Nomination statement

Schorlemmer, Taylor R., Kelechi G. Kalu, Luke Chigges, Kyung Myung Ko, Eman Abu Ishgair, Saurabh Bagchi, Santiago Torres-Arias, and James C. Davis. "Signing in four public software package registries: Quantity, quality, and influencing factors." In 2024 IEEE Symposium on Security and Privacy (SP), pp. 1160-1178. IEEE, 2024.

This exemplary empirical study of signing practices for open source software deserves strong consideration by the Distinguished Experts panel for this year's Best Science of Cyber Security paper award. Open source software is a significant component of many critical systems, and there is evidence that attackers have tried to insinuate malicious software into open source software supply chains. Without proper signatures on this software, provenance of open source software is impossible to assure, and a Software Bill of Materials (SBOM) loses its value. Consequently, understanding the quantity and quality of signed software in open source repositories, and the factors that influence the quantity and quality of signed open source software is a major research question. This paper provides both a model for how to conduct this kind of research and significant results that can (and should!) be used by those who develop open source software and those who operate repositories to move toward significantly increasing the fraction of open source software that has valid, verifiable signatures.

The paper’s combination of large-scale empirical data and causal inference methods offers both descriptive insights and explanatory power. The paper provides seven major findings. (1) Unless a registry mandates signing, engineers tend not to sign their packages. (2) When engineers do sign their packages, they often do so incorrectly – the paper reports signing failure rates of 24-76% in the Maven, PyPI, and HuggingFace registries. The most common failure mode is expired public keys; public key infrastructure remains problematic. (3) When registries provide dedicated tooling for signing, it does not cause more signing, but it does improve the quality of the resulting signatures. (4) Major cybersecurity events – software supply chain attacks, executive orders, and NIST guidance – have no statistically significant effect on software signing practices (neither quantity nor quality).

The results presented are significant in themselves: the authors carefully formulate their research questions, and discuss their results as well as threats to the validity of the work. In brief, it appears that good signing tools and a repository policy of requiring signatures can in fact promote quantity and quality of signatures but that external events, such as significant cybersecurity events, have had relatively little effect. The work not only stands on its own as a significant contribution but provides a strong foundation for extensions and future empirical work of this kind.

I wholeheartedly recommend this paper for the award.