Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 2 December 2019

K. Moriarty
DellEMC

31 May 2019

Coordinating Attack Response at Internet Scale 2 (CARIS2) Workshop Report
draft-moriarty-caris2-01

Abstract

The *Coordinating Attack Response at Internet Scale (CARIS) 2* workshop workshop [CARISEvent], sponsored by the Internet Society, took place 28 February and 1 March 2019 in Cambridge, Massachusetts, USA. Participants spanned regional, national, international, and enterprise 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 for CARIS 2 on scaling incident prevention and detection as the Internet industry moves to stronger and a more ubiquitous deployment of session encryption.

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

The Coordinating Attack Response at Internet Scale (CARIS) 2 workshop, sponsored by the Internet Society, took place 28 February ? 1 March 2019 in Cambridge, Massachusetts, USA. Participants spanned regional, national, international, and enterprise 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 for CARIS 2 on scaling incident prevention and detection as the Internet industry moves to stronger and a more ubiquitous deployment of session encryption. Considering the related initiative to from a research group [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 detect attacks using protocol design ideas 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.

2. Conventions

3. 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 included:

Visualizing Security Automation: Takeshi Takahashi, NICT, Japan

Automating Severity Determination: Hideaki Kanehara, NICT, Japan

OASIS's OpenC2, Draper and DoD

Automated IoT Security (PASC and PAVA): 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 towards the goals of the workshop. These sessions included:

Manufacturer Usage Description (MUD) [RFC8520]: Eliot Lear, Cisco

TF-CSIRT: Mirjam Kuhne, RIPE NCC

M2M Sharing Revolution, Scott Pinkerton, DoE ANL

Comparing OpenC2 with existing efforts, e.g. I2NSF: Chris Inacio

Alternate Sharing and Mitigation Models: Kathleen Moriarty, DellEMC

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

4. CARIS2 Goals

The goal of each CARIS workshop has been to focus on the challenge of scaling attack response because of the overall concern in industry on the lack of information security professionals to fill the job gap. Currently, there is a 2 million person deficit for security professionals worldwide and it's only expected to grow. The chair's belief is that this gap cannot be filled through training, but the gap requires measures to reduce the number of information security professionals needed through new architectures and research towards attack prevention. CARIS 2 was specifically focused on the industry shift towards the increased use of stronger session encryption (TLSv1.3, QUIC, TCPcrypt, 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/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.

5. Workshop Collaboration

Both CARIS workshops have brought together a unique set of individuals who have not previously had the chance to be in the same room or collaborate toward the goals of scaling attack response. This is important as the participants span various areas of Internet technology work, research, provide a global perspective, have access to varying data sets and infrastructure, and are influential in their area of expertise. The specific goals of the CARIS 2 workshop, contributions, and the participants 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, involving the full set of attendees and breakout groups.

5.1. Breakout 1 Results: Standardization and Adoption

The goal of this session was to consider points raised in the talks that preceded the breakout on hurdles for automating security controls, detection, and response as the teams presenting noted several challenges they still face today. The collaborative session worked toward identifying standard protocols and data formats that succeeded in achieving adoption and several that have failed or only achieved limited adoption. The breakout teams selected protocols that failed and were successful for group discussion and the results from their evaluation were interesting and could help advance work in these or related areas if considered further.

Wide adoption:

Secure Sockets Layer (SSL), now replaced by Transport Layer Security (TLS) protocol.

Observations: There was a clear need for session encryption at the transport layer to protect application data. eCommerce 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.

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. SNMP was also first to market, with no competition. The protocol facilitated an answer to a needed problem set: configuration, telemetry, and

network management. It's development considered the connection between the user, vendor, and developers. Challenges did surface for adoption of SNMPv1.1 and 1.2, 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 the 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.

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

Although STIX and IODEF are somewhat similar in their goals, the standards were selected for evaluation by two separate groups with some common findings.

Structured Threat Information eXpression (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.

Incident Object Description Exchange Format (IODEF) was specified by NRENs and CSIRTs and formalized in the IETF. 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 towards 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 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. The end user being everyone was 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.

5.2. Breakout 2 Results:Preventative Protocols and Scaling Defense

This next breakout followed the sessions on MUD, PAVA (Protocol for Automated Vulnerability Assessment), and PASC (Protocol for Automatic Security Configuration) which have themes of automation at scale. MUD was designed for 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 next breakout was focused on groups brainstorming on preventative measures and enabling vendors to deploy mitigations.

One group dove a bit deeper into *MUD and layer 2 (L2) discovery*. While the overall value of MUD, shifting the majority of control management to the vendor for a predictable platform scales well, the use of MUD and what traffic is expected for a particular device is 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, manufacturer, etc. There was some discussion on the use of 802.11 to improve connections. Several participants from this group planned 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 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 third breakout dove deeper into this topic.

The forth breakout group discussed *useful data for incident responders*. This built on the first breakout session. 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 event threat sharing format. When reporting back to the group, a successful solution used in the EU was mentioned, Malware Information Sharing Platform (MISP) [MISP]. This will be considered in their review of existing efforts to determine if anything new is needed.

5.3. Breakout 3 Results: *Incident Response Coordination*

Incident response coordination currently does not scale. This breakout session focused on brainstorming on 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 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 provides an initially a burst of wind, but eventually leads to complacency

Creating Drag:

* Lack of clear KPIs
* Too many standards
* Regional border 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
* Over-centralization of the Internet
* New technologies and protocols
* Changes in application layer configurations (e.g. browser resolvers)

5.4. Breakout 4 Results: Monitoring and Measurement

The fourth breakout 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 co-author 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 pro-privacy work to defeat measurement objectives.

This breakout 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 4 highest then explored by the breakout groups. Groups agreed to work further on the proposed ideas.

IP 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 for analysts and researchers can consult to, for instance, determine to lifetimes of IPv6 addresses or the prefix length at

which the address is expected to be stable over time. We propose either introduce a new database (compare PeerdingDB) or extending existing databases (e.g., the RIRs'), to contain such information and allowing arbitrary queries. The prefix information would either be provided by networks, who 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 so 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, the lead to contact if interested to help develop this further is Dave Plonka.

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 in connections, would be human readable, builds on zero trust networking, and avoids consolidation supporting legitimate new players. The group planned to research visual aspects and underlying principles as they begin work on this idea. SNARC has a similar theme to the IP reputation/ BGP ranking idea mentioned above. An RFC would help customers and design team on existing solutions. They planned 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. If interested to work further with this team, the lead contact is: Trent Adams.

Logging

The breakout 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. If the log-injecting compiler were used this would increase logging for those applications and improve the uniformity of logged activity. Increasing logging capabilities at the endpoint is necessary as the shift towards increased use of encrypted transport continues. The lead for contact if interested to develop this further is Nalini Elkins.

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. This 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 to derive methods to validate the authenticity without identifying the source of traffic. They are also concerned with scaling issues. If interested to work further with this team, the lead contact is: William Weinstein.

5.5. Taxonomy and Gaps Session

At the start of day 2, Kirsty Paine and Mirjam Kuhne 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 selecting language to be used in the proposed SMART group. Several taxonomies were presented for review and discussion. The topic remains open, 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 SAAG, 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? (problem to be solved; possible tooling required)

* Attack attribution had varying degrees of interest. Some felt the public sector cared more about attribution; not about individuals, but the possible motivations behind an attack and likely other 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.

6. Next Steps

The next steps from the CARIS workshop are twofold. The research initiatives spawned from the second CARIS 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 IRTF SMART mailing list and in a proposed research group.

The second initiative will be planning for the next CARIS workshop. This is likely to be coupled with the FIRST Conference in 2020 geared around a topic important to incident responders to assist with scale as it relates directly to problems of interest to that community.

7. Summary

Wrapping up the workshop, we reviewed the list of agreed projects to get a feel for actual interest in follow up now that a larger set had been generated, giving participants a chance to reassess commitments to better have them match actual outcomes. The highest ranking projects in terms of interest to drive the ideas forward included the following:

* Traffic fingerprinting

* SNARC

* Attack coordination solutions/automated security

* Cryptographic Rendezvous

* L2 discovery

8. Security Considerations

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

9. IANA Considerations

This memo includes no request to IANA.

10. Contributors

Thank you to each of the CARIS 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 Kuhlewind, ETH Zurich, CH, Mirjam Kuhne, RIPE NCC, NL, Carlos Martinez, LACNIC, UY, Kathleen M. Moriarty, Dell EMC (Chair), Kirsty Paine, NCSC, UK, and Takeshi Takahashi, NICT, JP.

Thank you to Megan Hyland, DellEMC, for her review and guidance on the breakout session format and tools to enable successful collaboration.

Thank you to the minute takers, Akashaya Khare and Thinh Nguyen, DellEMC OCTO Cambridge Dojo team.

11. References

11.1. Informative References

[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>.

[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>.

11.2. URL References

[CARISEvent]
Internet Society, "CARIS Event Information and Accepted Papers https://www.internetsociety.org/events/caris2", May 2019.

[MISP] MISP-project.org, "Malware Information Sharing Platform https://www.misp-project.org/", May 2019.

[PlonkaBergerCARIS2]
CARIS2, "CARIS2 Paper Submission,", May 2019.

[PlonkaBergerKIP]
Arxiv, "kIP: a Measured Approach to IPv6 Address Anonymization https://arxiv.org/abs/1707.03900", 2017.

[SMART] IRTF, "Stopping Malware and Researching Threats https://datatracker.ietf.org/group/smart/about/", May 2019.

Appendix A. Change Log

Note to RFC Editor: if this document does not obsolete an existing RFC, please remove this appendix before publication as an RFC.

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Author's Address

Kathleen M Moriarty
DellEMC
176 South Street
Hopkinton, MA 01748
United States

Email: kathleen.moriarty.ietf@gmail.com