2026 56th Annual IEEE International Conference on Dependable Systems and Networks - Supplemental Volume (DSN-S 2026)

Charlotte, North Carolina, USA
22-25 June 2026

IEEE Catalog Number: CFP26V03-POD
ISBN: 979-8-3195-3247-3

IEEE
Advancing Technology for Humanity

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:	CFP26V03-POD
ISBN (Print-On-Demand):	979-8-3195-3247-3
ISBN (Online):	979-8-3195-3246-6
ISSN:	2833-2903

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

CURRAN ASSOCIATES INC. proceedings .com

TABLE OF CONTENTS

From Security Standards to Adversary Emulation: An AI-Assisted Framework for ICS Security Assessment 1
Nima Faraji, Giovanni Mazzeo, Roberto Nardone, Francesco Brancati, Luigi Romano

Event Burst Trigger: An Availability Backdoor Attack on Event-Based SNN Object Detection 8
Jaesun Baek, Chanwook Lee, Eun-Kyu Lee

Knowledge Transfer from LLMs to Provenance Analysis: Semantic-Augmented APT Detection 15
Fei Zuo, Junghwan Rhee, Yung Ryn Choe, Chung Hwan Kim

Redefining Page Offlining: Real-World Memory Failure Modes with Preemptive Strategies 22
Ankit Saxena, Jongmin Lee, Youngwook Ko, Stuti Patel, Grant Gambetta, DanKyu Kang, Yoon Kim Ilgyu Jung, Paul Cho

Harnessing Large Language Models for Seed Generation in Greybox Fuzzing 29
Yunhang Zhang, Wenxuan Shi, Xinyu Xing, Jun Xu

SNAS: A Multi-Layer Defense-in-Depth Architecture for Secure Egress in Sandboxed Workloads 36
Niranjan Kumar Sharma, S Muralidhar, Samy Boshra-Riad, Mike Halcrow, Yuxiong He Nitya Kumar Sharma, Shawn Xia, Haowei Yu, Elliott Brossard, Derek Denny-Brown, Choden Konigsmark Bhanu Prakash, Brandon Baker, Andong Zhan

VeriExist: Blockchain-Assisted Existence-Verifiable Protocol in User-Centric Data Sharing 43
Rui Han, Weizhong Qiang, Weijie Chen, Yulin Tang

AEGIS: Preventing Cross-Domain Resource Abuse in MCP 50
Shriti Priya, Teryl Taylor, Frederico Araujo

A Polyglot OpenTelemetry Framework for Cloud-Native Reliability and Performance Engineering in Connected-Vehicle Platforms 57
Gaurav Saxena, Niladri Bihari Mohanty

A Secure, Manifest-Based Framework for Delegated Privilege Promotion 64
Rajarshi Chowdhury, Akshay Shah

The Anatomy of Silent Data Corruption: GPU Error Pattern Study and Modeling Guidance 71
Chung-Hsuan Tung, Yanxiang Huang, Nirmal Saxena, Philip Shirvani, Saurabh Hukerikar, Twinkle Jain Abhishek Tyagi, Sanjay Gongalore

CRISP: End-to-End Fingerprint Recognition Leveraging Hardware-Anchored Liveness Detection 78
Subin Ki, Kyuseung Han, Hyuk Kim, Taewook Kang, Kwang-Il Oh, Hyeonguk Jang, Sukho Lee, Jae-Jin Lee Jaehyoung Lee, Woojoo Lee, Jinho Lee

LLM4BP: A Collaborative LLM Framework for Insider-Resilient BPMN Workflows in Industrial Control Systems 85
Antonio Iannaccone, Chao Feng, Roberto Nardone, Burkhard Stiller

AIReSim: A Discrete Event Simulator for Large-Scale AI Cluster Reliability Modeling 92
Karthik Pattabiraman, Mihir Patel, Fred Lin

Decentralized Root Cause Analysis for Cloud-Native Microservices - Experience with Distributed PageRank in Production 99
Shabbir Hussain, Hari Charan Panjwani

Harpocrates-in-Action: Using Simulation-Based Functional Test Programs to Catch Marginal Defects in CPUs 106
Nikos Karystinos, Odysseas Chatzopoulos, Dimitris Gizopoulos, Ashish Kumar, Sankar Gurumurthy Sudhanva Gurumurthi, Vilas Sridharan

EgoLoc: Secure Ego-Agent Selection for Robust Collaborative Perception Against Privilege Hijacking 113
Kyeomhyeok Woo, Eun-Kyu Lee, Jiwoo Lee

A Configurable Logic Architecture Based on Micro-Electromechanical Relays for Nuclear Power Safety Systems 120
Coby Cockrell, Aidan Collins, Ben Smith, Umit Ozgur, Vitaliy Avrutin, Matt Gibson, Abhi Rajagopala Carl Elks

Cross-Referencing Whitepaper Claims and Smart Contract Code for Automated ICO Auditing 127
Chenglong Fu, Chi Zhang, Mengtian Li

A Robustness Metric and Evaluation Approach for Machine Learning Components in Aviation 134
Casidhe Hutchison, Tomasz Swierzewski, Joshua Neighbor, Alex Johnson, Todd Grossaint, Xusheng Luo Changliu Liu

An IPsec Tunnel Recovery Protocol to Overcome QKD Key Desynchronization 141
Hong-Sung Kim, Hyong-Shik Kim

PeerLumen: A Bitcoin P2P Honeypot Telescope 148
Wenjun Fan, Fei Du, Jungho Kim, Sang-Yoon Chang

Enterprise-Scale RAS in IBM Z Memory Subsystem 155
Patrick Meaney, David Cadigan, Gary Van Huben, Stephen Powell, Kelvin Wong, Colleen Weller Gary Tressler, Hailey Nicholson, Marc Gollub, Alper Buyuktosunoglu

HBM Errors at Scale: An Industrial Study 163
Shreya Varshini, Qintian Huang, Xun Jiao, Mihir Patel, Abhinav Pandey, Fred Lin

Structural Compliance, Governed Adaptation: Dependable Evolution in the eAyroll Governmental Blockchain Platform 169
Attila Klenik, Imre Kocsis, Sándor Pap, Péter Tóth, László Gönczy, Ferenc Vágujhelyi, András Pataricza

[Tutorial] Confidential Computing: Fundamentals, Platforms, and Research 176
Zhongshu Gu, Julian Stephen, Enriquillo Valdez

Fault Injection for ML Applications: from DNNs to LLMs 178
Karthik Pattabiraman, Bo Fang, Lishan Yang, Abraham Chan, Yu Sun

Preventing Human Errors for Dependable Computer System Design 180
Fuqun Huang

[Tutorial] Securing Your Airspace: The Interplay Between Cybersecurity, Sensing, and Safety 182
Alireza Famili, Amir Fakhim Babaei, Tolga O. Atalay, Angelos Stavrou

RESIST: $\underline{\text{Res}}$ilience Against $\underline{\mathrm{S}}$ilent $\underline{\mathrm{S}}$ilicon $\underline{\mathrm{T}}$hreats — A Cross-Layer Perspective (Tutorial) 184
Peter W. Deutsch, Vincent Quentin Ulitzsch, Mengjia Yan MIT, Sudhanva Gurumurthi, Vilas Sridharan Harish Dixit, Sriram Sankar Meta Platforms

Towards Security by Performance: Rethinking Cyber Risk Assessment for Distributed ML Systems 186
Alessandro Palma, Gianluca Capozzi, Silvia Bonomi, Giuseppe Antonio Di Luna

Network-Aware Differential Privacy 191
Zhou Li, Yu Zheng, Tianhao Wang, Sang-Woo Jun

Fast or Slow? Human-Inspired Self-Evolving Framework for Resilient AI Systems 197
Haoran Qiu, Phuong Cao, Shengkun Cui, Archit Patke, Ravishankar K. Iyer

Trust, but Verify: ByzTwin-Range, a Digital Twin Cyber-Range for Byzantine Faults 203
Tadeu Freitas, João Soares, Rolando Martins

In-DRAM Signature Generation Using Simultaneous Multiple-Row Activation: An Experimental Study of Off-The-Shelf DRAM Chips 208
Umut Başer, İsmail Emir Yüksel, F. Nisa Bostancı, Konstantinos Sgouras, Ataberk Olgun Emre Hakan Demirli, Zhiheng Yue, Harsh Songara, Oğuz Ergin, Onur Mutlu

Taming the Airwaves: Resilient Spectrum Sharing with BFT 215
Xiaoxuan Qin, Prashant Krishnamurthy, Amy Babay

RO: Reliable Orchestration of LLM Programs via Deterministic-Generative Separation 220
Saurabh Jha

An Experimental Paradigm for Studying Cognitive Mechanisms Behind Software Faults 225
Fuqun Huang, Owen Wright

Dynamic Wakeup with Collision Costs 230
Umesh Biswas, Maxwell Young

Balancing Mutation Entropy and Session Continuity in SDN-Based Moving Target Defense 233
Souhail Chakkour, Umesh Biswas, Charan Gudla

Contention Resolution with Evictions 236
Umesh Biswas, Vini Chaudhary, Shafqat Hasan, Maxwell Young

Securing Blockchain Applications by Transaction Analysis 239
Rui Xi

Dependable Real-Time Guarantees in Converged 5G and Time-Sensitive Networks 242
Simon Egger, Frank Dürr, Christian Becker

Data Quality and Fairness in Federated Learning 245
Sudlow James, Bouchenak Sara

Agentic Reinforcement Learning for Adaptive Moving Target Defense in Cyber-Physical Systems 248
Mohammad Zakaria Haider, Mohammad Ashiqur Rahman

Automating Research and Development of Dependable Federated Learning Systems 251
Dmitrii Korobeinikov, Leon Reznik

FGF: A Fact-Graph Feedback Loop Framework for Verifiable LLM-Based Decision Making 254
Subi Kim, Jimin Ryu, Jieun Kang, YongIk Yoon

Adaptive Symbolic Execution with LLMs for Evasive Malware Analysis 256
Ajay Kumara Makanahalli Annaiah, Shahid Ali, William Setters

Co-Evolving SQL Injection Attacks and Modsecurity Rules via Agentic AI 258
Vahid Babaey, Arun Ravindran

Why are Anomalies Rare in RSS? Diagnosing Strict Serializability Violations in Spanner-RSS 260
Sejong Kim, Yon Dohn Chung

Towards Automated Detection of Latent Software Robustness Failures Using Large Language Models 262
José Flauzino, Marco Vieira, Elias P. Duarte

Optimal Fast-ReRoute: A Tool for Generating and Visualizing Arborescence-Based Routes 264
Leon Okida, Elias P. Duarte, André Vignatti

Making Binary Decisions in Decentralized Autonomous Organizations 266
Nuno Braz, Miguel Correia

Memory-Bandwidth Side Channels for Input Recovery in Edge-Deployed LLM Inference 268
Davis Ranney, Yashaswini I Makaram, Yunsi Fei

Does the Knight and Leveson Hypothesis Hold for LLM-Based Code Generation? 271
Rodrigo Pato Nogueira, Karthik Pattabiraman, Marco Vieira, João R. Campos

The Overlooked Role of False Negatives in Reentrancy Detection 273
Alvise Spano, Dalila Ressi, Sabina Rossi, Michele Bugliesi

LUMENating: Insights from an LLM-Guided System for Activating Hidden Behaviors in Malware 275
Kevin Valakuzhy, Abhinav Vemulapalli, Miuyin Yong Wong, Douglas M. Blough, Mustaque Ahamad Fabian Monrose

Mind the Driver: A Fault Injection Methodology for AI Systems 277
Marcello Cinque, Luigi De Simone, Nike Di Giacomo

On the Feasibility of Side-Channel Information for Malicious Traffic Detection in IoT Networks 279
Shadman Sakib Arnob, Ahmad Patooghy, Islam Obaidat

Edge-Based Privacy Auditing for Trustworthy Humanoid Robot Interaction 281
Bontu Habtamu Geremew, Sharun Akter Khushbu, Shaohu Zhang

Evaluating Security Best Practices in the GitHub Actions Documentation 283
Lukas Boschanski, Marco Vieira

Demystifying Differentially Private Quantization in Federated Learning 285
Shan Huang, Zhijing Ye, Jiamin Wang, Zhaorui Zhang, Jie Fu, Wendy Hui Wang, Sheng Di, Xiaodong Yu

Investigating the Relationship Between Autonomous Driving System Disengagements and Crashes 287
Dev Patel, Arpan Swaroop, Saurabh Bagchi