2013 LASER Workshop

Learning from Authoritative Security Experiment Results

Greg Shannon
CERT® Program
Carnegie Mellon University
Pittsburgh, PA, USA

Laura Tinnel, David Balenson
SRI International
Arlington, VA, USA

Tadayoshi Kohno
Department of Computer Science and Engineering
University of Washington,
Seattle, WA, USA

Christoph Schuba
Oracle
Redwood Shores, CA, US

Carrie Gates
CA Labs
Computer Associates Technologies
Islandia, NY, USA

Ed Talbot
Livermore, CA, USA

***Abstract*—Poster: With the increasing importance of computer security, the goal of this workshop is to help the security community quickly identify and learn from both success and failure. The workshop focuses on research that has a valid hypothesis and reproducible experimental methodology, but where the results were unexpected or did not validate the hypotheses, where the methodology addressed difficult and/or unexpected issues, or where unsuspected confounding issues were found in previous work. The specific security results of experiments are of secondary interest for this workshop.**

Keywords—computer; security; experiment; methodology; design; reproducibility; validation; verification; confounds; mistakes; misinterpretation; failure; mitigations

I. Introduction

Journals and conferences typically publish papers that report successful experiments that extend our knowledge of the science of security or assess whether an engineering project has performed as anticipated. Some of these results have high impact; others do not. Unfortunately, papers reporting on experiments with unanticipated results that the experimenters cannot explain, experiments that are not statistically significant, or engineering efforts that fail to produce the expected results are frequently not considered publishable because they do not appear to extend our knowledge. Yet, some of these "failures" may actually provide clues to even more significant results than the original experimenter had intended. The research is useful, even though the results are unexpected.

Useful research includes a well-reasoned hypothesis, a well-defined method for testing that hypothesis, and results that either disprove or fail to prove the hypothesis. It also includes a methodology documented sufficiently so that others can follow the same path. When framed in this way, "unsuccessful" research furthers our knowledge of a hypothesis and a testing method. Others can reproduce the experiment itself, vary the methods, and change the hypothesis, as the original result provides a place to begin.

As an example, consider an experiment assessing a protocol utilizing biometric authentication as part of the process to provide access to a computer system. The null hypothesis might be that the biometric technology does not distinguish between two different people; in other words, that the biometric element of the protocol makes the approach vulnerable to a masquerade attack. Suppose the null hypothesis is not rejected - there is still value in publishing this result. First, it might prevent others from trying the same biometric method. Second, it might lead them to further develop the technology—to determine whether a different style of biometrics would improve matters, or if the environment in which authentication is being attempted makes a difference. For example, a retinal scan may be a failure in recognizing people in a crowd, but successful where the users present themselves one at a time to an admission device with controlled lighting, or when multiple "tries" are included. Third, it might lead to modifying the encompassing protocol so as to make masquerading more difficult for some other reason.

Equally important is research designed to reproduce the results of earlier work. Reproducibility is key to science as a way to validate earlier work or to uncover errors or problems in earlier work. Failure to reproduce the results leads to a deeper understanding of the phenomena that the earlier work uncovers.

Finally, many discussions about papers, proposals, and projects seek to explore previously tried strategies that failed, usually because published work does not exist. Old ideas are often pursued because the community is not aware of the prior failure. The workshop provides a venue that can help resolve this gap in the security community's research literature.

Suggested topics for the workshop include, but are not limited to:

1. Unsuccessful research in experimental security

2. Methods and designs for security experiments

3. Experimental confounds, mistakes, and mitigations

4. Successes and failures reproducing experimental techniques and/or results

5. Issues in hypothesis and methods development (e.g., realism, fidelity, scale)

II. Purpose of This Poster

The goals of talking about this workshop at the Oakland conference are to (1) stimulate interest in the LASER workshop, (2) stimulate interest in the methods and failures of cyber security research, and (3) encourage other researchers to begin writing papers in a structured format that is contains details sufficient for peers to determine the validity of and repeat the experiments described in the paper.

III. Organizing Committee:

Laura Tinnel (SRI International), General Chair
Greg Shannon (CMU/CERT), Program Co-Chair
Tadayoshi Kohno (U Wash), Program Co-Chair
Christoph Schuba (Oracle), Proceedings
Carrie Gates (CA Technologies), Treasurer
David Balenson (SRI International), Local Arrangements
Ed Talbot (Consultant), Publicity

IV. Program Committee:

Greg Shannon (CMU/CERT), Co-Chair
Tadayoshi Kohno (U Wash), Co-Chair
David Balenson (SRI)
Matt Bishop (UC Davis)
Joseph Bonneau (Google)
Pete Dinsmore (JHU/APL)
Debin Gao (Singapore Management University)
Carrie Gates (Computer Associates Technologies)
Alefiya Hussain (USC/ISI)
Carl Landwehr (George Washington University)
Sean Peisert (UC Davis and LBNL)
Angela Sasse (University College London)
Christoph Schuba (Oracle)
Ed Talbot (Consultant)
Laura Tinnel (SRI)
Kami Vaniea (Michigan State)
Charles Wright (Portland State)

V. Submission Guidelines

Both full papers and structured abstracts are solicited. Full papers follow a typical pattern of submission, review, notification, pre-conference version, conference presentation, and final post-conference version. One-page structured abstracts serve two purposes: (1) to enable authors to receive early feedback prior to investing significant effort writing papers, and (2) to provide all attendees a forum to share an abstract of their work before the workshop.

Abstracts will be reviewed by at least two PC members with comments returned in 5-10 days; submissions before June 27 will receive an "encouraged," "neutral," or "discouraged" indication for submission of a full paper based on the abstract. The pre-submission feedback is for the author's use only. All abstracts deemed relevant by the PC will be available on the laser-workshop.org website before the conference, but they will not be part of the proceedings.

A structured abstract is typically 200-500 words and less than one page. It includes at least these elements: background, aim, method, results, and conclusions. See the workshop website for more details. The abstracts for full papers should be similarly structured.

Full paper submissions should be 6–10 pages long including tables, figures, and references. All submissions should use the ACM Proceedings format: http://www.acm.org/sigs/publications/proceedings-templates (Option 2, if using LaTeX). At least one author from every accepted full paper must plan to attend the workshop and present. All papers and abstracts must be submitted via OpenConf https://www.openconf.org/laser2013.

VI. Important Dates:

March 1	Start rolling consideration of 1-page structured abstracts
June 27	Full papers due
August 27	Authors notified of accepted/rejected full papers
September 23	Pre-conference versions of full papers due
September 30	End rolling consideration of 1-page structured abstracts
October 16-17	2013 LASER Workshop
November 15	Post-conference versions of full papers due

VII. For Further Information

See www.laser-workshop.org for full and up-to-date details on the workshop. Send questions to info@laser-workshop.org.

VIII. Sponsorship

The LASER workshop is funded in part by NSF Grant #1143766 and by the Applied Computer Security Associates (ACSA).