Mythos attack on 3rd-round PQC algorithm candidate puts it out of commission

Mythos helped to find a new meet-in-the-middle technique that relies on a Möbius Bridge, a more sophisticated fingerprinting algorithm used in meet-in-the-middle attacks. Using it, Green said, the code Mythos produced was able to reduce the number of required inputs to 2^{89}. Anthropic said that savings can reduce the time required for such attacks by 200- to 800-fold.

The ability to produce that many inputs makes the attack beyond reach outside of the laboratory. Further, the actual speed-up is unknown, since the weakened AES algorithm tested used only seven rounds. Specification-compliant AES, Green said, uses 10, 12, or 14 rounds, depending on key size.

Anthropic is careful to explicitly spell out most of these caveats. The Monday blog post goes on to argue, however, that the results are nonetheless meaningful and could ultimately fundamentally disrupt the process of cryptanalysis, or the adversarial testing of cryptosystems.

“The cybersecurity community is now grappling with the fact that language models are able to discover so many bugs that the standard human processes (like vulnerability triage, verification, and remediation) struggle to keep up,” Anthropic wrote. “We predict that the same will soon be true in academic cryptography research. As language models increasingly produce novel research outputs autonomously, human researchers may become bottlenecked on studying and validating these results for technical validity, novelty, and utility.”

Not mentioned in Anthropic’s report is whether its researchers used Mythos to attack more tested cryptosystems, such as elliptic curve cryptography and RSA. Attack improvements against these systems would be more impressive. By achieving the most impressive result against an algorithm still in its infancy, it’s not clear how much of an advantage Mythos truly provided. There’s no way of knowing if researchers using conventional cryptanalysis techniques were already close to discovering the same attack.

Ultimately, the lesson from the research is simple. AI-assisted cryptanalysis remains untested, and providers of these platforms have a vested interest in exaggerating their benefits. At the same time, there’s growing evidence that LLMs may provide significant advantages in finding cryptographic weaknesses. It would be a mistake to conclude that LLMs won’t one day play an important role in the race between securing and compromising our most vital assets.

The headline and body of this story have been updated to reflect the withdrawing of HAWK.

https://arstechnica.com/security/2026/07/mythos-uncovers-crypto-weaknesses-that-went-unknown-for-years/