

Cybersecurity CFP Field Guide

Practical guidance for making the real work visible to a conference review committee.

What a reviewer needs to believe

A strong cybersecurity conference proposal lets a reviewer answer four questions quickly: What did you discover or build? How do you know? What is different from existing work? What will the audience be able to do with it?

The topic alone is not the contribution. "Cloud security," "AI attacks," or "supply-chain risk" names an area. Your proposal needs to name the result inside that area. Reviewers cannot award credit for research that remains implied.

This guide is independent and unofficial. It draws on public conference guidance and patterns observed across historical Black Hat CFP outcomes. It does not reproduce private submissions or individual reviews, and it does not speak for Black Hat or any other review board.

Lead with the result

An effective title names a target, technique, result, or useful tension. It gives the reviewer something more concrete than a broad category.

Weak: "Securing the Modern Cloud"

Stronger: "Cross-Tenant Secrets Through Mis-scoped Workload Identities"

The stronger title has not tried to sound dramatic. It tells the reader where the work lives and what failed.

The opening of the abstract should establish the claim before supplying background. In the first few sentences, identify the target, what you did, the result, and why it matters. If the reviewer must reach the outline to learn whether the research produced a finding, the abstract is not doing its job.

Use numbers where the work produced numbers. Name affected versions where versions matter. Say whether a tool, dataset, exploit, detection method, or proof of concept will be available. Never invent precision merely to sound complete.

Separate novelty from impact

A dramatic target can make ordinary work sound novel. Reviewers separate the two. A missing authorization check against a famous product may be important without introducing a new technique.

Name the closest prior work and finish this sentence: "Unlike that work, ours..." The answer may be a new primitive, scale, target class, dataset, exploit chain, measurement, or operational result. Honest incremental research can be strong when the increment is clear and useful.

Earlier publication or presentation requires context, not concealment. Identify what appeared, when, and what the proposed talk adds. Repeating the same recorded talk is different from extending a paper with new results or presenting a substantially changed technique.

Novelty cannot be proven by adjectives. "First ever," "unprecedented," and "novel" invite verification. Cite or describe the search you performed and narrow the claim to what the evidence supports.

Make the outline prove the talk exists

A detailed outline is evidence. It should contain information that could only come from someone who did the work.

Each major section should tell the reviewer what happens there: the target and version, the method, the observation, the result, or the lesson. "Background," "methodology," "demo," and "conclusion" are navigation labels, not content.

A useful sequence is:

- Scope and target: what was tested and why it matters.
- Method: how access, measurement, reversing, or experimentation worked.
- Finding: the mechanism, evidence, limitations, and affected versions.
- Demonstration: what the audience will see and the fallback if a live demo fails.
- Generalization: what transfers beyond one product or incident.
- Disclosure and current status.
- Audience takeaways.

Timings can demonstrate that the content fits the slot, but timings do not substitute for substance.

Show practical value

Cybersecurity programs can reward research that helps practitioners even when the underlying technique is incremental. State what an attendee can reproduce, test, detect, build, or decide afterward. "Understand the risks" is not yet a takeaway.

If the work produces a tool, connect it to the research. A talk whose main purpose is demonstrating a tool may belong in a tool-focused program such as Arsenal. A research talk can include a tool when the tool enables reproduction or applies the underlying technique.

Remove the sales motion

A commercial relationship is not automatically disqualifying. A product pitch is. Disclose the relationship and make the evidence carry the conclusion.

Replace claims such as "industry-leading protection" with measurements, design details, comparative results, or clearly stated limits. If the proposed defense happens to be the speaker's product, explain alternatives and conditions under which the product does not help.

Establish speaker credibility

Credibility means access to the work. It can come from discovering the issue, building the system, responding to the incident, maintaining the project, publishing related work, or operating the environment at meaningful scale.

Prior stage experience can help a board assess delivery, but first-time speakers should not imitate experience they do not have. A short, verifiable account of relevant work is stronger than a long list of titles.

Treat disclosure as part of the research

For new vulnerabilities, state who was notified, relevant dates, embargo status, patch availability, and what will be safe to present. Explain uncertainty. Coordinated disclosure demonstrates maturity and helps reviewers judge whether the material can be delivered at the event.

Use AI within the event's rules

Check the target event's current policy. As of October 2026, Black Hat's published terms prohibit LLM-generated submission text while allowing AI to edit or refine author-written material and help review prior art. Start from your own claims and prose. Use an assistant to find ambiguity, stress-test evidence, or tighten language. Review every change and remove anything the research does not support.

Before you submit

- The title names a target, technique, result, or useful tension.
- The abstract states the contribution in its opening sentences.
- The novelty statement identifies the closest work and the difference.
- The outline contains the actual method, evidence, and result.
- Takeaways use verbs such as test, reproduce, detect, build, or apply.
- Tools and demos support the research rather than replace it.
- Commercial relationships and previous publication are disclosed.
- Vulnerability status and affected versions are current.
- Every number, CVE, link, and accomplishment has been checked.
- The proposal matches the current CFP, track, format, and AI-use rules.

High scores and polished prose do not guarantee selection. Program balance, competition, timing, and expert judgment remain part of the decision. The purpose of a pre-check is to make sure the board is judging the real work rather than guessing what the submission meant.

Sources and currency

This edition was checked on October 2, 2026. Current Black Hat program descriptions, review criteria, tool-program guidance, coordinated-disclosure information, and submission terms are published at <https://blackhat.com/call-for-papers.html>. The event's current form and terms control when they differ from this guide.