

2026 IEEE Symposium on Security and Privacy (SP 2026)

**San Francisco, California, USA
18-21 May 2026**

Pages 1–674

IEEE Catalog Number: CFP26020-POD
ISBN: 979-8-3315-6066-9



Copyright © 2026, IEEE

All Rights Reserved

Copyright and Reprint Permissions:

Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

*** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.

IEEE Catalog Number:	CFP26020-POD
ISBN (Print-On-Demand):	979-8-3315-6066-9
ISBN (Online):	979-8-3315-6065-2
ISSN:	1081-6011

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA

Phone:	(845) 758-0400
Fax:	(845) 758-2633
E-mail:	curran@proceedings.com
Web:	www.proceedings.com

TABLE OF CONTENTS

LLM Unlearning Should Be Form-Independent	1
<i>Xiaotian Ye, Mengqi Zhang, Shu Wu</i>	
The Person Behind the Sound: Demystifying Audio Private Attribute Profiling Via Multimodal Large Language Models	19
<i>Lixu Wang, Kaixiang Yao, Xinfeng Li, Dong Yang, Haoyao Li, XiaoFeng Wang, Wei Dong</i>	
StepStone: LLM-Based GPU Kernel Driver Fuzzing via User-Space Libraries	39
<i>Xiaochen Zou, Juefei Pu, Arrdya Srivastav, Jonathan Cox, Zhengchuan Liang, Yuan Tan, Xingyu Li</i> <i>Yilin Zhu, Zhiyun Qian</i>	
Keytar: Practical Keystroke Timing Attacks and Input Reconstruction	57
<i>Mufan Qiu, Lihuan Chuang, Dohhyun Kim, Huaizhi Qu, Tianlong Chen, Andrew Kwong</i>	
Blinding Post-Quantum Hash-and-Sign Signatures	77
<i>Charles Bouillaguet, Thibault Feneuil, Jules Maire, Matthieu Rivain, Julia Sauvage, Damien Vergnaud</i>	
SoK: All You Ever Wanted to Know About Bootloader Security but Were Afraid to Ask	95
<i>Connor Glosner, Aravind Machiry</i>	
Descriptors of Exposure: Undermining Tor Anonymity Through Exploiting Descriptor Flood	117
<i>Chunmian Wang, Junzhou Luo, Zhen Ling, Yue Zhang, Shan Wang, Ming Yang, Guangchi Liu, Xinwen Fu</i>	
VerfCNN, Optimal Complexity zkSNARK for Convolutional Neural Networks	135
<i>Wenjie Qu, Yanpei Guo, Yue Ying, Jiaheng Zhang</i>	
DY* Unchained: Now with Composable Security Proofs and Precise Compromise Scenarios	153
<i>Th��ophile Wallez</i>	
One Tap to Hijack Them All: A Security Analysis of the Google Fast Pair Protocol	171
<i>Sayon Duttagupta, Seppe Wyns, Nikola Antonijevi��, Dave Sing����, Bart Preneel</i>	
Responsible Disclosure is a Two-Way Street: Empirically Measuring the Responsible Disclosure Contract in the Firmware Ecosystem	189
<i>Hui Jun Tay, Souradip Nath, Arvind S Raj, Abhay Bhat, Ishan Bansal, Audrey Dutcher, Moritz Schloegel</i> <i>Adam Doup��, Tiffany Bao, Yan Shoshitaishvili, Ruoyu Wang</i>	
Fuzzing the Physical Space: Physics-Aware Testing of Black-Box Industrial Control Systems	207
<i>Burak Sahin, David Oyg��nblik, Mingxuan Yao, Yizhi Huang, Brendan Saltaformaggio, Saman Zonouz</i>	
SoK: Critical Evaluation of Quantum Machine Learning for Adversarial Robustness	225
<i>Saeefa Rubaiyet Nowmi, Jesus Lopez, Md Mahmudul Alam Imon, Shahrooz Pouryousef</i> <i>Mohammad Saidur Rahman</i>	
Sok: Evaluating Jailbreak Guardrails for Large Language Models	245
<i>Xunguang Wang, Zhenlan Ji, Wenxuan Wang, Zongjie Li, Daoyuan Wu, Shuai Wang</i>	
URLcoat: Exploiting Web Search Capability to Jailbreak Large Language Models	265
<i>Yiheng Sun, Linkang Du, Zhou Su, Yuntao Wang, Han Liu</i>	
MetaBreak: Jailbreaking Online LLM Services via Special Token Manipulation	284
<i>Wentian Zhu, Zhen Xiang, Wei Niu, Le Guan</i>	

SoK: Robustness in Large Language Models against Jailbreak Attacks	304
<i>Feiyue Xu, Hongsheng Hu, Chaoxiang He, Sheng Hang, Hanqing Hu, Xiuming Liu, Yubo Zhao Zhengyan Zhou, Bin Benjamin Zhu, Shi-Feng Sun, Dawu Gu, Shuo Wang</i>	
Parasites in the Toolchain: A Large-Scale Analysis of Attacks on the MCP Ecosystem	324
<i>Shuli Zhao, Qinsheng Hou, Zihan Zhan, Yanhao Wang, Yuchong Xie, Yu Guo, Libo Chen, Shenghong Li Zhi Xue</i>	
Ensemble Conformal Predictor (EnCP): A New Conformal Predictor with Robustness Guarantees Against Data Poisoning Attacks	342
<i>Yuxin Yang, Qiang Li, Runyang Feng, Liren Shan, Binghui Wang</i>	
Exploiting Leaderboards for Large-Scale Distribution of Malicious Models	361
<i>Anshuman Suri, Harsh Chaudhari, Yuefeng Peng, Ali Naseh, Alina Oprea, Amir Houmansadr</i>	
GraphRAG Under Fire	381
<i>Jiacheng Liang, Yuhui Wang, Changjiang Li, Tanqiu Jiang, Rongyi Zhu, Neil Gong, Ting Wang</i>	
AI Wrote My Paper and All I Got was This False Negative:* Measuring the Efficacy of Commercial AI Text Detectors	399
<i>Seth Layton, Bernardo B. P. Medeiros, Kevin Butler, Patrick Traynor</i>	
Your Compiler is Backdooring Your Model: Understanding and Exploiting Compilation Inconsistency Vulnerabilities in Deep Learning Compilers	419
<i>Simin Chen, Jinjun Peng, Yixin He, Junfeng Yang, Baishakhi Ray</i>	
DREAM: Scalable Red Teaming for Text-to-Image Generative Systems via Distribution Modeling	438
<i>Boheng Li, Junjie Wang, Yiming Li, Zhiyang Hu, Leyi Qi, Jianshuo Dong, Run Wang, Han Qiu, Zhan Qin Tianwei Zhang</i>	
On the (In)Security of Loading Machine Learning Models	458
<i>Gabriele Digregorio, Marco Di Gennaro, Stefano Zanero, Stefano Longari, Michele Carminati</i>	
Evaluating Concept Filtering Defenses against Child Sexual Abuse Material Generation by Text-to-Image Models	476
<i>Ana-Maria Crețu, Klim Kireev, Amro Abdalla, Wisdom Obinna, Raphael Meier, Sarah Adel Bargal Elissa M. Redmiles, Carmela Troncoso</i>	
Hijacking Large Audio-Language Models via Context-Agnostic and Imperceptible Auditory Prompt Injection	496
<i>Meng Chen, Kun Wang, Li Lu, Jiaheng Zhang, Tianwei Zhang</i>	
Adversarial Hubness in Multi-Modal Retrieval	515
<i>Tingwei Zhang, Fnu Suyu, Rishi Jha, Collin Zhang, Vitaly Shmatikov</i>	
Recovering and Rehosting Mobile Local LLM Conversations and Contexts via Memory Forensics	531
<i>Haichuan Xu, David Oygenblik, Runze Zhang, Mingxuan Yao, Muhammad Ibrahim Brendan Saltaformaggio</i>	
Nonlocalizable Jamming with Curving Beams	550
<i>Caroline Jane Spindel, Edward Knightly</i>	
WRATH: Turning Watermark Robustness Against Itself via a Watermark-Agnostic Black-Box Invalidation Attack	565
<i>Nan Jiang, Juan Hu, Bangjie Sun, Terence Sim, Jun Han</i>	

Are LLM-Enhanced Graph Neural Networks Robust Against Poisoning Attacks?	584
<i>Yuhang Ma, Jie Wang, Zheng Yan</i>	
Breaking Free from Ivory Tower: Evaluating and Enhancing Real-world Chinese Underground Adversarial Jargon Detection	603
<i>Zhifan Jiang, Mingxuan Liu, Yue Qin, Baojun Liu</i>	
WebCloak: Characterizing and Mitigating Threats From LLM-Driven Web Agents as Intelligent Scrapers	622
<i>Xinfeng Li, Tianze Qiu, Yingbin Jin, Lixu Wang, Hanqing Guo, Xiaojun Jia, XiaoFeng Wang, Wei Dong</i>	
GHost in the Shell: A GPU-to-Host Memory Attack and its Mitigation	642
<i>Sihyun Roh, Woohyuk Choi, Jaeyoung Chung, Yoochan Lee, Suhwan Song, Byoungyoung Lee</i>	
Demystifying and Exploiting ASLR on NVIDIA GPUs	658
<i>Ruofan Zhu, Ganhao Chen, Wenbo Shen, Lyuye Zhang, Dakun Shen, Rui Chang, Yanan Guo</i>	
Phoenix: Rowhammer Attacks on DDR5 with Self-Correcting Synchronization	675
<i>Diego Meyer, Patrick Jattke, Michele Marazzi, Salman Qazi, Daniel Moghimi, Kaveh Razavi</i>	
GPUBreach: Privilege Escalation Attacks on GPUs Using Rowhammer	694
<i>Chris S. Lin, Yuqin Yan, Guozhen Ding, Joyce Qu, Joseph Zhu, David Lie, Gururaj Saileshwar</i>	
GDDRHammer: Greatly Disturbing DRAM Rows — Cross-Component Rowhammer Attacks From Modern GPUs	714
<i>Yichang Hu, Noah Brown, Yuhang Chen, Joshua Bakita, Tianlong Chen, Daniel Genkin, Andrew Kwong</i>	
GeForge: Hammering GDDR Memory to Forge GPU Page Tables for Fun and Profit	733
<i>Junpeng Wan, Yanan Guo, Zhi Zhang, Zhuo Li, Dave Jing Tian, Zhenkai Zhang</i>	
Defeating Transient Execution Attacks by Limiting Secret Reachability Through Register Hiding and ShadowCFI	752
<i>Daniël Trujillo, Jagadish Kotra, David Kaplan, Mengjia Yan</i>	
Transient Architectural Execution: From Weird Gates to Weird Programs	771
<i>Ping-Lun Wang, Fraser Brown, Riccardo Paccagnella, Eyal Ronen, Riad S. Wahby, Yuval Yarom</i>	
SeqAss: Using SeqUential Associative Caches to Mitigate Conflict-Based Cache Attacks with Reduced Cache Misses and Performance Overhead	788
<i>Wei Song, Zhidong Wang, Jinchu Han, Da Xie, Hao Ma, Peng Liu</i>	
Trevex: A Black-Box Detection Framework for Data-Flow Transient Execution Vulnerabilities	806
<i>Daniel Weber, Fabian Thomas, Leon Trampert, Ruiyi Zhang, Michael Schwarz</i>	
AESpoly: Symmetric-Key Cryptographic Designs Using Instruction-Level Parallelism Between AES and Polynomial Hash	826
<i>Yukihito Hiraga, Yusuke Naito, Yu Sasaki, Takeshi Sugawara</i>	
Crucible: Retrofitting Commodity CPUs with Vulnerabilities via Transparent Software Emulation	846
<i>Tristan Hornetz, Lukas Gerlach, Michael Schwarz</i>	
RISCy Cache Coherence: Timer-Free Architectural Cache Attacks via Instruction/Data Cache Incoherence	862
<i>Fabian Thomas, Michael Schwarz</i>	

“I Wonder if These Warnings are Accurate”: Security and Privacy Advice in Nine Majority World Countries	880
<i>Collins W. Munyendo, Veronica A. Rivera, Jackie Hu, Emmanuel Tweneboah, Amna Shahnawaz Karen Sowon, Dilara Keküllüoğlu, Marcos Silva, Yue Deng, Mercy Omeiza Gayatri Priyadarsini Kancherla, Maria Rosario Niniz Silva, Abhishek Bichhawat, Maryam Mustafa Francisco Marmolejo-Cossio, Elissa M. Redmiles, Yixin Zou</i>	
Privacy Perspectives and Practices of Chinese Smart Home Product Teams	900
<i>Shijing He, Yaxiong Lei, Xiao Zhan, Chi Zhang, Juan Ye, Ruba Abu-Salma, Jose Such</i>	
Searching for a Farang: Collective Security Among Women in Pattaya, Thailand	919
<i>Taylor Robinson, Rikke Bjerg Jensen</i>	
Toward Inclusive Security and Privacy for Deaf and Hard-of-Hearing People: A Community-Based Interview Study	939
<i>Mindy Tran, Xinru Tang, Adryana Hutchinson, Adam J. Aviv, Yixin Zou</i>	
XDup: Privacy-Preserving Deduplication for Humanitarian Organizations Using Fuzzy PSI	959
<i>Tim Rausch, Sylvain Chatel, Wouter Lueks</i>	
Human-Centered Threat Modeling in Practice: Lessons, Challenges, and Paths Forward	978
<i>Warda Usman, Yixin Zou, Daniel Zappala</i>	
LISA: A Scale-Optimized and Psychometrically-Validated Instrument for the Lightweight Assessment of Organizational Information Security Awareness in Heterogeneous Organizations	997
<i>David Langer, Jan Tolsdorf, Luigi Lo Iacono</i>	
Perceived Privacy Risk and Mitigation Post-Roe	1017
<i>Alan F. Luo, Phoebe Moh, Cora Sula, Michelle L. Mazurek, Nora McDonald</i>	
Breaking the Illusion: Automated Reasoning of GDPR Consent Violations	1035
<i>Ying Li, Wenjun Qiu, Faysal Hossain Shezan, Kunlin Cai, Michelangelo Van Dam, Lisa Austin, David Lie Yuan Tian</i>	
Understanding and Analyzing Privacy Risks in Mobile Consent-Management Platforms	1054
<i>Jingzhou Ye, Fares Alharbi, Luyi Xing, Xueqiang Wang</i>	
Convenience at a Cost: the Security Risks of Template-Based Development in the App-in-App Ecosystem	1073
<i>Yizhe Shi, Zhemin Yang, Yifan Yang, Yunteng Yang, Yifan Yang</i>	
When VR Meets BCI: (Un)Observable Brainwave-Aware Privacy Reconstruction in the Metaverse via Unrestricted Inbuilt Motion Sensors	1091
<i>Tao Ni, Zehua Sun, Qingchuan Zhao, Wei-Bin Lee, Cong Wang</i>	
Secret State Leakage Attacks and Their Impacts on EMV Contactless Payment Apps	1110
<i>Jesse Chen, Rubin Yuchan Yang, Ahmad Musa, Syed Rafiul Hussain, Omar Chowdhury, Sazzadur Rahaman</i>	
Navigating Developers' Quagmire: LLM-Enabled Privacy Compliance Analysis for SDK Integrations	1129
<i>Zhaojie Hu, Xueqiang Wang</i>	
LLMThief: Evaluating Configuration Leaking Risks in Commercial LLM App Stores	1149
<i>Pinji Chen, Jinlong Jiang, Jianjun Chen, Feiran Qin, Minghao Zhang, Jiahe Zhang, Haixin Duan Kaiwen Shen, Hui Jiang</i>	

LEAKYLINKS: Measuring the Security and Privacy Risks of URL Scanning Services	1168
<i>Ali Mustafa, Jannis Rautenstrauch, Florian Hantke, Shubham Agarwal, Stefano Calzavara, Ben Stock</i>	
Concretely-Efficient Multi-Key Homomorphic Secret Sharing and Applications	1188
<i>Kaiwen He, Sacha Servan-Schreiber, Geoffroy Couteau, Srinivas Devadas</i>	
No Honor Among Crooks: Non-Transferable Anonymous Tokens from Betrayability	1208
<i>David Kretzler, Yong Li</i>	
Revisiting PQ Wireguard: A Comprehensive Security Analysis with a New Design Using Reinforced KEMs	1228
<i>Keitaro Hashimoto, Shuichi Katsumata, Guilhem Niot, Thom Wiggers</i>	
Can I Get More? An Incremental Inference Attack on Encrypted SQL	1248
<i>Xiaoqian Sun, Ruiqi He, Yang Zhang, Siyi Lv, Guiyun Qin, Fangzhou Yi, Zheli Liu, Xiaofeng Chen</i>	
Starfighters-On the General Applicability of X-Wing	1268
<i>Deirdre Connolly, Kathrin Hövelmanns, Andreas Hülsing, Stavros Kousidis, Matthias Meijers</i>	
A Leakage-Free Framework for Private Set Operations	1287
<i>Wenhao Wu, Yuyue Chen, Bowen Shen, Peng Yang, Ximing Fu, Zoe Lin Jiang, Junbin Fang</i>	
Hardware Trojans from Invisible Inversions: On the Trojanizability of Standard Cell Libraries	1306
<i>Kolja Dorschel, René Walendy, Lukas Plätz, Thorben Moos, Christof Paar, Steffen Becker</i>	
A Maliciously-Secure Post-Quantum OPRF from Crypto Dark Matter	1325
<i>Diego F. Aranha, Aron van Baarsen, Adam Blatchley Hansen, Kent Nielsen, Peter Scholl</i>	
Scalable Accountable Byzantine Agreement and Beyond	1345
<i>Pierre Civit, Daniel Collins, Vincent Gramoli, Rachid Guerraoui, Jovan Komatovic, Manuel Vidigueira Pouriya Zarbafian</i>	
Practical Asynchronous Distributed Key Reconfiguration and Its Applications	1365
<i>Hanwen Feng, Yingzi Gao, Yuan Lu, Qiang Tang, Jing Xu</i>	
Lattice-Based Threshold Blind Signatures	1385
<i>Sebastian Faller, Guilhem Niot, Michael Reichle</i>	
A Full Threshold NIST PQC-Compliant Framework for Distributed Trust in Federal Public Key Infrastructure	1405
<i>Kiarash Sedghighadikolaei, Changqi Sun, Thang Hoang, Bechir Hamdaoui, Attila A. Yavuz</i>	
Mad-Dag: Protecting Blockchain Consensus From MEV	1425
<i>Roi Bar-Zur, Ittay Eyal, Aviv Tamar</i>	
Robot: Robust Threshold BBS+ in Two Rounds	1444
<i>Guofeng Tang, Tian Qiu, Bowen Jiang, Haiyang Xue, Guomin Yang, Man Ho Au, Robert H. Deng Kwok-Yan Lam</i>	
New Constructions of Functional Adaptor Signatures: Broader Functions and Improved Efficiency	1462
<i>Nikhil Vanjani, Garrett Greiner, Sri AravindaKrishnan Thyagarajan, Pratik Soni</i>	
Practical Multi-Party Private Set Intersection with Reducible Zero-Sharing	1481
<i>Yewei Guan, Hua Guo, Man Ho Au, Jiarong Huo, Jin Tan, Zhenyu Guan</i>	

Single-Server Stateful PIR with Verifiability and Balanced Efficiency	1497
<i>Pranav Shriram Arunachalam, Ling Ren</i>	
ZELDA: Efficient Multi-Server Preprocessing PIR With Unconditional Security	1515
<i>Ashrut Ghoshal, Mingxun Zhou, Bo Peng, Elaine Shi</i>	
Verifiable PIR with Small Client Storage	1535
<i>Mayank Rathee, Keewoo Lee, Raluca Ada Popa</i>	
Euston: Efficient and User-Friendly Secure Transformer Inference with Non-Interactivity	1555
<i>Xinwen Gao, Shaojing Fu, Lin Liu, Zhuotao Liu, Yuchuan Luo, Yongjun Wang</i>	
InsPIRe: Communication-Efficient PIR with Server-Side Preprocessing	1575
<i>Rasoul Akhavan Mahdavi, Sarvar Patel, Joon Young Seo, Kevin Yeo</i>	
LatORAM: ORAMs from Lateral Stashes and Delayed Shuffling	1595
<i>Sarvar Patel, Giuseppe Persiano, Joon Young Seo, Kevin Yeo</i>	
VIA: Communication-Efficient Single-Server Private Information Retrieval	1615
<i>Chenyang Liu, Xukun Wang, Zhifang Zhang</i>	
Automated Formal Analysis of Signal's Double Ratchet: Attacks, Fixes and Security Proofs	1634
<i>Vincent Cheval, Charlie Jacomme, Jessica Richards</i>	
Banshee: Target Switch Attacks on Gimbal-Stabilized Visual Tracking Systems via Acoustic Injection	1651
<i>Jiarui Li, Joseph Brewington, Qingzhao Zhang, Z. Morley Mao</i>	
Rain: Transiently Leaking Data from Public Clouds Using Old Vulnerabilities	1670
<i>Mathé Hertogh, Dave Quakkelaar, Thijs Raymakers, Mahesh Hari Sarma, Marius Muench, Herbert Bos, Erik Van der Kouwe</i>	
Chypnosis: Undervolting-based Static Side-channel Attacks	1687
<i>Kyle Mitard, Saleh Khalaj Monfared, Fatemeh Khojasteh Dana, Robert Dumitru, Yuval Yarom, Shahin Tajik</i>	
SoK: Systematizing a Decade of Architectural Rowhammer Defenses Through the Lens of Streaming Algorithms	1706
<i>Michael Jaemin Kim, Seungmin Baek, Jumin Kim, Hwayong Nam, Nam Sung Kim, Jung Ho Ahn</i>	
RadKey: An LLM-Guided RF Backscatter System for Through-Wall Keystroke Inference	1725
<i>Qijun Wang, Chunqi Qian, Huacheng Zeng</i>	
Beyond Indistinguishability: Measuring Extraction Risk in LLM APIs	1744
<i>Ruixuan Liu, David Evans, Li Xiong</i>	
Agentic Concolic Execution	1763
<i>Zhengxiong Luo, Huan Zhao, Dylan Wolff, Cristian Cadar, Abhik Roychoudhury</i>	
C-Verifier: Understanding and Formally Verifying Cross-Service Flaws in AWS Cognito	1782
<i>Zhen Chen, Ze Jin, Le Gong, Kexin Chen, Xiangyi Zeng, Qixu Liu</i>	
VMSCAPE: Exposing and Exploiting Incomplete Branch Predictor Isolation in Cloud Environments	1800
<i>Jean-Claude Graf, Sandro Rüegg, Ali Hajiabadi, Kaveh Razavi</i>	
Practical Covert Channel Across Isolated Browser Instances via GPU Command Queue Contention	1818
<i>Jinhong Liu, Zifeng Kang, Song Li, Yinzhi Cao</i>	

Detecting Privilege Escalation in Polyglot Microservices via Agentic Program Analysis	1837
<i>Penghui Li, Hong Yau Chong, Yinzhi Cao, Junfeng Yang</i>	
State of Browser Process-Isolation: The Same-Site Weakness	1856
<i>Fabian Kilger, Hannah Fischer, Adrian Staeves, Robin Marchart, Josef Schönberger, Fabian Franzen</i>	
KeyChaser: Unveiling API Keys in Browser Extensions	1875
<i>Shijin Chen, Willy Susilo, Yudi Zhang, Fuchun Guo</i>	
Site Isolation is Dead: How Site Isolation is Broken in Agentic Browsers and Extensions	1894
<i>Suyoung Lee, Seongho Keum, Changoo Lee, Dongwon Shin, Sanghyun Hong, Byoungyoung Lee, Sooel Son</i>	
The Secrets Must Not Flow: Scaling Security Verification to Large Codebases	1912
<i>Linard Arqint, Samarth Kishor, Jason R. Koenig, Joey Dodds, Daniel Kroening, Peter Müller</i>	
Mechanized Safety and Liveness Proofs for the Mysticeti Consensus Protocol Under the LiDO-DAG Framework	1932
<i>Longfei Qiu, Jingqi Xiao, Zhong Shao</i>	
Towards Practical Zero-Knowledge Proof for PSPACE	1952
<i>Ashwin Karthikeyan, Hengyu Liu, Kuldeep S. Meel, Ning Luo</i>	
Coral: Fast Succinct Non-Interactive Zero-Knowledge CFG Proofs	1972
<i>Sebastian Angel, Sofia Celi, Elizabeth Margolin, Pratyush Mishra, Martin Sander, Jess Woods</i>	
Dory: Streaming PCG with Small Memory	1992
<i>Xiaojie Guo, Hanlin Liu, Zhicong Huang, Hongrui Cui, Wenhao Zhang, Cheng Hong, Xiao Wang</i>	
<i>Kang Yang, Yu Yu</i>	
CAVERN: Efficient Honest-Majority Maliciously Secure $(2+1)$ -PC for $\mathbb{Z}_{2^n}$ via DPF	2011
<i>Yang Liu, Liang Feng Zhang</i>	
Sort, Sweep, Mirror: Batch Private Interval Lookup with Logarithmic Cost	2027
<i>Andes Y. L. Kei, Lucien K. L. Ng, Jack P. K. Ma, Sherman S. M. Chow</i>	
Vega: Low-Latency Zero-Knowledge Proofs over Existing Credentials	2041
<i>Darya Kaviani, Srinath Setty</i>	
The Pipes Model for Latency and Throughput Analysis	2060
<i>Andrew Lewis-Pye, Kartik Nayak, Nibesh Shrestha</i>	
Generate-then-Verify: Reconstructing Data from Limited Published Statistics	2079
<i>Terrance Liu, Eileen Xiao, Adam Smith, Pratiksha Thaker, Zhiwei Steven Wu</i>	
Setting the Course, but Forgetting to Steer: Analyzing Compliance with GDPR's Right of Access to Data by Instagram, TikTok, and Youtube	2094
<i>Sai Keerthana Karnam, Abhisek Dash, Antariksh Das, Sepehr Mousavi, Stefan Bechtold</i>	
<i>Krishna P. Gummadi, Animesh Mukherjee, Ingmar Weber, Savvas Zannettou</i>	
Practical Anonymous Two-Party Gradient Boosting Decision Tree	2113
<i>Chenyu Huang, Fan Zhang, Minxin Du, Sherman S. M. Chow, Huangxun Chen, Huaming Rao</i>	
<i>Danqing Huang, Bo Qian, Peng Chen</i>	

Consumer Beware! Exploring Data Brokers' CCPA Compliance	2132
<i>Elina van Kempen, Isita Bagayatkar, Pavel Frolikov, Chloe Georgiou, Gene Tsudik</i>	
Private Data Imputation	2147
<i>Addelkarim Kati, Florian Kerschbaum, Marina Blanton</i>	
It's a Feature, Not a Bug: Secure and Auditable State Rollback for Confidential Cloud Applications	2166
<i>Quinn Burke, Anjo Vahldiek-Oberwagner, Michael Swift, Patrick McDaniel</i>	
Understanding Data Collection, Brokerage, and Spam in the Lead Marketing Ecosystem	2186
<i>Yash Vekaria, Nurullah Demir, Konrad Kollnig, Zubair Shafiq</i>	
CBUE: Conclusion Based Utility Evaluation for Differentially Private Categorical Data	2206
<i>Furkan Sarikaya, Shaohua Lu, Johes Bater, Mark Hempstead</i>	
Fast Deterministically Safe Proof-of-Work Consensus	2223
<i>Ali Farahbakhsh, Giuliano Losa, Youer Pu, Lorenzo Alvisi</i>	
Hadal: Centralized Label DP without a Trusted Party	2242
<i>James Choncholas, Stanislav Peceny, Amit Agarwal, Mariana Raykova, Baiyu Li, Karn Seth</i>	
Privacy-Conscious Algorithm Design Via PAC Privacy	2261
<i>Mayuri Sridhar, Xiaochen Zhu, Srinivas Devadas</i>	
Decomposition-Based Optimal Bounds for Privacy Amplification via Shuffling	2281
<i>Pengcheng Su, Haibo Cheng, Ping Wang</i>	
Making Privacy Public: Toward a Differential Privacy Deployment Registry	2301
<i>Priyanka Nanayakkara, Elena Ghazi, Salil Vadhan</i>	
Auditing Apple's DifferentialPrivacy.framework: Implementation Bugs, Misconfigurations, and Practical Risks	2320
<i>Rishav Chourasia, Ergute Bao, Uzair Javaid, Xiaokui Xiao</i>	
Sparse Estimation Under Local Differential Privacy at All Privacy Levels	2340
<i>Puning Zhao, Qingqing Ye, Shaowei Wang, Jun Feng, Sheng Yue, Zhen Chen, Xiaochun Cao</i>	
Shared Spotlight Meridian: Distributed Sparse Pseudorandom Functions for Scalable Federated Learning	2359
<i>Youlong Ding, Peihua Mai, Jingqi Zhang, Sherman S. M. Chow, Minxin Du, Yan Pang</i>	
Jigsaw: Doubly Private Smart Contracts	2378
<i>Sanjam Garg, Aarushi Goel, Dimitris Kolonelos, Rohit Sinha</i>	
CHORUS: Secret Recovery with Ephemeral Client Committees	2397
<i>Deevashwer Rathee, Emma Dauterman, Allison Li, Raluca Ada Popa</i>	
ARES: Scalable and Practical Gradient Inversion Attack in Federated Learning Through Activation Recovery	2417
<i>Zirui Gong, Leo Yu Zhang, Yanzun Zhang, Viet Vo, Tianqing Zhu, Shirui Pan, Cong Wang</i>	
On the Detectability of Active Gradient Inversion Attacks in Federated Learning	2436
<i>Vincenzo Carletti, Pasquale Foggia, Carlo Mazzocca, Giuseppe Parrella, Mario Vento</i>	

Toward Efficient Membership Inference Attacks Against Federated Large Language Models: A Projection Residual Approach	2456
<i>Guilin Deng, Silong Chen, Yuchuan Luo, Yi Liu, Songlei Wang, Zhiping Cai, Lin Liu, Xiaohua Jia Shaojing Fu</i>	
STIR/SHAKEN: A Cocktail of Cryptographic Clumsiness	2475
<i>Joshua Brown, Paul Grubbs, Matthew Hardeman</i>	
A Liveness Attack to Ethereum PoS with No Additional Cost	2495
<i>Mingfei Zhang, Rujia Li, Xueqian Lu, Sisi Duan</i>	
Sealing the Window: Efficient Tamper Protection for Provenance Logs*	2514
<i>Sagar Mishra, R. Sekar</i>	
Efficient Fuzzy Private Set Intersection from Secret-Shared OPRF	2532
<i>Xinpeng Yang, Meng Hao, Chenkai Weng, Robert H. Deng, Yonggang Wen, Tianwei Zhang</i>	
Weighted Batched Threshold Encryption With Applications to Mempool Privacy	2552
<i>Amit Agarwal, Kushal Babel, Sourav Das, Babak Poorebrahim Gilakaye, Arup Mondal, Benny Pinkas Peter Rindal, Aayush Yadav</i>	
GoSSamer: Lightweight and Linear-Communication Asynchronous (Dynamic Proactive) Secret Sharing and the Applications	2572
<i>Xinxin Xing, Yizhong Liu, Boyang Liao, Jianwei Liu, Bin Hu, Xun Lin, Yuan Lu, Tianwei Zhang</i>	
UltraProofs: Scalable Reed-Solomon Code Commitment	2592
<i>Yanpei Guo, Alex Luoyuan Xiong, Wenjie Qu, Jiaheng Zhang</i>	
Scalable Registration-Based Encryption from Lattices	2610
<i>Michael Klooß, Russell W. F. Lai, Jan Niklas Siemer, Monisha Swarnakar</i>	
From Perfect to Approximate Hints: Efficient LWE Secret Recovery Leveraging Low Hamming Weight	2629
<i>Minki Hhan, Ga Hee Hong, Jiseung Kim, Changmin Lee, JeongHwan Lee</i>	
International Students and Scams: At Risk Abroad	2644
<i>Katherine Zhang, Arjun Arunasalam, Pubali Datta, Z. Berkay Celik</i>	
Lost in Translation: Text Message Spoofing via Email	2660
<i>Sumanth Rao, Ye Shu, Stefan Savage, Aaron Schulman, Geoffrey M. Voelker, Enze Liu</i>	
From “Be Careful” to “Here's Why”: Investigating User Reasoning with Context-Specific SMS Scam Warnings	2679
<i>Elijah Bouma-Sims, Enze Liu, Alexandra Xinran Li, Lorrie Faith Cranor</i>	
Towards Automating Data Access Permissions in AI Agents	2699
<i>Yuhao Wu, Ke Yang, Franziska Roesner, Tadayoshi Kohno, Ning Zhang, Umar Iqbal</i>	
COSSETER: GitHub Actions Permission Reduction Using Demand-Driven Static Analysis	2718
<i>Greg Tystahl, Jonah Ghebremichael, Siddharth Muralee, Sourag Cherupattamoolayil, Antonio Bianchi Aravind Machiry, Alexandros Kapravelos, William Enck</i>	
EyeSpy: Inferring Eye Gaze via Side-Channel Attacks Against Foveated Rendering	2736
<i>Paul Maynard, Harris Amjad, Camila Molinares, Bo Ji, Brendan David-John</i>	

SaTor: Exploring Satellite Routing in Tor to Reduce Latency	2756
<i>Haozhi Li, Tariq Elahi</i>	
A Context Is Worth a Thousand Lies: Evading Intrusion Detectors via Intelligent Context Distortion	2775
<i>Magdy Nasr, Vansh Rastogi, Azadeh TabibanB</i>	
MadeYouReset: Exploiting HTTP/2 Server-Side Resets for Large-Scale DoS	2793
<i>Gal Bar Nahum, Anat Bremler Barr, Yaniv Harel</i>	
Guardians of the Air: In-Device Detection of 5G Control-Plane Threats	2811
<i>Tianwei Wu, Abdullah Al Ishtiaq, Tianchang Yang, Yilu Dong, Kai Tu, Zeyu Song, Ridwanul Hasan Tanvir Md Toufikuzzaman, Shagufta Mehnaz, Syed Rafiul Hussain</i>	
The Threat Landscape of IP Leasing in the RPKI Era	2831
<i>Weitong Li, Yongzhe Xu, Taejoong Chung</i>	
Batch Me If You Can: Coverage-Guided RPKI Fuzzing at Scale	2850
<i>Haya Schulmann, Niklas Vogel</i>	
The Fault in Our Drafts: Vulnerabilities in RPKI Specification and Software	2867
<i>Oliver Jacobsen, Tobias Kirsch, Haya Schulmann, Niklas Vogel, Michael Waidner</i>	
Camveil: Unveiling Security Camera Vulnerabilities Through Multi-Protocol Coordinated Fuzzing	2886
<i>Fuchen Ma, Yuqiao Yang, Yuanliang Chen, Yanyang Zhao, Ting Chen, Yu Jiang</i>	
Designing Transport-Level Encryption for Datacenter Networks	2902
<i>Tianyi Gao, Xinshu Ma, Suhas Narreddy, Eugenio Luo, Steven W. D. Chien, Michio Honda</i>	
Fizzle: A Framework for Deterministic and Reproducible Network Fuzzing	2920
<i>Nathaniel Bennett, Tyler Tucker, Carson Stillman, William Enck, Patrick Traynor, Kevin Butler</i>	
The Battle of Metasurfaces: Understanding Security in Smart Radio Environments	2937
<i>Paul Staat, Christof Paar, Swarun Kumar</i>	
SatBleed: Security of Commoditized Communication Modules in Satellites	2956
<i>Ulysse Planta, Julian Rederlechner, Martin Strohmeier, Mathias Fischer, Ali Abbasi</i>	
RIS-CLA: Reviving CSI-Based Continuous Location Authentication With Reconfigurable Intelligent Surfaces	2974
<i>Yan Zhang, Jiawei Li, Yizhou Wang, Dianqi Han, Yanchao Zhang, Aditya Shekhawat Georgios Trichopoulos</i>	
MUSICSHIELD: Protection for Musicians in the Era of Generative AI	2990
<i>Syed Irfan Ali Meerza, Jian Liu</i>	
MAYA: Addressing Inconsistencies in Generative Password Guessing Through a Unified Benchmark	3009
<i>William Corrias, Fabio De Gaspari, Dorjan Hitaj, Luigi V. Mancini</i>	
Credential Extraction Attacks Against Compromised Credential Checking Services of Password Managers	3029
<i>Yihe Duan, Ding Wang, Yutong Li</i>	

Can Foundation LLMs Accurately Estimate Password Strength and Provide Appropriate Password Feedback?	3049
<i>Madison Pickering, Garrison Hinson-Hasty, Luca Dovichi, Helena Williams, Nathaniel Kim, Aybala Esmer Blase Ur</i>	
MoPE: A Mixture of Password Experts for Improving Password Guessing	3069
<i>Mingjian Duan, Ming Xu, Shenghao Zhang, Weili Han</i>	
zkFuzz: Foundation and Framework for Effective Fuzzing of Zero-Knowledge Circuits	3089
<i>Hideaki Takahashi, Jihwan Kim, Suman Jana, Junfeng Yang</i>	
Single-Server Private Outsourcing of zk-SNARKs	3109
<i>Kasra Abbaszadeh, Hossein Hafezi, Jonathan Katz, Sarah Meiklejohn</i>	
Language-Agnostic Detection of Computation-Constraint Inconsistencies in ZKP Programs Via Value Inference	3125
<i>Arman Kolozyan, Bram Vandenbogaerde, Janwillem Swalens, Lode Hoste, Stefanos Chaliasos Coen De Roover</i>	
APEX: Accurate Parallel Expressive Homomorphic Execution for Encrypted Databases	3145
<i>Wei Chen, Qi Hu, Siu-Ming Yiu, Heming Cui</i>	
Secure Lookup Tables: Faster, Leaner, and More General	3164
<i>Chongrong Li, Pengfei Zhu, Yun Li, Zhanpeng Guo, Jingyu Li, Yuncong Hu, Zhicong Huang, Cheng Hong</i>	
Efficient Arithmetic-and-Comparison Homomorphic Encryption with Space Switching	3184
<i>Erwin Eko Wahyudi, Yan Solihin, Qian Lou</i>	
Code-Based Scalable Collaborative SNARKs	3201
<i>Christodoulos Pappas, Dimitrios Papadopoulos, Charalampos Papamanthou</i>	
Optimistic Asynchronous Dynamic-Committee Proactive Secret Sharing	3220
<i>Bin Hu, Jianwei Liu, Zhenliang Lu, Qiang Tang, Zhuolun Xiang, Zongyang Zhang</i>	
Nebula: Proving Machine Executions via Folding Schemes	3239
<i>Arasu Arun, Srinath Setty</i>	
Consistent Estimation of Numerical Distributions Under Local Differential Privacy by Wavelet Expansion	3259
<i>Puning Zhao, Zhikun Zhang, Bo Sun, Li Shen, Liang Zhang, Shaowei Wang, Zhe Liu</i>	
Breaking the Barrier for Asynchronous MPC with a Friend	3278
<i>Banashri Karmakar, Aniket Kate, Shravani Patil, Arpita Patra, Sikhar Patranabis, Protik Paul, Divya Ravi</i>	
Best of Both Worlds: Effective Foreign Bridge Identification in V8 Embedders for Security Analysis	3297
<i>Georgios Alexopoulos, Thodoris Sotiropoulos, Zhendong Su, Dimitris Mitropoulos</i>	
Decor: Delegated Computation on Randomness for Secure Evaluation of Nonlinear Functions	3317
<i>Haris Smajlović, Kyle Sheng, Timos Antonopoulos, Ruzica Piskac, Hyunghoon Cho</i>	
Bridge: High-Order Taint Vulnerabilities Detection in Linux-Based IoT Firmware	3336
<i>Jiaqian Peng, Puzhuo Liu, Yicheng Zeng, Kai Cheng, Yongji Liu, Yun Yang, Hongsong Zhu</i>	
SFA-Miner: Mining Path-Sensitive API Usage Patterns Via Symbolic Finite Automata	3356
<i>Jiasheng Jiang, Mingwei Zheng, Qingkai Shi, Xiangyu Zhang</i>	

ENCHTABLE: Unified Safety Alignment Transfer in Fine-Tuned Large Language Models	3374
<i>Jialin Wu, Kecen Li, Zhicong Huang, Xinfeng Li, Xiaofeng Wang, Cheng Hong</i>	
QuickSafe: Targeted Hardening Against Memory Corruption	3393
<i>Johannes Blaser, Floris Gorter, Klaus v. Gleissenthall, Herbert Bos</i>	
SpecAuditor: Generating Audit Specifications for LLM-Driven Bug Detection	3412
<i>Miaoqian Lin, Hao Chen</i>	
BACHunter: Detecting Broken Access Control Vulnerabilities in Intelligent Connected Vehicles	3430
<i>Yanbang Sun, Xiaohong Li, Quanzhou Wang, Hebo Leng, Guangzheng Yao, Zhihua Xie, Qiang Hu Junjie Wang</i>	
Oxidizer: Toward Concise and High-fidelity Rust Decompilation	3449
<i>Yibo Liu, Zion Leonahenahe Basque, Arvind S Raj, Chavin Udomwongsa, Chang Zhu, Jie Hu Changyu Zhao, Fangzhou Dong, Adam Doupé, Tiffany Bao, Yan Shoshitaishvili, Ruoyu Wang</i>	
NanoTag: Systems Support for Efficient Byte-Granular Overflow Detection on ARM MTE	3468
<i>Mingkai Li, Hang Ye, Joseph Devietti, Suman Jana, Tanvir Ahmed Khan</i>	
Parasol Compiler: Pushing the Boundaries of FHE Program Efficiency	3487
<i>Rick Weber, Ryan Orendorff, Ghada Almashaqbeh, Ravital Solomon</i>	
The First Large-Scale Systematic Study of Python Class Pollution Vulnerability	3506
<i>Zhengyu Liu, Jiacheng Zhong, Jianjia Yu, Muxi Lyu, Zifeng Kang, Yinzhi Cao</i>	
Cottontail: Large Language Model-Driven Concolic Execution for Highly Structured Test Input Generation	3525
<i>Haoxin Tu, Seongmin Lee, Yuxian Li, Peng Chen, Lingxiao Jiang, Marcel Böhme</i>	
Catch Me If You Can: Detector-Resistant Evasion via Semantics-Preserving Command Re-Realization	3544
<i>Muhammad Shoaib, Hare Sudhan Muthusamy, Tareq Alkhatib, Wajih Ul Hassan</i>	
INSIGHT: Automatic Generation of Explanations for Efficient Identification of Hardware Bugs and Underspecifications	3564
<i>Vincent Quentin Ulitzsch, Alessandro Bertani, Peter W. Deutsch, David Langus Rodriguez, Kelly Xu Aarti Gupta, Sharad Malik, Mengjia Yan</i>	
Interplay of Efficient Model Checking and Secure Processor Design: A Case Study on Secure Speculation	3583
<i>Tingzhen Dong, Qinhan Tan, Kunpeng Wang, Thomas Bourgeat, Yuheng Yang, Sharad Malik, Yu-Wei Fan Mengjia Yan</i>	
CiRCLE: Recovering Complex Data Structures in Binaries Beyond Fragmentation	3603
<i>Zeyu Gao, Junlin Zhou, Songtao Yang, Chao Zhang</i>	
No Password, No Problem? A Large-Scale Field Study of Passkey Adoption and Usage	3623
<i>Tobias Reitinger, Günther Pernul</i>	
Usable Anonymity in Reproductive Health Privacy	3643
<i>Qirong Song, Yanlai Wu, Rie Helene Hernandez, Yao Li, Yubo Kou, Xinning Gui</i>	
The Passkey Promise: A Comparative Usability Study of MFA Methods	3663
<i>Erwin Kupris, Thomas Schreck</i>	

2FiA: Towards WiFi Sensing-Based Authentication with Unique Biometrics	3683
<i>Li Bofan, Ye Zhankai, Yu Weikuan, Tang Yongning, Liu Xiu</i>	
SmuFuzz: Enable Deep System Management Mode Fuzzing in Fully Featured UEFI Runtime Environment	3700
<i>Jianqiang Wang, Yi Xiang, Meng Wang, Qinying Wang, Ali Abbasi, Thorsten Holz</i>	
Jazzer: Coverage-Guided Fuzzing for Semantic Vulnerabilities in the Java Ecosystem	3718
<i>Sergej Dechand, Tobias Wienand, Fabian Meumertzheim, Peter Samarin, Simon Resch, Khaled Yakdan, Thorsten Holz, Flavio Toffalini</i>	
Beyond Nodes vs. Edges: A Multi-View Fusion Framework for Provenance-Based Intrusion Detection	3738
<i>Fan Yang, Binyan Xu, Di Tang, Kehuan Zhang</i>	
Specializing Language Models for Textual Fuzzing via Reinforcement Learning	3758
<i>Jiayi Lin, Liangcai Su, Junzhe Li, Chenxiong Qian</i>	
PILOT: Command-Line Interface Fuzzing Via Path-Guided, Iterative Large Language Model Prompting	3775
<i>Momoko Shiraishi, Yinzhi Cao, Takahiro Shinagawa</i>	
Contextualizing Sink Knowledge for Java Vulnerability Discovery	3794
<i>Fabian Fleischer, Cen Zhang, Joonun Jang, Jeongin Cho, Meng Xu, Taesoo Kim</i>	
deepSURF: Detecting Memory Safety Vulnerabilities in Rust Through Fuzzing LLM-Augmented Harnesses	3814
<i>Georgios C. Androutsopoulos, Antonio Bianchi</i>	
CenRL: A Framework for Performing Intelligent Censorship Measurements	3834
<i>Hieu Le, Armin Huremagic, Kevin Wang, Roya Ensafi, Ram Sundara Raman</i>	
Prrr: Personal Random Rewards for Blockchain Reporting	3853
<i>Hongyin Chen, Yubin Ke, Xiaotie Deng, Ittay Eyal</i>	
CenAlert: Amplifying User Voices to Rally Censorship Investigation	3872
<i>Aaron Ortwein, Anna Ablove, Armin Huremagic, Luqin Chang, Vinicius Fortuna, Roya Ensafi</i>	
Promoguardian: Detecting Promotion Abuse Fraud with Multi-Relation Fused Graph Neural Networks	3891
<i>Shaofei Li, Xiao Han, Ziqi Zhang, Minyao Hua, Shuli Gao, Zhenkai Liang, Yao Guo, Xiangqun Chen, Ding Li</i>	
Hidden Secrets in the arXiv: Discovering, Analyzing, and Preventing Unintentional Information Disclosure in Source Files of Scientific Preprints	3909
<i>Jan Pennekamp, Johannes Lohmöller, David Schütte, Joscha Loos, Martin Henze</i>	
Revelio: Blurred Images Can Still Disclose Your Identity	3929
<i>Haoyu Zhai, Shuo Wang, Pirouz Naghavi, Qingying Hao, Gang Wang</i>	
Fine-Grained Kernel Auditing Using Augmented Syscall Reference Behavior Analysis and Virtualized Selective Tracing	3948
<i>Chuqi Zhang, Spencer Faith, Feras Al-Qassas, Theodorus Februnto, Zhenkai Liang, Adil Ahmad</i>	
Fractal: An Operating System Designed for Microarchitecture Reverse Engineering	3967
<i>Joseph Ravichandran, Mengjia Yan</i>	

Stop Starving or Stuffing Me: Boosting Firmware Fuzzing Efficiency with On-Demand Input Delivery	3983
<i>Shandian Shen, Wei Zhou, Keming Zhao, Peng Liu, Chung Hwan Kim, Le Guan</i>	
PufferDoS: Efficient and Effective Attack String Generation for Regular Expression Denial of Service Vulnerabilities	4002
<i>Shangzhi Xu, Ziqi Ding, Xiao Cheng, Yuekang Li, Nan Sun, Benjamin Turnbull, Shuangxiang Kan, Siqi Ma</i>	
HEAP LOCALIZATION: Cache Side-Channel Based Linux Kernel Heap Exploit Techniques	4021
<i>Yoochan Lee, Sihyun Roh, Hyuk Kwon, Byoungyoung Lee, Thorsten Holz</i>	
NetPanic: the Attack Surface You Can't Syscall	4038
<i>Tianshuo Han, Zong Cao, Zhen Dong, Xiapu Luo, Zhenyu Song, Jian Liu</i>	
APIECHO: Training-Less Anomaly Detection via Intra-API Behavioral Comparison for Web Applications	4054
<i>Yihao Peng, Yiming Wu, Du Wu, Shouling Ji, Hai Wan, Xibin Zhao</i>	
PortGPT: Towards Automated Backporting Using Large Language Models	4071
<i>Zhaoyang Li, Zheng Yu, Jingyi Song, Meng Xu, Yuxuan Luo, Dongliang Mu</i>	
Behind the Curtain: How Shared Hosting Providers Respond to Vulnerability Notifications	4091
<i>Giada Stivala, Rafael Mrowczynski, Maria Hellenthal, Giancarlo Pellegrino</i>	
TRIGFUZZ: Triggering Conditions Guided Directed Fuzzing	4110
<i>Yiyang Chen, Nuoqi Gui, Long Wang, Longfei Chen, Xuanqing Shi, Xi Cao, Chao Zhang</i>	
Death Is Not the End: a Longitudinal Study on the Impact of Automatic Updates on Container Vulnerability Lifespans	4129
<i>Simge Tekin, Octavian Suciu, Sungsu Kwag, Yonghwi Kwon, Tudor Dumitras</i>	
Web Application Vulnerability Repair Via Context-Aware Fault Localization and Directed Differential Fuzzing	4147
<i>Chenlin Wang, Wei Meng</i>	
Papers, Please: A First Look at Age Verification on the Web	4167
<i>Shreyas Minocha, Isaac Sheridan, Harry Oppenheimer, Paul Pearce, Michael A. Specter</i>	
Who Taught the Lie? Responsibility Attribution for Poisoned Knowledge in Retrieval-Augmented Generation	4185
<i>Baolei Zhang, Haoran Xin, Yuxi Chen, Zhuqing Liu, Biao Yi, Tong Li, Lihai Nie, Zheli Liu, Minghong Fang</i>	
When AI Meets the Web: Prompt Injection Risks in Third-Party AI Chatbot Plugins	4205
<i>Yigitcan Kaya, Anton Landerer, Stijn Pletinckx, Michelle Zimmermann, Christopher Kruegel Giovanni Vigna</i>	
PromptLocate: Localizing Prompt Injection Attacks	4225
<i>Yuqi Jia, Yupei Liu, Zedian Shao, Jinyuan Jia, Neil Zhenqiang Gong</i>	
LLMs in the SOC: An Empirical Study of Human-AI Collaboration in Security Operations Centres	4244
<i>Ronal Singh, Shahroz Tariq, Fatemeh Jalalvand, Mohan Baruwat Chhetri, Surya Nepal, Cecile Paris Martin Lochner</i>	
Incalmo: an Autonomous Llm-Assisted System for Red Teaming Multi-Host Networks	4264
<i>Brian Singer, Keane Lucas, Lakshmi Adiga, Meghna Jain, Lujo Bauer, Vyas Sekar</i>	

PromptCOS: Towards Content-Only System Prompt Copyright Auditing for LLMs	4283
<i>Yuchen Yang, Yiming Li, Hongwei Yao, Enhao Huang, Shuo Shao, Yuyi Wang, Zhibo Wang, Dacheng Tao Zhan Qin</i>	
AttnTrace: Contextual Attribution of Prompt Injection and Knowledge Corruption	4302
<i>Yanting Wang, Runpeng Geng, Ying Chen, Jinyuan Jia</i>	
Hollow-LLM Attack: Computationally Trivial Weights in Zero-Knowledge Verification of LLM Inference	4321
<i>Chen Gong, Beijie Liu, Mengyuan Li</i>	
Leafblower: a Leakage Attack Against Tee-Based Encrypted Databases	4338
<i>Zachary Espiritu, Seny Kamara, Tarik Moataz, Valentin Ogier</i>	
AEX-NStep: Probabilistic Interrupt Counting Attacks on Intel SGX	4358
<i>Nicolas Dutly, Friederike Groschupp, Ivan Puddu, Kari Kostinen, Srdjan Čapkun</i>	
Battering RAM: Low-Cost Interposer Attacks on Confidential Computing via Dynamic Memory Aliasing	4373
<i>Jesse De Meulemeester, David Oswald, Ingrid Verbauwhede, Jo Van Bulck</i>	
TEE.Fail: Breaking Trusted Execution Environments via DDR5 Memory Bus Interposition	4390
<i>Jalen Chuang, Alex Seto, Nicolas Berrios, Stephan Van Schaik, Christina Garman, Daniel Genkin</i>	
PrintSpy: Pixel-Level Eavesdropping on Commodity Laser Printers via Electromagnetic Side Channels	4409
<i>Wenhao Li, Jiarong Yang, Mingda Han, Xiuzhen Cheng, Pengfei Hu, Cong Wang</i>	
TÄMU: Emulating Trusted Applications at the (GlobalPlatform)-API Layer	4427
<i>Philipp Mao, Li Shi, Marcel Busch, Mathias Payer</i>	
Goldilocks and the Three P-States: Mitigating Hertzbleed with Formal Leakage Guarantees	4443
<i>Inwhan Chun, Christine Guo, Riccardo Paccagnella</i>	
TDXRay: Microarchitectural Side-Channel Analysis of Intel TDX for Real-World Workloads	4461
<i>Tristan Hornetz, Hosein Yavarzadeh, Albert Cheu, Adria Gascon, Lukas Gerlach, Daniel Moghimi Phillipp Schoppmann, Michael Schwarz, Ruiyi Zhang</i>	
Investigating the Impact of Dark Patterns on LLM-Based Web Agents	4479
<i>Devin Ersoy, Brandon Lee, Ananth Shree Kumar, Arjun Arunasalam, Muhammad Ibrahim, Antonio Bianchi Z. Berkay Celik</i>	
Demystifying the (In)Security of OAuth-Based Account Linking in Connector Ecosystems	4499
<i>Kaixuan Luo, Xianbo Wang, Pui Ho Adonis Fung, Wing Cheong Lau</i>	
When Designers Meet GenAI: Understanding the Role of Prompt-to-Design Generators in Privacy Dark Patterns	4519
<i>Jingzhou Ye, Zhaojie Hu, Yao Li, Xueqiang Wang</i>	
SoK: After Decades of Web Tracker Detection, What's Next?	4539
<i>Wolf Rieder, Philip Raschke, Thomas Cory, Christian Rene Sechting, Aditya Kumar, Axel Küpper</i>	
Your Eyes Won't Lie: Snooping Online Voting Privacy from User Webcam	4565
<i>Zeyu Deng, Jingwei Zhang, Chen Wang</i>	

Weaponizing Reflectivity for Pointcloud Deception with Forged Invisible Geometries	4582
<i>Hengwei Chen, Menglan Hu, Tianyue Zheng</i>	
PLaTypus: Restricting Cross-Module Transitions to Mitigate Code-Reuse Attacks	4600
<i>Apostolos Chatzianagnostou, Marcos Bajo, Christian Rossow</i>	
Crashing Through Defenses: Exploiting Segfaults and Chaining Around Intel CET	4620
<i>Marcos Bajo, Ritvik Goyal, Apostolos Chatzianagnostou, Christian Rossow</i>	
BreakFAST: Confused Deputy Attack on Infinity Fabric to Break AMD SEV-SNP	4640
<i>Philipp Giersfeld, Benedict Schlüter, Shweta Shinde</i>	
One Char to Rule Them All: Systematically Exploring and Exploiting DNS Silent Vulnerabilities in Domain Name Resolution	4657
<i>Fasheng Miao, Xiang Li, Changqing An, Wenbin Xu, Jilong Wang</i>	
Resolve the Unresolved: Systematic Work Profiling for DNS Resolvers	4677
<i>Liwen Xu, Huayi Duan, Zechao Cai, Adrian Perrig</i>	
Poisoned by the Host: Large-Scale Measurement of Host Name Poisoning in Web Applications	4694
<i>Rui Yang, Haoyu Wang, Zhicheng Sun, Zhengyu Liu, Yinzhi Cao</i>	
Knocking on the Front Door: An LLM-Guided Systematic Analysis of DNS Query Processing Vulnerabilities	4713
<i>Yuqi Qiu, Xiang Li, Zheli Liu</i>	
APT to Disagree: A Comparative Analysis of Attribution in Commercial TI	4733
<i>Aksel Ethembabaoglu, Rolf Van Wegberg, Yury Zhauniarovich, Michel Van Eeten</i>	