

35th USENIX Security Symposium

August 12–14, 2026, Baltimore, MD, USA

Sponsored by USENIX, the Advanced Computing Systems Association



The USENIX Security Symposium brings together researchers, practitioners, system programmers, and others interested in the latest advances in the security and privacy of computer systems and networks. The 35th USENIX Security Symposium will be held on August 12–14, 2026, in Baltimore, MD, USA.

Summary of main changes from previous editions

1. USENIX Security 2026 will no longer feature major revisions; papers will at most undergo a two-week shepherding process.
2. Authors may submit at most seven papers per cycle to the conference.
3. Mandatory registration for all papers one week before the submission deadline, including all authors, title, tentative abstract, and topics for the paper. The list of authors cannot be changed after registration.
4. Every author on a submission must use their HotCRP account to individually confirm compliance with the submission terms.

Important Dates

Cycle 1

- Mandatory registration (title, authors, tentative non-blank abstract, and topics) due: **Tuesday, August 19, 2025, 11:59 pm AoE**
- Paper submissions due: **Tuesday, August 26, 2025, 11:59 pm AoE**
- Early reject notification: **Tuesday, October 7, 2025**
- Rebuttal period: **November 6–13, 2025**
- Notification to authors: **Wednesday, December 4, 2025**
- Shepherded approval due: **Thursday, December 18, 2025**
- Final papers due: **Thursday, January 15, 2026**

Cycle 2

- Mandatory registration (title, authors, tentative non-blank abstract, and topics) due: **Thursday, January 29, 2026, 11:59 pm AoE**
- Paper submissions due: **Thursday, February 5, 2026, 11:59 pm AoE**
- Early reject notification: **Tuesday, March 17, 2026**
- Rebuttal period: **April 16–23, 2026**
- Notification to authors: **Thursday, May 14, 2026**
- Shepherded approval due: **Thursday, May 28, 2026**
- Final papers due: **Thursday, June 11, 2026**

Symposium Topics

Refereed paper submissions are solicited in all areas relating to systems research in security and privacy. This topic list is not meant to be exhaustive; USENIX Security is interested in all aspects of computing systems security and privacy. Papers without a clear application to security or privacy of computing systems, however, will be considered out of scope and may be rejected without full review at the discretion of the chairs.

- System security
 - Operating systems security
 - Distributed systems security
 - Cloud computing security
- Network security
 - Intrusion and anomaly detection and prevention
 - Network infrastructure security
 - Denial-of-service attacks and countermeasures
 - Wireless security
 - Analysis of network and security protocols
- Security and Privacy for Web, Mobile, and Emerging Technologies
 - Web
 - Mobile
 - IoT
 - VR/AR
 - Games
- Software Security
 - Malware analysis
 - Forensics and diagnostics for security
 - Automated security analysis of source code and binaries
 - Program analysis
 - Fuzzing and vulnerability discovery
- Security of ML
 - Read the [Submitting ML Work to USENIX Security](#) section below to prevent desk rejection
 - USENIX Security is a systems security venue. AI-focused contributions must be relevant to the broader systems security community
 - This topic is not meant for papers that propose/evaluate/improve AI-based techniques for a security task. Such papers should select the topic most closely related to the task.
- Privacy of ML
 - Read the [Submitting ML Work to USENIX Security](#) section below to prevent desk rejection
 - USENIX Security is a systems security venue. AI-focused contributions must be relevant to the broader systems security community
 - This topic is not meant for papers that propose/evaluate/improve AI-based techniques for a privacy task. Such papers should select the topic most closely related to the task.



- Privacy and Anonymity
 - Privacy metrics
 - Anonymity
 - Privacy-preserving computation
 - Privacy attacks
 - Surveillance and censorship
- Human Aspects
 - Usable security and privacy
 - Security and privacy law, policy, and/or ethics
 - Security education and training
 - Understanding, measuring, quantifying, and protecting users from: information manipulation, mis/disinformation, harassment, extremism, and abuse via qualitative and quantitative methods
- Hardware security
 - Secure computer architectures
 - Embedded systems security
 - Cyber-physical systems security
 - Methods for detection of malicious or counterfeit hardware
 - Side channels
 - Automated security analysis of hardware designs and implementation
- Applications of cryptography
 - Analysis of deployed cryptography and cryptographic protocols
 - Cryptographic implementation analysis
 - New cryptographic protocols with real-world applications
 - Blockchains and distributed ledger security

Special Submissions

Replication

USENIX Security solicits submission of replication papers, which verify, refute, or refine prior technical results. We encourage submissions that not only replicate studies but also offer meta-analyses that assess the replicability of research. Additionally, while replication studies often replicate original findings, we also value novel investigations into why certain studies fail to replicate. Papers that critically examine the conditions under which replication is feasible, or those that propose innovative methods to enhance the reliability of scientific findings, are also welcome. Papers with this specific purpose must be flagged accordingly in the HotCRP submission system.

Systematization of Knowledge

USENIX Security solicits the submission of Systematization of Knowledge (SoK) papers, which have been very valuable to help our community to clarify and put into context complex research problems.

It is important to stress that SoK papers go beyond simply summarizing previous research (like in a survey); they also include a thorough examination and analysis of existing approaches, identify gaps and limitations, and offer insights or new perspectives on a given, major research area.

While both SoK and survey papers may involve summarizing existing research, the key difference is that an SoK paper provides a more structured and insightful overview, which might also involve new experiments to replicate and compare previous solutions. For examples, please see the list of SoK papers at <https://oaklandsok.github.io/>.

Papers with this specific purpose must have their title prefixed with “SoK: “.

Bulk Submissions

Each author can submit at most seven papers per cycle to USENIX Security 2026. For authors submitting more than seven papers, the chairs will only retain the seven papers with the lowest submission numbers. This rule is enforced per author, i.e., authors may have papers rejected even if they stay under the limit in cases where a co-author violates the requirement. Once the registration deadline has passed, the submission’s author list may not be changed. Even if one of the first seven papers is later withdrawn, this does not change the reject status of the paper(s) submitted 8th+ as of the submission deadline. Any attempts to bypass this rule (e.g., by having multiple HotCRP accounts) will be escalated to the Steering Committee.

Submission Policies and Instructions

Executive Summary

All papers must adhere to the following rules (outlined more verbosely below). Papers found in violation of any of these requirements will be desk-rejected without review.

- All papers MUST be registered in the respective HotCRP instance one week before the submission deadline, including the full list of authors, the title, the (tentative, non-blank) abstract, and the topics of the paper.
- All papers MUST have a discussion of research ethics. This MUST be in a separate appendix called “Ethical Considerations.”
- All papers MUST discuss Open Science. This MUST be in a separate appendix called “Open Science” that clearly lists where the artifacts necessary for evaluating the contributions of their submission are located. These artifacts must be available at the time of submission.
- All papers MUST NOT reveal the identity of the authors (e.g., names, self-references with “we”, funding acknowledgements, Github repositories).
- All papers MUST comply with the unaltered USENIX Security LaTeX template. Any attempts to remove whitespace (e.g., negative vspaces, savetrees, titlesec, removing author blocks, etc.) are strictly forbidden. While usage of the Word template is not strictly forbidden, any non-compliant formatting will lead to a desk rejection.
- All authors MUST provide their ORCID IDs through the profile page on HotCRP by the registration deadline.
- All authors MUST acknowledge the submission terms on HotCRP by the submission deadline.
- Authors MUST NOT submit more than seven papers per cycle.
- At least one author of each accepted paper MUST register for and commit to attending the conference.

Submission Deadlines

USENIX Security ‘26 submissions deadlines are as follows:

Cycle 1

- Paper registrations due: **Tuesday, August 19, 2025**
- Paper submissions due: **Tuesday, August 26, 2025**

Cycle 2

- Paper registrations due: **Thursday, January 29, 2026**
- Paper submissions due: **Thursday, February 5, 2026**

All papers accepted in either Cycle will appear in the proceedings of USENIX Security ‘26. All submissions should be made online via their respective submission systems on the Call for Papers page. We do not accept any other form of submission or updates outside of the submission systems.

Submitted papers should describe original, scientifically sound work produced by the co-authors. All submissions will be judged on originality, relevance, correctness, and clarity. Submissions

should be finished, complete papers. We may desk-reject papers that have severe editorial problems (broken references, egregious spelling or grammar errors, missing figures, etc.), are submitted in violation of the Submission Instructions outlined below, are outside of the scope of the symposium, or are deemed clearly of insufficient quality to appear in the program.

Submissions must be in PDF format. Please make sure your submission can be opened using Adobe Reader. Please make sure your submission, and all embedded figures, are intelligible when printed in grayscale.

Formatting Requirements & Page Limit

In Brief:

- Papers can include up to 13 pages of body text
- Papers must include “Ethical Considerations” and “Open Science” appendices. These appendices must appear immediately following the main body, before the references/appendices and may each one page of text.
- Overall length:
 - The initial submission can contain unlimited appendices and references
 - Final camera-ready papers can contain a maximum of 20 total pages

Submissions must be typeset on U.S. letter-sized pages in two-column format in 10-point Times Roman type on 12-point leading (single-spaced), in a text block 7” x 9” deep. The main body of the submission must not exceed 13 pages, immediately followed by the “Ethical Considerations” and “Open Science” appendices (each of which may consume up to one page). Authors must use the [USENIX Security templates and style files](https://www.usenix.org/sites/default/files/usenixsecurity2026_latex_templates.zip) (https://www.usenix.org/sites/default/files/usenixsecurity2026_latex_templates.zip) when preparing the paper for submission. We strongly encourage the use of the unaltered USENIX LaTeX template. While usage of the Word template is not strictly forbidden, any erroneous formatting will lead to a desk rejection. Failure to adhere to the page limit and formatting requirements are grounds for rejection.

At submission time, there is no limit on the length of the bibliography and appendices but reviewers are not required to read any appendices. These appendices may be included to assist reviewers who may have questions that fall outside the stated contribution of the paper on which your work is to be evaluated, or to provide details that would only be of interest to a small minority of readers. The paper must be self-contained without appendices.

Papers must not attempt to “squeeze space” by exploiting underspecified formatting criteria (e.g., columns) or through manipulating other document properties (e.g., page layout, spacing, fonts, figures and tables, headings). Papers that, in the chair’s assessment, make use of these techniques to receive an unfair advantage, will be rejected, even if they comply with the above specifications. We offer [several examples](#) of observed techniques that have or could lead to rejection. Authors should seek to meet page limits through the modification of content alone. Any other techniques (whether appearing in these examples or not) may result in rejection.

For the final version of their paper, authors may fill at most 18 pages with the main paper, mandatory “Ethical Considerations” and “Open Science” appendices, the bibliography, and any additional appendices. In total, the final PDF must therefore not exceed 20 pages. Papers accepted conditional to shepherd approval are allowed to follow the final version’s page limits.

Ethical Considerations

Reviewers will be asked to evaluate the ethics of all submissions. Authors are expected to complete a stakeholder-based ethics analysis (see the [Ethics Guidelines](#) section below) or justify and complete an alternative approach to considering ethics. All submissions are hence required to have a dedicated appendix after the main body of the paper (but before any references) called “Ethical Considerations”. Authors are expected to thoroughly study the [Ethics Guidelines](#) section below and laid out by the conference and follow these guidelines when preparing their statement.

Authors should be prepared to answer these questions in the conference submission portal at the time of registration:

- “I attest that I read the ethics discussions in the conference call for papers, the detailed submissions instructions, and the ethics guidelines.”
- “I attest that the research team considered the ethics of this research, that the authors believe the research was done ethically, and that the team’s next-step plans (e.g., after publication) are ethical.”
- “I attest that the submission has a clearly-marked appendix of up to one page on ethical considerations that complies with the ethics guidelines.”

At the chairs’ discretion, papers may be desk rejected for missing or inadequate ethics statements. This includes, but is not limited to, missing the required statement, missing the information required to be contained in the statement, and incorrect naming of the appendix.

Open Science Policy and Artifact Evaluation

In 2025, USENIX Security introduced a new open-science policy, aiming to enhance the reproducibility and replicability of scientific findings. In 2026, **authors are required to openly share their research artifacts at submission time**. This initiative is part of a broader commitment to foster open-science principles, emphasizing the sharing of artifacts such as datasets, scripts, binaries, and source code associated with research papers.

In accordance with this policy, all USENIX Security papers must make their compliance with the open-science policy explicit: the “Open Science” appendix must list all artifacts necessary to evaluate the contribution of the paper and make clear how the review committees can access each artifact. All artifacts must be available at the time of submission. **Paper acceptance is conditional on public availability of these artifacts**. Artifacts will be used as part of the reviewers’ evaluation process. Thus, during submission, any artifacts and the links to reach them must be anonymized.

In case artifacts cannot be shared (e.g., for licensing restrictions, to prevent adversarial methods from being used adversarially before defenses are deployed, or to prevent harm from study subjects), the statement must explain the reasons for this omission.

Following conditional paper acceptance, there will be an opportunity to provide a non-anonymized camera-ready link(s) to the paper’s artifacts. Artifact availability at the camera-ready link(s) will be re-verified by the Artifact Evaluation Committee.

Authors of accepted papers are encouraged to register their artifacts to also be checked for functionality and reproducibility. Authors that indicate a desire to have their artifacts evaluated for functionality and reproducibility will undergo that evaluation after acceptance.

Anonymous Submission

The review process will be anonymous. Papers must be submitted in a form suitable for anonymous review:

- The title page must not contain any author names or affiliations.
- Authors should carefully review figures and appendices (especially survey instruments) to ensure affiliations are not accidentally included.
- When referring to your previous work, do so in the third person, as though it were written by someone else. Anonymous references are only allowed in the (unusual) case that a third-person reference is infeasible, and after approval of the chairs.
- Authors may include links to websites that contain source code, tools, or other supplemental material. Neither the link in the paper nor the website itself may suggest the authors' identities (e.g., the website must not contain the authors' names or affiliations). We require artifacts to be available during paper review, but suggest authors make use of anonymized repositories (e.g., Anonymous 4OpenScience) or repositories that allow for the creation of anonymized links (OSF) to avoid leaking their identity.

Papers that are not properly anonymized may be rejected without review or during the review process if detected later on in the process.

While submitted papers must be anonymous, authors may choose to give talks about their work, post a preprint of the paper online, disclose security vulnerabilities to vendors or the public, etc., during the review process. Authors should take care to avoid timing of public announcements, use of USENIX templates, and other behaviors that appear to intentionally aim to inform reviewers of their identity.

Submission Instructions and Procedures

Registration and Additional Requirements for Submissions

All papers to be submitted to the conference must be registered with the title, the full list of authors, a tentative abstract, and topics one week before the submission deadline. Any papers not registered at that time or those missing any of the aforementioned requirements will not be considered for review.

All authors are required to 1) set up or update their HotCRP profile to include their Orcid and 2) confirm the submission terms within HotCRP. In exceptional cases of unavailability of a co-author, the chairs may be contacted with an explanation of why the requirements cannot be fulfilled. In all other cases, papers for which at least one author fails to comply with the requirements will not be considered for review.

Conflicts of Interest

The program co-chairs require cooperation from both authors and program committee members to prevent submissions from being evaluated by reviewers who have a conflict of interest. During the submission process, we will ask authors to identify members of the program committee with whom they share a conflict of interest. This includes anyone

1. who shares an institutional affiliation with an author at the time of submission (including secondary affiliations and consulting work)
2. who was the advisor or advisee of an author at any time in the past
3. with whom the author has collaborated (including on-going research conversations, large grants) or published in the current year or the two years preceding it
4. who is affiliated with a party that funds your research

5. who can identify the authors (e.g., because of previous conversations about the work)

6. with whom you have a close personal relationship

For other forms of conflict, authors must contact the chairs and explain the perceived conflict. In addition to selecting program committee conflicts when submitting, we recommend that all authors ensure they have up-to-date HotCRP profiles listing all known conflicts.

Program committee members who have conflicts of interest with a paper, including program co-chairs, will be excluded from the evaluation and discussion of the paper.

Final versions of accepted submissions should include all sources of funding in an acknowledgments section. Authors should also disclose any affiliations, interests, or other facts that might be relevant to readers seeking to interpret the work and its implications. Authors may wish to consider the [2023 IEEE S&P Financial Conflicts Policy](https://www.ieee-security.org/TC/SP2023/financial-con.html) (<https://www.ieee-security.org/TC/SP2023/financial-con.html>) for example.

To prevent retroactive conflicts of interest, all authors must be declared at registration time.

In case of missed conflicts, the PC chairs may update conflicts based on external information (such as DBLP). In addition, in egregious cases of missed conflicts, authors will be reported to USENIX.

Confidentiality of Submissions

The program committee and external reviewers are required to treat all submissions as confidential. All papers must only be read for the purposes of evaluation for publication. However, the program co-chairs or designated committee members may share submissions outside the program committee to allow chairs of other conferences to identify dual submissions.

Internet Defense Prize

The Internet Defense Prize recognizes and rewards research that meaningfully makes the internet more secure. Created in 2014, the award is funded by Meta and offered in partnership with USENIX to celebrate contributions to the protection and defense of the internet. Successful recipients of the Internet Defense Prize will provide a working prototype that demonstrates significant contributions to the security of the internet, particularly in the areas of prevention and defense. This award is meant to recognize the direction of the research and not necessarily its progress to date. The intent of the award is to inspire researchers to focus on high-impact areas of research. The USENIX Security Awards Committee—selected by the Program Chairs among the symposium Program Committee members—independently determines the prize, to be distributed by USENIX.

Simultaneous Submission and Plagiarism

Simultaneous submission of the same work to multiple venues, submission of previously published work, and plagiarism constitute dishonesty or fraud. This includes cases in which preliminary reviews from another venue indicate likely rejection but the paper is still considered under submission to said venue. Authors should relate their submission to any other relevant submissions of theirs in other venues that are under review at the same time as their submission to the Symposium. These citations to simultaneously submitted papers should be anonymized; non-anonymous versions of these citations must, however, be emailed to the program co-chairs at sec26chairs@usenix.org. Failure to point out and explain overlap with published or simultaneously submitted papers will be grounds for rejection.

USENIX, like other scientific and technical conferences and journals, prohibits these practices and may take action against authors who have committed them. See the [USENIX Conference Submissions Policy](https://www.usenix.org/conferences/author-resources/submissions-policy) (<https://www.usenix.org/conferences/author-resources/submissions-policy>) for details.

All submitted papers are considered to be under review for USENIX Security '26 until authors are notified of a decision by the program committee or the program co-chairs approve a request for withdrawal. Papers may no longer be withdrawn once the author response phase has started. Any CfP violations after the response phase has begun will not remove the paper from consideration at USENIX until the official notification date.

Embargo Requests

Authors may request an embargo for their papers by the deadline dates listed below. All embargoed papers will be released on the first day of the conference, Wednesday, August 12, 2026.

- Cycle 1 deadline for embargo requests:
Thursday, February 12, 2026
- Cycle 2 deadline for embargo requests:
Thursday, July 9, 2026

If your accepted paper should not be published prior to the event, please notify production@usenix.org after you submit your final paper.

Outcomes and Conference Expectations

Review Process Outcomes

The reviewing process will result in one of the following three outcomes:

- **Accepted**
In this case, reviewers either see no need for any required edits—the paper could be published as is—or trust the authors to make any required edits.
- **Accepted on Shepherd Approval**
These papers offer sufficient contributions to be in the USENIX Security program, but the reviewers feel that some changes are needed. Reviewers expect that these changes are highly unlikely to reduce their enthusiasm regarding the paper being published. These can include text changes, clarifications, and explicit discussion of limitations.
- **Rejected**
This outcome indicates that the paper is not currently appropriate for publication at USENIX Security. Reviews for such papers will attempt to identify major changes that could move the paper toward a publishable state.

Papers rejected from either cycle of USENIX Security '25 may be submitted. Papers rejected from the first cycle of USENIX Security '26 may not be submitted to the second cycle of USENIX Security '26. Restrictions on submissions rejected from the second cycle will be decided by the USENIX Security '27 PC co-chairs.

Conference Attendance and Publishing Accepted Papers

Papers that have been formally reviewed and accepted will be presented during the Symposium and published in the Symposium Proceedings. By submitting a paper, if the paper is accepted, at least one of the authors will register to attend the conference at full price (i.e., not the student rate) and to present the paper; USENIX members at the Advocate level and higher may apply their membership discounts to their registrations. If an author plans to present more than one paper, one full-price registration will still be required for each paper. If the conference registration fee will pose a hardship for the presenter of the accepted paper, please contact conference@usenix.org.

A major mission of the USENIX Association is to provide for the creation and dissemination of new knowledge. USENIX allows authors to retain ownership of the copyright in their works, requesting only that USENIX be granted the right to be the first publisher of that work. See our [sample consent form](https://www.usenix.org/sites/default/files/consent_author_proceedings.pdf) (https://www.usenix.org/sites/default/files/consent_author_proceedings.pdf) for the complete terms of publication.

Papers accepted during the first reviewing cycle will be published on the USENIX Security website after the camera ready deadline for the first reviewing cycle. Papers accepted during the second reviewing cycle will be published on the first day of the symposium.

See the Paper Submission Policies and Instructions page for more information.

Presentation of Papers

Motivated by rising conference costs and increasing numbers of submitted and accepted papers, USENIX Security '25 is implementing a new approach to presenting accepted papers and fostering interactions at the conference. At this time, the exact format for the presentations at USENIX Security '26 has not been decided and will follow based on the experience of presenters and attendees of USENIX Security '25.

Contact Information

Specific questions about submissions may be sent to the program co-chairs at sec26chairs@usenix.org. The chairs will respond to individual questions about the registration process if contacted at least a week before the registration deadline and will respond to individual questions about the submission process if contacted at least a week before the submission deadline.

Further questions? Contact your program co-chairs, sec26chairs@usenix.org, or the USENIX office, submissionspolicy@usenix.org.

Submitting ML Work to USENIX Security

If you are using machine learning to solve a security or privacy problem, mark the paper's primary field as the field of the problem you're solving. For instance, a paper proposing a novel ML method for Intrusion Detection should be tagged with "Network Security" as the primary field not Security & Privacy of ML as the primary field; a paper proposing a novel attack against an ML pipeline for phishing website detection should be tagged with "Web Security" as the primary field; a paper evaluating previously-proposed ML techniques for malware analysis should be tagged with "Software Security" as the primary field.

If you are working on the security or privacy of machine learning, this document will help you decide whether USENIX is an appropriate venue for your submission.

We note that USENIX Security remains a systems security venue. Therefore, ML-focused contributions must be relevant to the broader systems security community. **Papers deemed out of scope will be desk rejected.**

Relevant works may investigate novel security and privacy attacks and defenses, and their implications, across the ML lifecycle, including data curation for pre-training and post-training, the training phases (pre-training and post-training), and deployment. Examples of attacks include data poisoning, model poisoning, backdoors, adversarial examples, prompt injection, jailbreaks, model inversion, membership inference, and model stealing. Defenses may involve redesigning ML algorithms and pipelines to prevent attacks, developing methods to detect

attacks, advancing post-attack forensic analysis and recovery techniques, and understanding the implications of defenses on model or pipeline performance.

All papers submitted should provide a threat model that clearly articulates the (i) envisioned attacker(s), (ii) threat surfaces (e.g., system components including but not limited to the underlying machine learning algorithm), (iii) generality and (iv) practicality of the attack.

Papers on robustness need to have a clear security flavour, presenting an adversary that aims to, deliberately, affect the performance of the system arbitrarily or in a targeted manner. Works that primarily focus on improving ML functionality or efficiency (e.g., robustness to noise on the data or spurious artifacts) are not in scope.

Examples

To assist in selecting the primary field of your paper, we provide below a list of ML-related papers accepted to recent editions of USENIX Security, and explain which primary field in USENIX Security '26 should be selected:

- Paper: [KnowPhish: Large Language Models Meet Multimodal Knowledge Graphs for Enhancing Reference-Based Phishing Detection](https://www.usenix.org/conference/usenixsecurity24/presentation/li-yuexin) (https://www.usenix.org/conference/usenixsecurity24/presentation/li-yuexin) → Primary Topic: Web Security (The paper proposes the usage of LLMs for Phishing Website Detection).
- Paper: [Uncovering the Limits of Machine Learning for Automatic Vulnerability Detection](https://www.usenix.org/conference/usenixsecurity24/presentation/risse) (https://www.usenix.org/conference/usenixsecurity24/presentation/risse) → Primary Topic: Software Security (The paper evaluates the application of ML for the task of vulnerability detection, which pertains to software security).
- Paper: [PentestGPT: Evaluating and Harnessing Large Language Models for Automated Penetration Testing](https://www.usenix.org/conference/usenixsecurity24/presentation/deng) (https://www.usenix.org/conference/usenixsecurity24/presentation/deng) → Primary Topic: Systems Security (The paper studies the application of LLMs for penetration testing of systems.)
- Paper: [Formalizing and Benchmarking Prompt Injection Attacks and Defenses](https://www.usenix.org/conference/usenixsecurity24/presentation/liu-yupe) (https://www.usenix.org/conference/usenixsecurity24/presentation/liu-yupe) → Primary Topic: Security of ML (The paper evaluates various Prompt Injection Attacks and Defenses, which are specific vulnerabilities of ML-based systems)

Ethics Guidelines

As stated in the CFP all papers must contain an ethics section on the supplemental 14th page. By default, this section should detail the authors' stakeholder-based ethics analysis or should justify the use of an alternative approach to considering the ethics of their work.

All authors are expected to read this document in its entirety as it reflects USENIX Security's perspective on ethics. Authors are strongly encouraged to familiarize themselves with [The Menlo Report](https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf) (https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf) and the [2023 USENIX Security paper, "Ethical Frameworks and Computer Security Trolley Problems"](https://www.usenix.org/conference/usenixsecurity23/presentation/kohno) (https://www.usenix.org/conference/usenixsecurity23/presentation/kohno).

Note: this document is a starting point for ethical considerations, the guidelines herein are not exhaustive; authors are encouraged to think critically about ethics and not be limited to the guidelines in this document.

Stakeholder-Based Ethics Analysis

- **Stakeholders:** You are expected to consider all possible stakeholders (people, including the research team and society at large, and entities including companies) that may be impacted by your research. You are expected to detail how each stakeholder may have been impacted by the research procedures you undertook and how those stakeholders may be impacted by the publication of your research now and in the future.
- **Impacts:**
 - **Ethical principles:** You are expected to articulate the ethical principles you considered. A starting point is considering the principles in The Menlo Report in the context of each identified stakeholder: "Beneficence", "Respect for Persons", "Justice", and "Respect for Law and Public Interest".
 - **Harms:** There are at least two broad categories of potentially negative outcomes from the research and publication process: tangible harms (e.g., financial loss or exposure to psychologically disturbing content) and violations of human rights even if there are no directly tangible harms (e.g., the violation of a participants' right to informed consent or the violation of users' right to privacy via the study of data that users expect and desire to be private).
- **Mitigations:** You are expected to detail both mitigated and unmitigated (potential) harms of your work. You are expected to detail the steps taken to mitigate harms.
- **Decision:** You are expected to articulate why the decision to proceed with the research and the decision to publish the research was reached, respectively. One approach to reaching such a decision is to weigh harms against benefits (consequentialist ethics); see the "Beneficence" principle in The Menlo Report and the discussion of consequentialist ethics in the 2023 USENIX Security paper, "Ethical Frameworks and Computer Security Trolley Problems." An alternative or additional approach is to focus on avoiding the violation of individuals' (deontological ethics); see the "Respect for Persons" principle in The Menlo Report and the discussion of deontological ethics in the above-cited 2023 USENIX Security paper.
 - In some cases, the use of different principles in reaching an ethics decision will lead to the same conclusion for what is "right", e.g., the "Beneficence" and "Respect for Persons" analyses would agree. In other cases, the analyses may lead to different conclusions. If multiple analyses lead to the same conclusions, then documenting all those analyses will provide greater confidence in the ethics of the research. If different analyses lead to different conclusions on the ethics of the research, then the authors should clearly articulate how and why they chose the path they did, even if some principles would have led to a different decision. In some cases, researchers may need to make assumptions about the likelihood of different outcomes or the likely impacts of different decisions; in such cases, the authors are encouraged to articulate and justify all assumptions they make.
 - **When considering ethics, researchers and reviewers must acknowledge that, sometimes, the most ethical path is not to do the research or not to publish the research after it is complete.**

FAQ

When should I start considering Ethics?

As early as possible in the research process. By proactively considering ethics early, it is sometimes possible to avoid more challenging and complicated ethical questions in the future.

Is the goal to justify the decision I want to make?

No! Given that different approaches to ethical considerations can lead to different decisions, authors should not pick the decision that they want to make and then find the ethics argument that supports it. Rather, authors should be as objective as possible and ask themselves: How would someone not involved in the research evaluate the ethics of the research?

If I follow a procedure from past work, is saying that good enough?

No! Authors should also not simply look at past “similar” works and assume that the ethics analyses for those past works apply, directly, to their new works. Different situations may have subtle differences that, upon closer investigation, lead to significantly different conclusions. And, as the community learns more about the ethical implications of actions, and as technologies and societies and knowledge change, what might have been ethical in the past may no longer be ethical. For example, in the past, an action might have resulted in significant benefits that outweighed the harms but now, given knowledge from past results and/or differences in technologies, the benefits of the same action today might not outweigh the harms.

Is IRB approval a sufficient alternative?

No! IRBs are not expected to understand computer security research well or to know about best practices and community norms in our field, and so IRB approval does not absolve researchers from considering ethical aspects of their work.

Authors should leverage all available resources when making ethical decisions including the IRB, however IRB approval is not a substitute for ethical analysis nor is IRB approval sufficient to guarantee that the PC will not have additional concerns with respect to potential negative outcomes associated with the research. Hence, the discussion of IRB approval (if relevant) will likely only be a subset of the ethics statement.

If authors do not have access to an IRB but are doing human subjects-related research for which IRB approval might be required elsewhere and of other researchers, please clearly state that you do not have access to an IRB.

Domain-Specific Concrete Examples

Below, we consider in more depth several *examples* of ethical considerations that have come up in the past. These should be viewed as *examples*, however, and *not* an exhaustive list of potential concerns or considerations. And, as noted above, different situations, even if in many ways similar, may have unique considerations.

Disclosures. Vulnerabilities, if known to adversaries, can expose people to negative outcomes, such as harms or rights violations. Publicly disclosing vulnerabilities before they have been privately disclosed to the responsible parties, and hence before they have been mitigated, can therefore expose people to negative outcomes. Adversaries or others can also independently discover vulnerabilities. The potential for independent adversary discovery means that *knowing about vulnerabilities but not disclosing them* to the responsible parties can also result in exposing people to negative outcomes. Additionally, in some

cases it can take the responsible parties time to develop mitigations. Therefore, once a vulnerability has been discovered, *it is important to initiate the mitigation process as early as possible*. Specifically, absent strong and convincing reasons otherwise, we expect researchers to disclose vulnerabilities as soon as they are discovered. If the researchers believe that a different timeline is the most ethical in their situation, they should present clear and convincing arguments for that different timeline. The arguments should clearly articulate why a delayed disclosure is in the best interest of users or people in general, e.g., most supportive of these people’s wellbeing or least likely to violate their rights. *Submissions that fail to disclose prior to submission and that do not present convincing ethical arguments for delaying disclosure may be rejected*.

Often, the most direct path for vulnerability mitigation is to disclose the vulnerability to the responsible party, e.g. the manufacturer. In some cases, for example when the vulnerability is widespread or the mitigation process involves coordination with many organizations, the most ethical course of action may be to leverage organizations that coordinate vulnerability disclosure, such as the [CISA](https://www.cisa.gov/coordinated-vulnerability-disclosure-process) (<https://www.cisa.gov/coordinated-vulnerability-disclosure-process>) in the United States, rather than or in addition to disclosing to affected parties directly.

- Clearly state to whom you have disclosed vulnerabilities and the outcomes of those disclosures, if any.

Experiments with live systems without informed consent.

Researchers testing live services (e.g., for vulnerabilities) such as web services or APIs that give access to otherwise non-public algorithms or models must also consider ethics. Such experiments should only be performed after carefully analyzing the potential negative outcomes to the service provider, which may include cost (of CPU cycles or of human effort) or corrupting system state, and to end users who are using the same service provider for non-research purposes. That similar experiments might have been performed in the past does not automatically justify performing them again. Researchers should identify ways to minimize even small risks of negative outcomes, including by considering alternate methods (even if these are more difficult to carry out) and scaling down experiments.

- Papers describing such experiments should describe the analysis that lead to a particular methodology being used, justify its necessity and justify mitigations taken as well as what risks remained unmitigated and why it was not possible to mitigate them.

Terms of service. If experiments violate terms of service, the justification for violating them should be discussed in the ethics section.

Deception. In most cases, participants should be fully informed of the purposes and risks (among other things) of participating in experiments. If deception is used, the necessity of doing so should be carefully considered and the decision to use deception should be discussed in the ethics section. In general, participants in a deception study should be debriefed afterward to explain the necessity of the deception, even when the deception was mild.

Wellbeing for team members. In some cases, research activities have the potential to negatively impact team members. For example, research on hate speech could expose team members to disturbing content and negatively impact their psychological wellbeing. Or, crawling morally questionable websites from a home network could cause an ISP to (incorrectly) make inferences about the researcher that may not be true or that may be undesirable to the researcher.

- Research teams are expected to articulate how they considered the wellbeing of their researchers and the steps taken to mitigate identified risks as well as what risks remained unmitigated.

Innovations with both positive and negative potential outcomes.

Technologies that can positively impact one stakeholder group may negatively impact those same or other stakeholder groups. For example, advancements in anonymity systems could positively impact people that need anonymity under repressive regimes or excessive surveillance. At the same time, the mere use of those technologies could create negative impacts to those same people if the use of such technologies is detectable and hence subjects those individuals to additional scrutiny. And, those same anonymity technologies could also be used by illicit actors to conceal their activities. Likewise, advances in program analysis that can facilitate more rapid vulnerability finding could be used by both defenders and by adversaries. And, as an additional example, new insights into how and why some people become vulnerable to phishing could be used by both defenders and adversaries. Thus, researchers should think broadly about both the positive and negative potential impacts of their research throughout the research process, including during project selection and publication.

Retroactively identifying negative outcomes. While research teams should strive to proactively identify and address all ethics-related concerns before commencing their research and proactively address any new concerns that arise about the project's next steps during the research, in some cases research teams may discover post facto that their past research activities had unexpected and previously unknown (to the researchers) negative outcomes. Handling such situations is always difficult. While one might think that an appropriate response is to ignore and simply not talk about those past activities, doing so does not change the fact that negative outcomes did happen. In general, we believe that research teams should take ownership of any past negative outcomes that their research created, document such outcomes, and discuss what steps, if any, the researchers have taken to remediate those past negative outcomes and/or ensure that the potential for such negative outcomes are proactively addressed in the future, both for themselves and as guides for future researchers.

While this discussion provides a possible path forward for projects that retroactively identify negative outcomes and ethics-related concerns, simply following the suggestion above, making remedies, and documenting plans for the future does not guarantee that the PC will not have ethics-related concerns sufficient for paper rejection. Additionally, the following is explicitly considered unethical: identifying prior to the research that an activity might have negative outcomes, doing the activity anyway, and then documenting how researchers might avoid such negative outcomes in the future.

The law. In addition to considering ethics, we encourage authors to fully consider the legality of their research.

Symposium Organizers

Program Co-Chairs

Elissa Redmiles, Georgetown University
Ben Stock, CISPA Helmholtz Center for Information Security

Ethics Committee Co-Chairs

Eleanor Birrell, Pomona College
Tadayoshi Kohno, University of Washington

Artifact Evaluation Committee Co-Chairs

Aurore Fass, CISPA Helmholtz Center for Information Security
Deepak Kumar, University of California, San Diego
Christian Wressnegger, Karlsruhe Institute of Technology

Program Committee

TBA

Steering Committee

Michael Bailey, Georgia Institute of Technology
Kevin Butler, University of Florida
Joe Calandrino, Federal Trade Commission
Srdjan Capkun, ETH Zurich
William Enck, North Carolina State University
Rachel Greenstadt, New York University
Casey Henderson-Ross, USENIX Association
Nadia Heninger, University of California, San Diego
Thorsten Holz, Ruhr-Universität Bochum
Tadayoshi Kohno, University of Washington
Franziska Roesner, University of Washington
Kurt Thomas, Google
Patrick Traynor, University of Florida
Carmela Troncoso, EPFL