic-transport-33, 13 December 2020, <<https://tools.ietf.org/html/draft-ietf-quic-transport-33>>.
- [RFC4949]** Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC7970]** Danyliw, R., "The Incident Object Description Exchange Format Version 2", RFC 7970, DOI 10.17487/RFC7970, November 2016, <<https://www.rfc-editor.org/info/rfc7970>>.
- [RFC8073]** Moriarty, K. and M. Ford, "Coordinating Attack Response at Internet Scale (CARIS) Workshop Report", RFC 8073, DOI 10.17487/RFC8073, March 2017, <<https://www.rfc-editor.org/info/rfc8073>>.
- [RFC8446]** Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8520]** Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.
- [RFC8548]** Bittau, A., Giffin, D., Handley, M., Mazieres, D., Slack, Q., and E. Smith, "Cryptographic Protection of TCP Streams (tcpcrypt)", RFC 8548, DOI 10.17487/RFC8548, May 2019, <<https://www.rfc-editor.org/info/rfc8548>>.
- [SMART]** IRTF, "Stopping Malware and Researching Threats (smart)", <<https://datatracker.ietf.org/group/smart/about/>>.

Acknowledgements

Thank you to each of the CARIS2 workshop participants who brought their ideas, energy, and willingness to collaborate to advance attack response at Internet scale.

A big thank you to each member of the program committee for your review of program materials, papers, and guidance on the workshop format: Mat Ford (Internet Society, UK); Jamie Gillespie (APNIC, AU); Chris Inacio (CERT/CC, US); Mirja Kühlewind (ETH Zurich, CH); Mirjam Kühne (RIPE NCC, NL); Carlos Martinez (LACNIC, UY); Kathleen M. Moriarty, Chair (Dell EMC); Kirsty Paine (NCSC, UK); and Takeshi Takahashi (NICT, JP).

Thank you to Megan Hyland (Dell EMC) for her review and guidance on the format of breakout sessions and tools to enable successful collaboration.

Thank you to the minute takers, Akashaya Khare and Thinh Nguyen (Dell EMC OCTO Cambridge Dojo team).

Author's Address

Kathleen M. Moriarty
Center for Internet Security
31 Tech Valley Drive
East Greenbush, NY 12061
United States of America
Email: kathleen.moriarty.ietf@gmail.com