ETH zürich



Zurich Information Security and Privacy Center (ZISC)

Annual Review 2026

ETH Zurich, Zurich Information Security and Privacy Center (ZISC)

Introduction



Prof. Srdjan Capkun
Chair of the center



Dr. Kari Kostiainen
Director of the center

The information security landscape is going through a dramatic disruption. Within the last few years, **AI models** have gained increasing capabilities, and in some domains they are surpassing human abilities. Take the example of coding—leading AI models can analyze large codebases and write software very efficiently. This provides a huge productivity boost for software development teams. However, adversaries can use the same models to discover new software vulnerabilities and deploy sophisticated exploits at significantly lower cost and effort than before.

AI agents, which are essentially AI models connected to software tools and given Internet access, have great potential to automate tasks and thus make our lives easier. However, such agents can also cause unintended damage to their owner or to third parties while trying to carry out their assigned task step by step. Beyond software implementations, AI models have strong capabilities for analyzing cryptographic protocols that enable secure online communication. Furthermore, AI models can tackle challenging mathematical problems comparable to those that underpin modern cryptography.

Such recent developments mean that individuals, companies, organizations, and entire societies need to improve their security posture and cyber resilience. In part, this means carefully planned and thoroughly executed traditional cybersecurity work. But it also demands research to better understand these new and emerging risks, and it calls for research into novel security mechanisms and defenses.

While news about AI developments and associated threats is prominently reported in the media, it is certainly not the only source of information security disruption. The prospect of practical **quantum computing** means organizations need to take migration to post-quantum-secure systems seriously. At the same time, further research into post-quantum cryptographic primitives is needed. Recent developments in geopolitics call for greater **technological sovereignty**. Individuals, organizations, and countries need more control over where their data is stored and how the computation is protected. We also need better control over communication and more sovereign networking solutions.

Besides such emerging threats, organizations continue to struggle with **traditional information security problems:** ransomware remains a major headache for companies, phishing campaigns are prevalent, large data leaks continue to damage company reputations, and the increasing complexity of IT systems makes it difficult to apply best security practices. Although these IT security problems are well understood, the existing tools and technologies are insufficient to eliminate them.



The ZISC Center

The Zurich Information Security and Privacy Center (ZISC) was established in 2003 to bring academia and industry together to address tomorrow's information security challenges. Today, more than 20 years later, this mission remains more relevant than ever.

The ZISC center comprises 8 faculty members whose research groups, combined, include more than 100 researchers (mainly Ph.D. students and postdoctoral researchers) working in various aspects of information security, privacy, and cryptography. The center approaches its mission in the following ways:

1. We conduct **PhD-level research projects** in which we study emerging technologies to understand their novel security and privacy risks and threats, develop new security tools and solutions with strong guarantees, and tackle the fundamental open questions of information security and privacy. More information about selected research projects can be found in the Research Highlights section of this report.

2. **We educate** the next generation of academic researchers, information security experts, and IT professionals. In addition, we support educational programs in Swiss primary schools and gymnasiums. For more information on these activities, please refer to the Education section of this report.

3. We create an environment that enables promising young researchers to launch **startups** and commercialize their research. The companies founded by ZISC researchers are listed in the Startup Companies section of this report.

The center's main research areas include the security and privacy of AI technologies, confidential and sovereign computing, the foundations of cryptography, future Internet, secure positioning and localization, security protocol verification, and applied cryptography. More information can be found in the Main Research Areas section of this report. Research conducted at ZISC is supported by its Industry Partners.

Partnership Model

The typical way to engage with the ZISC center and its researchers is a long-term partnership. The ZISC partnership includes the following benefits:

1. **Customized research projects**. The primary component of the ZISC partnership is PhD-level research projects tailored to our partners' needs and interests. Every year, we conduct several such projects in collaboration with our industry partners. For more information on current research collaborations, please refer to the Research Projects section of this report.

2. **Expert advice**. We connect our partners to leading research experts (typically professors and postdocs) for technical discussions and practical advice.

3. **Networking**. The ZISC center organizes various events, including a bi-weekly lunch seminar with technical talks.

4. **Continuing education**. We also provide our partners access to cybersecurity continuing education programs. Two programs are available: Certificate of Advanced Studies (https://inf.ethz.ch/continuing-education/cas-cybersecurity.html) and Diploma of Advanced Studies (https://inf.ethz.ch/continuing-education/das-cybersecurity.html).

5. **OpenLab**. Our collaborative working environment, known as the ZISC OpenLab, enables us to host extended visits to ETH Zurich and facilitate various forms of collaboration.

Partners

The research activities of the ZISC center are supported by these partners



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

DIETER SCHWARZ STIFTUNG



Kanton Zug

Swissquote

THE SWISS LEADER IN ONLINE BANKING



Zürcher Kantonalbank

SWISS POST

We work in collaboration with the following organizations:



ICRC

ETH Foundation

ZISC Faculty Members

The ZISC center includes the following ETH faculty members:



Prof. Dr. David Basin leads the Information Security Group that performs research on methods and tools for the analysis and construction of safe and secure systems.



Prof. Dr. Srdjan Capkun leads the System Security Group, studying the design and the analysis of security protocols for wired and wireless networks and systems.



Prof. Dr. Dennis Hofheinz leads the Foundations of Cryptography group that designs and analyzes cryptographic building blocks and their use.



Prof. Dr. Kenny Paterson leads the Applied Cryptography Group whose research focus is on applied cryptography and communication security.



Prof. Dr. Adrian Perrig leads the Network Security Group whose research revolves around building secure and robust network systems – with a particular focus on the design of next-generation Internet architectures.



Prof. Dr. Shweta Shinde leads research in trusted computing and its intersection with system security, program analysis, and formal verification.



Prof. Dr. Florian Tramèr leads the Secure and Private AI Lab whose research currently focuses on understanding and improving the worst-case behavior of machine learning systems.



Prof. Dr. Kaveh Razavi leads the Computer Security Group whose main research goal is the construction of reliable and secure computing systems.

Affiliated Faculty Members

The ZISC center works in close collaboration with the following ETH faculty members:



Prof. Stelian Coros leads the Computational Robotics Lab whose research is about robots who understand the physical world and function as skilled co-workers and trusted social companions.



Prof. Hans Gersbach is a professor of Macroeconomics: Innovation and Policy at D-MTEC. His joint research with ZISC focuses on secure governance schemes through assessment voting and vote delegation.



Prof. Christian Holz leads the Sensing, Interaction and Perception Lab whose research covers topics ranging from technical computer-human interaction to wearable sensing and virtual reality.



Prof. Peter Müller leads the Chair of Programming Methodology where the main research objective is to enable programmers to develop correct software.



Prof. Laurent Vanbever leads the Networked Systems Group whose goal is to make the current and future networks, especially the Internet, easier to design, understand and operate.



Prof. Dennis Komm leads the Algorithms and Didactics group whose research focusses on the sustainable impartation of core concepts of computer science.

Administrative Team



Saskia Wolf
Administrative Assistant for Prof. Capkun, Prof. Basin, Prof. Shinde.



Vivien Klomp
Administrative Assistant for Prof. Capkun, Prof. Basin, Prof. Tramèr.

Research Highlights 2026

Large language models can deanonymize internet users at scale



Prof. Florian Tramèr

Pseudonymity is a cornerstone of how people interact online.
Yet, it has been known for decades that "anonymized" data rarely is: famously, Netflix ratings were enough to recover someone's public IMDb profile, and a few location points singled out most people in a mobile phone dataset.
But most pseudonymous online accounts have remained safe over the years. A core reason is economic rather than technical: linking a pile of unstructured forum posts to a real person would require either structured data amenable to algorithmic matching, or hours of manual work. Until now this cost was typically only worthwhile for high-value targets.

In our work this year, we show that large language models (LLMs) largely remove this economic barrier. An LLM can read a user's posts, summarize what they reveal about the user (demographics, profession, interests, writing style, etc), search for candidate identities over a large anonymity set, and reason about whether two profiles belong to the same person. Each of these steps used to require a human; now they cost a few cents to a few dollars per user.

We first asked the simplest question: if you give a frontier LLM agent an anonymized text, can it recover the author using web search?
To evaluate this, we used a dataset of 125 interview transcripts that Anthropic had published (anonymized) about how scientists use AI in their research.
Our agent re-identified at least nine of the scientists, from nothing but their description of their own work.

Such attacks demonstrate the potential of LLMs for deanonymization, but they are hard to study scientifically: ground-truth datasets are small, and it is unclear what the LLM contributes versus the search engine. We thus designed a modular three-stage pipeline, in the spirit of the classical Netflix Prize attack, that augments each step with an LLM: (1) an LLM first *extracts* identity-relevant features from raw text; (2) these features are *embedded* so that we can efficiently search over thousands or millions of candidates; (3) an LLM then *reasons* over the top candidates to pick (or reject) a match.
We evaluated this pipeline on three settings with verifiable ground truth: linking Hacker News accounts to LinkedIn profiles, linking users across different movie-discussion subreddits, and linking two halves of a single Reddit user's history split across time.
In every setting the LLM-based attack vastly outperforms classical baselines. When matching Hacker News profiles to LinkedIn, the classical baseline recovers essentially no users at 99% precision, while our attack has close to 50% recall. The LLM reasoning step is what makes the biggest difference, by automating the process of filtering out false-positives based on contextual cues.

Of course, deanonymization becomes harder as the pool of candidates grows. But we found that recall degrades only slowly with the size of the candidate pool, and extrapolating our results suggests that the attack would still find a sizeable fraction of users among a million candidates. More worryingly, even when only one in ten thousand queries has a true match in the candidate pool, the attack still confidently identifies a non-trivial fraction of those rare matchable users. The attack thus keeps working precisely in a hard, real-world regime where most people are safe but a few are not.

Our pipeline uses only public embedding models, standard LLM APIs, and off-the-shelf agent scaffolding, and costs a few dollars per target. So the implicit assumption that deanonymization is possible in theory but too expensive in practice, clearly no longer holds. Users who post under a persistent username should thus assume that their accounts can be linked to each other and to their real identity, unless they diligently avoid revealing any facts about themselves over time.

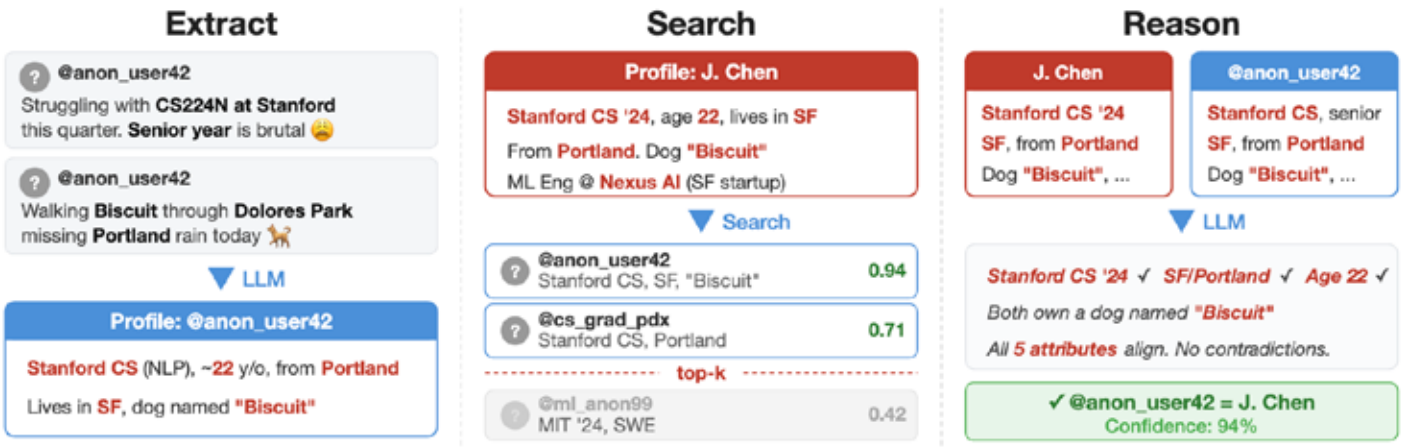
This is of course challenging, and building explicit defenses is also hard. Platforms could rate-limit access to user data and detect bulk scraping. LLM providers could also try to detect deanonymization attempts, but our attack decomposes into simple tasks (summarization, search, and ranking) that each look benign.

Further information

"Large-scale online deanonymization with LLMs"
Simon Lermen, Daniel Paleka, Joshua Swanson, Michael Aerni, Nicholas Carlini and Florian Tramèr
USENIX Security Symposium, 2026
https://arxiv.org/abs/2602.16800

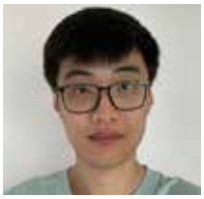
Researchers

Prof. Florian Tramèr,
Daniel Paleka,
Joshua Swanson,
Michael Aerni
(Secure and Private AI Lab, ETH Zurich)



Research Highlights 2026

Criminology: Refined Techniques for Compression Side-Channel Attacks



Yuanming Song

Compression Side Channel

It has been known for more than two decades that performing compression before encryption may leak information about the plaintext from the ciphertext length, giving rise to a compression side channel. Namely, an attacker can typically deduce the approximate length of the compressed plaintext from the ciphertext length, which intuitively reflects the redundancy of the plaintext. For example, compared to random data, repetitive data of the same length generally result in much shorter ciphertexts.

While seemingly innocuous, the compression side channel enables an attacker with the ability to adaptively embed queries in the plaintext to recover secrets from the plaintext. For example, to recover a secret byte following a known prefix "secret=" in the plaintext, an attacker can embed queries with a guess for the secret byte (e.g. "secret=1") to the plaintext. If the guess is correct, then the plaintext contains a one-byte longer repeated substring and likely results in a slightly shorter ciphertext. This technique allows the attacker to recover a secret string byte-by-byte. Prominent examples of compression side-channel attacks include CRIME and BREACH, which can recover secrets such as cookies by exploiting compression in TLS and HTTP respectively.

An Underestimated Threat

New compression-based attacks are still being regularly discovered, including against SSH, databases, encrypted messaging applications, VPNs, cloud backup mechanisms and password managers. However, even after vulnerability disclosures, affected vendors often do not implement any mitigation at all or only partially mitigate by introducing some level of noise, an approach theoretically justified in academic publications.

One reason to not mitigate is that compression side channels, like other side channels, are perceived to be "noisy" and so hard to exploit in practice; hence, the performance benefit of compression seems to outweigh its risks. Indeed, disabling compression would usually incur drastic overhead, while attacks typically require hundreds or thousands of queries to extract a short secret, and many more if noise is present, making them costly and prone to detection.

New Techniques

Our work demonstrates that compression side channels are much more dangerous than previously thought. We focus on DEFLATE, since it is the most commonly deployed compression algorithm and is used in popular compressors such as zlib and Gzip. We perform an analysis that goes deeper into DEFLATE's internals than prior work. Through this, we discover entirely novel techniques for amplifying compressed length differences: whereas previous work created length differences of a few bits or bytes, we can arrange differences of hundreds of bytes. In fact, using our most powerful techniques based on exploiting the hash tables used internally in DEFLATE, we can make compressed length differences grow linearly. These results enable us to obtain query-efficient, plaintext-recovering compression side-channel attacks that overcome existing noise-based countermeasures.

We couple these new amplification techniques with a novel abstraction for compression side-channel attacks, in which we build a deterministic state machine (with a single bit state) from the DEFLATE algorithm. We can manipulate the state of this machine through carefully crafted input strings that we call gadgets; these gadgets can be chained together to build what we call CRIME automata. They allow complex logical operations to be performed on the state, with controlled impact on the compressed output length. This enables us to modularly design efficient queries with advanced functionalities not anticipated in previous work. For example, our CRIME cascade automaton enables the attacker to identify the first match between a target string and multiple test strings, while amplifying compressed length differences to overcome noise-based countermeasures. Our designs are supported by experiments and a publicly

available codebase demonstrating the power, flexibility, and practical impact of our CRIME automata approach.

Implications

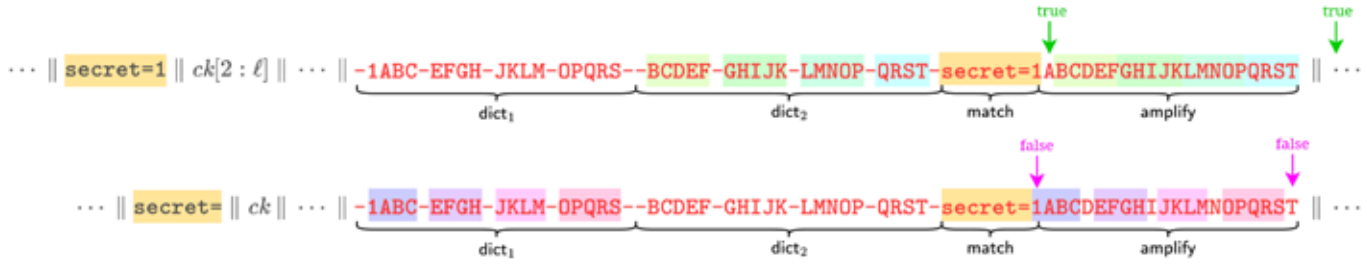
Our techniques could potentially be used to improve most existing compression side-channel attacks against DEFLATE, and they are likely to be especially effective against systems that are noisy or not highly interactive. This means that performing compression before encryption may pose greater risks than expected, and we urge vendors to mitigate the compression side channels in their systems. Apart from disabling compression, we are unaware of any other general-purpose mitigation that can prevent compression side-channel attacks, but existing mitigations such as compartmentalization and masking can help in some use cases.

Further information

Yuanming Song, Lenka Mareková, and Kenneth G. Paterson, "Criminology: Refined Techniques for Compression Side-Channel Attacks". In ACM CCS 2026. https://eprint.iacr.org/2026/1972

Researchers

Yuanming Song,
Dr. Lenka Mareková,
Prof. Kenny Paterson
Applied Cryptography Group, ETH



Research Highlights 2026

Authenticated Browsing: Defending Against Application-Layer Floods Without Breaking Confidentiality



Felix Stöger

The Confidentiality Dilemma

Banks, hospitals, and government services are attractive targets for denial-of-service (DoS) attacks, because the availability of their online services matters to their customers, patients, and citizens. The most damaging attacks today are application-layer floods. Large botnets—an army of compromised machines—send huge volumes of well-formed HTTPS requests over properly established TLS sessions. At the network level, these requests look exactly like legitimate traffic, so filters that work on packet headers cannot tell them apart.

The commercial services, such as CDNs and web application firewalls, mitigate such floods by terminating TLS and inspecting the decrypted requests. For organizations in regulated sectors, however, this is often not an option. Banking secrecy, medical confidentiality rules, and the GDPR restrict or outright prohibit giving a third party access to plaintext customer traffic. These organizations face an uncomfortable choice. They can keep their data confidential and rely on defenses that cannot see application semantics, or they can obtain request-aware protection by handing their plaintext to a mitigation provider.

We systematically analyzed existing DoS defenses, both deployed and academic. None of them simultaneously offers application-layer traffic differentiation, compatibility with strict confidentiality requirements, line-rate operation without per-connection state, and deployment without changes to the network. Our key observation is that two questions are coupled in today's defenses but do not need to be. Whether a packet comes from a legitimate user is a different question from what that packet contains.

Lifting Authentication into the Network

In collaboration with Zürcher Kantonalbank, one of our ZISC partners, we developed Authenticated Browsing (AB). AB answers the first question without ever asking the second. Its core idea is to reuse identity infrastructure that already exists and that users and operators already trust, such as commercial single sign-on services and governmental eID platforms. The result of this application-layer authentication is then carried down to the packet level.

Before her first session, a user, say Alice, links her bank account to one or more Identity Providers (IdPs) trusted by the bank through a standard OAuth 2.0 flow. When she later visits the bank's website, a small user agent on her device transparently fetches a short-lived key bundle from the IdP. From this bundle, it derives an authentication tag for every outgoing packet and attaches it in a compact 24-byte AB header.

Forward-deployed validators, for example at Internet exchange points or scrubbing centers, verify these tags using only symmetric cryptography. Under attack, validators reserve capacity for authenticated traffic and rate-limit or drop everything else, including the botnet's floods. TLS is never terminated, and the payload is never inspected.

Privacy by Design

A three-level key hierarchy makes this possible without revealing who the user is. For each protected application, the IdP derives two keys. The authentication key is forwarded to the validators, while the identity-resolution key never leaves the application. Each key bundle carries a pseudonymous user identifier computed with the identity-resolution key, which rotates every few minutes. Validators can therefore verify packets, but they can neither learn identities nor link a user's traffic across bundle epochs. Because validators derive the per-user key from the packet header alone, they need no per-user or per-connection state and never contact the IdP.

Only the protected service can resolve a pseudonym back to an account. This preserves accountability: if an authenticated account starts misbehaving, the service can identify it and revoke its linkage at the IdP. The attack surface thus shifts from the number of bots an adversary controls to the number of verified accounts it can compromise.

Performance at Line Rate

We implemented the validator as a highly optimized DPDK application that uses AES-NI hardware acceleration and a parallelizable MAC construction (AES-PMAC). On commodity servers, a single validator sustains more than 70 million

authenticated packets per second and saturates a 160 Gbps link. In our line-rate attack-triage experiments, not a single forged packet passed verification, and we reliably caught replayed packets without ever dropping a fresh legitimate one.

The end-to-end results show what this means for a protected service. Under a 160 Gbps forged-packet attack, a web server behind the validator kept serving its full load of about 20,000 requests per second, with at most 0.7 ms added to the 99.9th-percentile latency. Without the validator, the same attack reduced the server to fewer than 18 requests per second.

Internet measurements with RIPE Atlas show that the AB header survives on up to 86% of measured paths, and the QUIC variant sidesteps this issue entirely by carrying the header inside the UDP payload.

Outlook

Authenticated Browsing adds a third option to the binary choice that regulated services face today: application-layer user identification combined with network-layer verification, with payloads kept inside the service's own TLS trust boundary. Because AB requires no changes to routers, transport protocols, or application servers, services can adopt it incrementally. Next, we are integrating AB directly into QUIC stacks and adding per-account rate limiting at the validator to contain floods from compromised accounts.

Further Information

Source code: https://github.com/netsec-ethz/authenticated-browsing

Publication

Felix Stöger, Jelte van Bommel, Jonghoon Kwon, Patrick Bamert, Stefanie Boss, Alexander Eggerth, and Adrian Perrig. "Authenticated Browsing: Reconciling Application-Layer Flooding Defense with Data Confidentiality". Proceedings of the ACM on Networking, Vol. 4, CoNEXT4 (CoNEXT 2026). https://doi.org/10.1145/3856957

Researchers

Felix Stöger, Jelte van Bommel, Dr. Jonghoon Kwon, Alexander Eggerth, Prof. Adrian Perrig (Network Security Group, ETH Zurich)
Patrick Bamert (Zürcher Kantonalbank)
Stefanie Boss (University of Amsterdam)

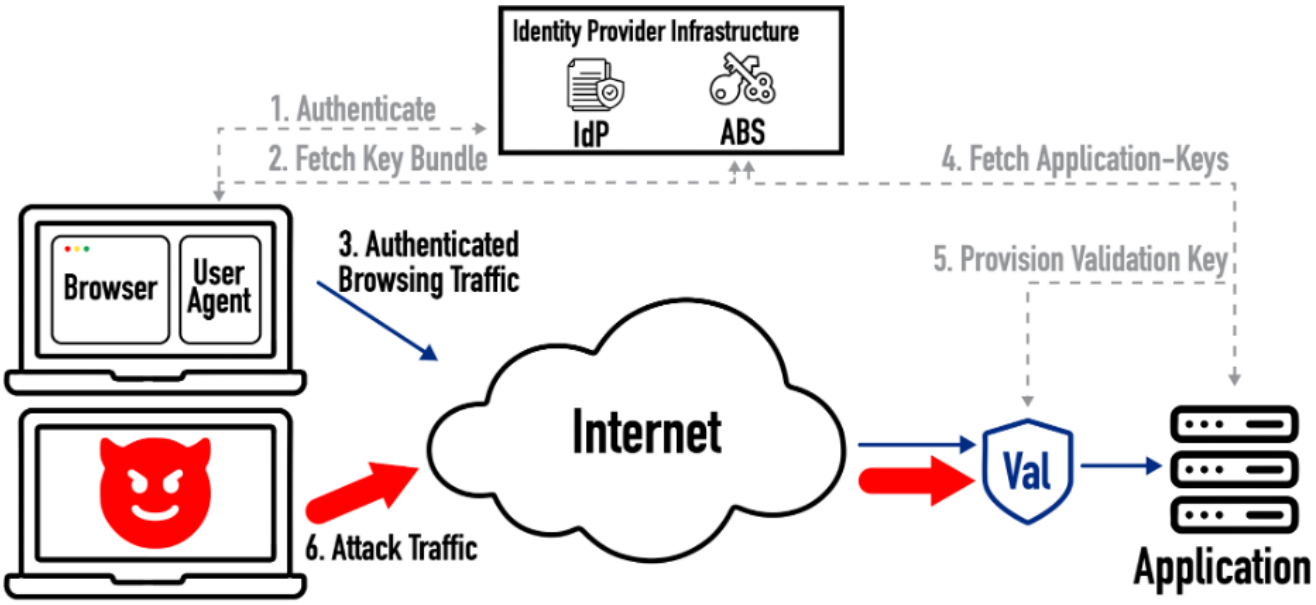


Figure 1: Overview of Authenticated Browsing. Users obtain short-lived key bundles from their Identity Provider and tag every packet. Forward-deployed validators forward authenticated traffic with priority and drop or rate-limit unauthenticated attack traffic, without terminating TLS.

Research Highlights 2026

Skipping Class: Algebraic Attacks on Poseidon2



Dr. Simon-P. Merz

Introduction

A zero-knowledge proof lets one party convince another that a statement is true, for instance that a computation was carried out correctly, without revealing anything beyond the truth of the statement. Inside such a proof every step of the computation must be expressed as additions and multiplications in a finite field, and each such operation has a computational cost. Conventional cryptographic hash functions shuffle bits and use XOR operations, both of which are expensive to express this way. Therefore, new hash functions optimised for this so-called arithmetisation have been designed to use only a small number of field multiplications. Poseidon is probably the most prominent of them. It is used in many academic benchmarks, runs in deployed blockchains, and is a candidate for use in Ethereum itself.

The price of using so few multiplications is that preimages or collisions for such a hash function can be expressed as a system of polynomial equations. Security therefore depends on how hard these equations are to solve. To invite deeper scrutiny of the Poseidon family of hash functions, the Ethereum Foundation runs the Poseidon cryptanalysis initiative, which offers bounties for breaking round-reduced versions of Poseidon2 (the latest version of the hash function when this research was done). We claimed multiple bounties and found weaknesses in the design of Poseidon2 that have since led to changes being made to the design.

Poseidon2

At the core of Poseidon2 is a permutation of vectors of field elements, typically of length 12, 16 or 24. The permutation repeats two operations. A linear layer multiplies the vector by a fixed matrix to mix the inputs, while a non-linear layer raises entries to a small odd power. One pass through the two operations is called a round. In the external rounds at the start and the end of hashing, every entry is raised to the power; in the cheap internal rounds in the middle, only the first entry is. The figure shows the whole permutation. To turn the permutation into a hash function, a mode of operation is needed. Poseidon2 is used in two different modes: in compression mode the input is added to the output of the permutation and the result is truncated; in sponge mode the message is absorbed block by block. Until now, the security of Poseidon2 has been assessed through the study of a single problem about the permutation, known as CICO for "Constrained Input, Constrained Output". Stated simply, the CICO problem is to find an input with prescribed values in certain positions that the Poseidon2 permutation maps to an output with prescribed values in certain other positions. To save multiplications, the designers of Poseidon2 replaced the matrices of the original Poseidon by structured ones. For vectors of 16 elements, the matrix consists of four copies of a fixed 4 x 4 block followed by a fixed permutation of the entries. This cuts the cost of the linear layer by up to 90 percent, and so is attractive in the context of zero-knowledge proof systems. It is precisely this additional structure that our attack exploits.

Skipping Rounds

Solving the polynomial equations for a preimage over all rounds at once is hopeless, since the degree of the equation system grows exponentially with the number of rounds. We therefore split the permutation into two parts, the first few rounds and the rest. For the first part we don't solve equations directly. Instead we describe a family of inputs by a few low-degree polynomials in free parameters, chosen so that all of them satisfy the constraints on the input and so that the state after the first rounds is again given by low-degree polynomials in the parameters. Equations are then set up only for the remaining rounds, in the free parameters. Once this reduced system is solved, the parameters are known and substituting them into the description of the family immediately gives a solution for the full permutation.

The block structure of the matrix makes suitable families easy to find. Vectors of specific forms can pass through the matrix unchanged, and raising their entries to an odd power preserves the pattern. With a few parameters we can steer part of the state into such patterns, and the mixing the matrix is supposed to provide does not happen. In this way we can skip more rounds than any previous attack, up to the first ten rounds for one recommended parameter set for a variant of Poseidon2.

The largest round-reduced instance of the Ethereum Foundation's 2025 bounty programme fell to this strategy: our attack ran in under eleven hours on 64 cores, using a few hundred lines of unoptimised Sage code. The previous record, on an instance about 56 times easier, had taken a different team 40 hours on 1000 cores with an optimised implementation.

Hash Function versus Permutation

The mode of operation links the CICO problem to the hash function. In sponge mode, for instance, the part of the state the message does not touch starts out as zeros and the output is prescribed, so a preimage of the hash function is a solution of the CICO problem, and if the CICO problem is hard the hash function is secure. We modelled the full hash function itself, mode of operation included, and found that both modes give an attacker more room than the CICO problem might suggest. In compression mode, the constrained inputs may be low-degree polynomials rather than constants; in sponge mode, the inputs to the second permutation are outputs of the first, which can be absorbed into the round skip. In both cases a preimage is cheaper to find than a solution of the CICO problem for the underlying permutation. These are the first such examples for Poseidon2, and they show that the CICO problem alone does not determine the security of the hash function.

For collisions (i.e. two inputs having the same hash value), both inputs can be parametrised and the round skip reaches further. For one recommended 128-bit parameter set for a variant of Poseidon2 our collision attack beats the previous algebraic attacks by more than 106 bits (i.e. it is 2^{106} times cheaper to run). In fact, it is the first algebraic collision attack on Poseidon2 that is faster than the corresponding preimage attack.

Consequences

Poseidon2 carries a large safety margin on top of the rounds needed to resist known attacks, and even with a speed-up of over 2^{100}, the recommended parameters stay above the intended 128-bit security level. The design did not survive unchanged, however. We reported our findings to Ethereum's Poseidon initiative before publication, and the designers raised the number of external rounds of a variant of Poseidon2. After the paper appeared, several design choices of Poseidon2, among them the structured matrices that our attacks exploit, were abandoned, and a new version of Poseidon is now under examination.

This work grew out of the Bachelor's thesis of Àlex Rodríguez García from Universitat Politècnica de Catalunya (UPC), conducted in the Applied Cryptography Group at ETH.

Further information

Merz, S.-P., Rodríguez García, À. "Skipping Class: Algebraic Attacks exploiting weak matrices and operation modes of Poseidon2(b)". Full version on the ePrint Archive: https://eprint.iacr.org/2026/306.pdf

Researchers

Simon-Philipp Merz - Applied Cryptography Group, ETHZ
Àlex Rodríguez García - Applied Cryptography Group, ETHZ and Universitat Politècnica de Catalunya

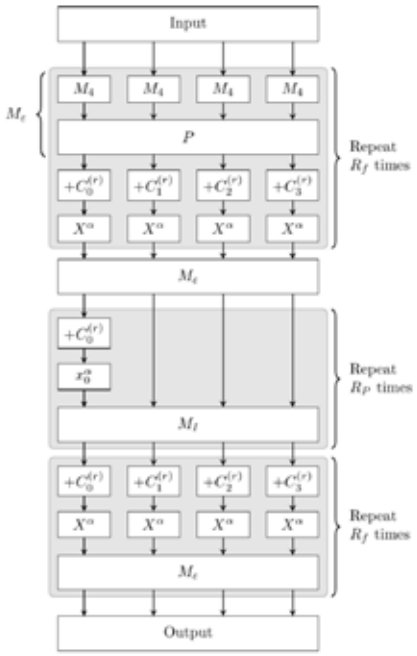


Figure 1: The Poseidon2 permutation on 16 field elements, in the form used in our attacks. The external rounds raise every entry to the power α, the internal rounds only the first. The matrix of the linear layer is written as four copies of a 4 x 4 block M4 followed by a fixed permutation P of the entries.

Research Highlights 2026

Bringing Confidential Computing to Android



Mark Kuhne

Securing Sensitive Workloads on Android

Modern smartphones run a wide variety of workloads. Besides ordinary apps, these include security-sensitive tasks such as payments, authentication, and software updates. Due to their sensitivity, these workloads require protection from other apps and even the operating system, since both may be compromised.

A mobile system such as Android cannot provide this protection by itself and instead relies on virtualization. The Android Virtualization Framework (AVF) runs sensitive workloads in different virtual machines (VMs), which a trusted hypervisor manages. This hypervisor uses software isolation to ensure that Android and the other VMs never share memory. However, this approach has flaws in how isolation is enforced, in how much privilege the hypervisor holds, and in its protection against physical attackers. In our recent work, Aster, we address these issues by leveraging confidential computing to separate trusted and untrusted AVF workloads.

Removing Trust from the Hypervisor

The security of AVF rests entirely on its hypervisor, which causes three problems. First, the isolation is fragile, as a single bug or misconfiguration in the hypervisor pages tables can expose VM memory. Second, the hypervisor is overprivileged. It can access the memory of Android and every VM, so an attacker who manages to compromise it can control all workloads. Third, VM memory is not encrypted or integrity-protected, so an attacker with physical access could read or modify sensitive data directly in memory.

Confidential computing can address all three problems. It protects data in use with hardware-based security mechanisms. Recent Arm CPUs support it through the Confidential Computing Architecture (CCA). In CCA, dedicated hardware enforces isolation, and trusted firmware, rather than the hypervisor, controls it. The firmware also checks all hypervisor requests and can encrypt each VM individually. With Aster, we bring these CCA guarantees to AVF.

Combining Confidential Computing with Android Virtualization

We start by aligning the security requirements of AVF and CCA. This is necessary because AVF and CCA each define their own security goals, and there is no single obvious way to satisfy both at once.

We therefore evaluate four designs that place Android and its VMs across different execution environments, called worlds, on an Arm processor. Our analysis shows that keeping Android in the normal world and executing VMs in the realm world of CCA offers the best trade-off between security and implementation effort. That is because realm world already supports VM lifecycle tasks without requiring larger re-implementations.

However, even this design does not satisfy all AVF requirements, so Aster extends CCA in five ways. First, it makes world isolation mutual, so VMs cannot access Android memory. Second, it provides each VM its own encryption context to defend against physical attacks. Third, it deploys and enforces launch policies, so only approved VMs can execute. Fourth, it integrates AVF attestation, so that VMs can prove they booted correctly. Finally, it adds rollback protection ensuring each VM executes with the correct state.

Impact and Evaluation

We implement Aster across the Android Kernel, Arm Trusted Firmware, and the Android virtualization software and evaluate it on an emulator and an Arm development board. Aster adds a one-time cost of 0.85 seconds to platform boot. At runtime, it slows down CPU, system and I/O operations by about 2-3%.

Four real-world case studies, including apps for key generation and isolated compilation, show overheads of 2.2%. Aster does not affect the performance of OS and app execution in normal world.

In summary, Aster is the first system that brings confidential computing to AVF while staying compatible with Android's security model. We hope that Aster paves path for future work to similarly lift-and-shift other hypervisor-based isolation mechanisms onto confidential computing.

Further information

Website: https://sovereign-smartphone.github.io/projects/aster/

Publication

Bringing Confidential Computing to Android

Mark Kuhne, Supraja Sridhara, Andrin Bertschi, Nicolas Dutly, Fabio Aliberti, Srdjan Capkun, Shweta Shinde

ACM International Conference on Mobile Systems, Applications, and Services (MobiSys 2026).

Researchers

Mark Kuhne (Secure & Trustworthy Systems Group, ETH)

Dr. Supraja Sridhara (Secure & Trustworthy Systems Group, ETH)

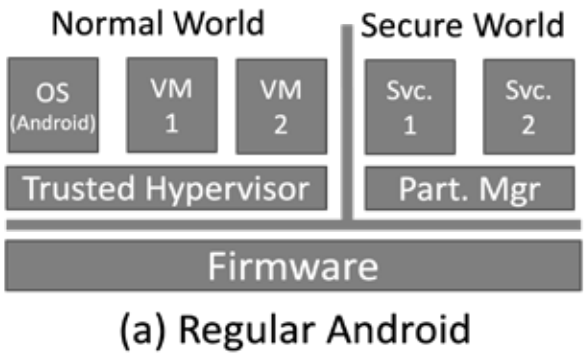
Andrin Bertschi (Secure & Trustworthy Systems Group, ETH)

Nicolas Dutly (System Security Group, ETH)

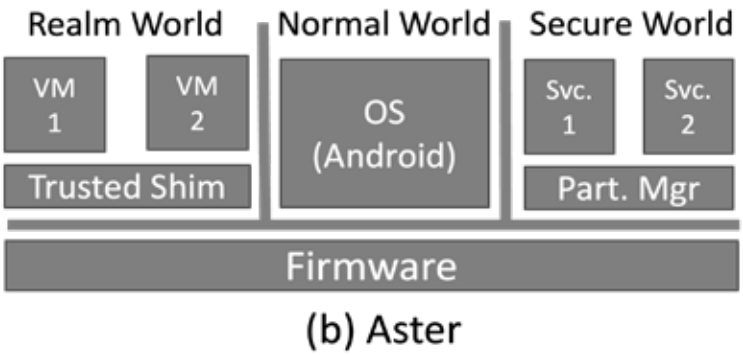
Fabio Aliberti (Secure & Trustworthy Systems Group, ETH)

Prof. Srdjan Capkun (System Security Group, ETH)

Prof. Shweta Shinde (Secure & Trustworthy Systems Group, ETH)



(a) Regular Android



(b) Aster

Research Highlights 2026

Formal Compliance by Design



Dr. François Hublet

Almost a decade after the adoption of the General Data Protection Regulation (GDPR) and over six years after that of its revised Swiss counterpart, the Federal Act on Data Protection (DSG), ensuring compliance with privacy laws remains a major challenge. Year after year, large-scale empirical studies have shown that many popular websites and applications violate basic data protection requirements. The fines imposed on non-compliant entities have had uncertain benefits for consumers, while legal uncertainty and regulatory complexity continue to put a heavy burden on businesses large and small. A promising approach for industries concerned about a high level of conformity is Compliance by Design, in which compliance requirements are integrated into a system's development to ensure that it fulfills legal requirements by construction.

Most existing research and standardization efforts address Compliance by Design only informally, proposing guidelines, auditing methodologies, and checklists that reduce the risk of non-compliance but cannot formally establish conformity. The Formal Compliance by Design paradigm [1] developed at the Information Security Group goes beyond such informal guidelines, offering methods that guarantee conformity with a formalized interpretation of legal provisions. This yields software systems (e.g., web applications) that provably adhere to an extensive set of requirements derived from the law, with all interpretive assumptions needed to bridge the gap between the informal legal text and formal system policies available for review by legal experts.

Realizing Formal Compliance by Design requires overcoming three main challenges. First, software systems and digital regulations are notoriously complex, making it hard to implement compliance safeguards correctly. Second, building compliant systems necessitates rigorous, interdisciplinary collaboration. On the one hand, assessing compliance requires legal experts who rely on specific interpretative methodologies to evaluate the legal constraints applicable to a given technical system. On the other hand, assessing compliance with these legal constraints typically requires a deep understanding of the system, which technical experts (e.g., developers) possess. Third, laws often rely on vague, abstract legal concepts, whereas system specifications are described in terms of explicit, concrete actions. To achieve Compliance by Design, a correspondence between legal concepts and system actions must thus be established and justified in a way that is convincing to both technical experts and legal practitioners.

In our latest work published at the 33rd ACM Conference on Computer and Communications Security (CCS 2026) [2], we provide the first foundations, methodology, and tools for Formal Compliance by Design, with web applications as our primary target. At the core of our methodology is Lex, a novel formalization language that helps experts formalize legal requirements and compile them into enforceable system policies. We call a policy enforceable when, by preventing or causing certain system actions at runtime, one can guarantee that the system always complies with the policy. This ensures that the system never violates the formalized, legally grounded requirements. Explicit documentation of all interpretative steps needed to produce such requirements streamlines internal and external audits, thereby reducing the risk of errors and helping businesses demonstrate due diligence.

In more detail, our methodology has six steps, shown in Figure 1. It involves a legal expert (lawyer) and a technical expert (developer) collaborating to design a policy-compliant system. First, the lawyer and developer pair up to formalize relevant parts of the law in the Lex language. Lex was developed jointly with legal colleagues and supports formalizing laws in a way that closely follows their wording, thereby limiting oversimplification. Then, the two experts refine this formalization by providing a mapping between concrete system actions and legal concepts. They then use the Lex compiler to convert the formalization and mapping into a policy (a large logical formula) that can be evaluated by one of our enforcement tools. Next, the developer instruments the system, making sure that it logs all

relevant actions and that any potentially violating action can be prevented. The system is then run together with our enforcement tools, ensuring compliance with the refined formalization. Finally, the lawyer and developer can audit the formalization, refinement, and instrumentation code at any time, and modify them if necessary. This is particularly important if inaccuracies are identified or regulatory changes occur.

We have applied our methodology and tools in two large case studies, enforcing GDPR requirements in two Python/Django applications. Our 3,200-line GDPR formalization covers a broad set of system-relevant aspects of the law. To the best of our knowledge, our approach is the first to enforce such an extensive set of GDPR requirements in software. We find that both applications can be instrumented for compliance with our formalization while incurring modest additional user-side latency and maintaining a clear separation of concerns between compliance and functional code.

Formal Compliance by Design is not intrinsically restricted to web applications or the enforcement techniques on which Lex relies. Accordingly, we are currently exploring multiple extensions of this approach that would expand the range of legal requirements that can be verified. These include developing techniques that statically verify complex legal properties, such as data minimization; extending our development methodology to support rigorous validation of legal interpretations through a neurosymbolic combination of logical engines and large language models; and expanding our portfolio of formalized provisions beyond the GDPR. As advances in agentic AI systems create opportunities to reduce the cost of verification, we believe that deploying Formal Compliance by Design at scale offers a promising path for industries seeking to demonstrate high compliance standards through verifiable safeguards.

References

[1] François Hublet. Formal Compliance by Design. Information Security and Cryptography. Springer, 2027 (forthcoming).

[2] François Hublet et al. “Lex: Turning Laws into Enforceable Security Policies”. In: Proceedings of the Conference on Computer and Communications Security. 2026.

Researchers

Dr. François Hublet, Prof. David Basin
Information Security Group, ETH Zurich
Prof. Srđan Krstic
University of Copenhagen

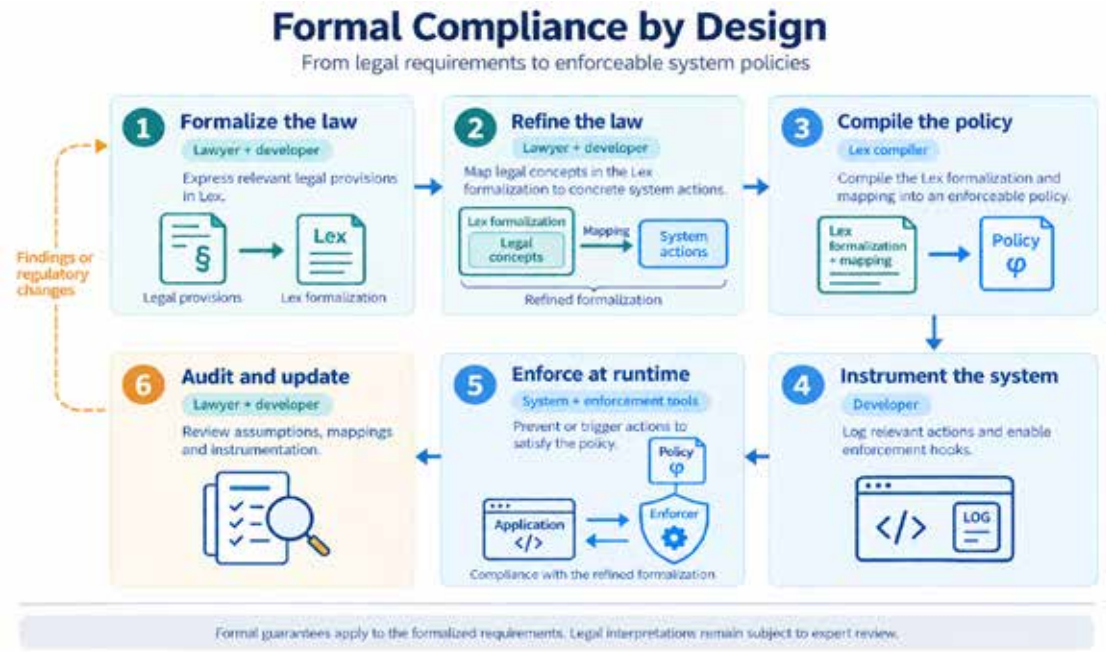


Figure 1

Research Highlights 2026

It's TEEtime: Secure Interrupt Isolation for Normal-world Enclaves on Arm



Friederike Groschupp

Smartphones handle some of our most sensitive information, from private conversations to authentication credentials. Apps rely on the phone's operating system to execute, access hardware, and interact with users. Protecting the sensitive information that apps process therefore usually requires trusting the operating system. For individuals and organizations handling confidential information, this creates a difficult dependency: even a carefully designed application cannot prevent a compromised operating system from reading or manipulating its sensitive data.

Trusted execution environments (TEEs) offer a way to reduce this dependency by isolating sensitive software from the operating system. However, protecting an app's computation and memory is only part of the challenge. Apps also need to interact with the physical world and the user through peripherals such as displays, touchscreens, and cameras. If the operating system still controls these interactions, it can observe or manipulate sensitive data. For example, protecting the cryptographic keys of a messenger that uses end-to-end encryption does little to preserve confidentiality if the operating system can read messages as they appear on the screen.

The TrustZone TEE for Arm platforms already provides mechanisms for protected peripheral access. However, TrustZone protection is typically available only to a few applications that have been chosen or approved by the smartphone manufacturer, as deploying applications there expands the trusted computing base (TCB) of the whole system. Approaches that isolate applications outside Arm TrustZone's secure world offer greater flexibility and deployment options, but they do not support the isolation of peripherals.

With TEEtime, we extend these isolated execution environments, called domains, with direct and protected peripheral access. A domain can contain a small application or a richer software stack with its own operating system. TEEtime uses existing Arm hardware protection mechanisms and the trusted firmware to isolate domains, including the phone's main operating system, from one another. This allows isolated apps to control assigned peripherals without relying on the main operating system to mediate that access.

A central technical challenge is protecting peripheral interrupts: hardware notifications that tell software when an event occurs, such as a touchscreen input or the arrival of a network packet. An attacker who can observe or manipulate another domain's interrupts may learn about its activity or interfere with its execution. Protecting peripheral memory alone does not prevent these attacks because the interrupt controller is shared across domains. TEEtime introduces fine-grained interrupt isolation. Each domain can configure and handle only the interrupts it owns, and those interrupts are delivered only to that domain.

A further challenge is that peripherals, unlike memory, generally cannot be divided freely between domains. Several applications may need the same touchscreen or network interface, yet each peripheral needs a single owner that controls its configuration and interrupts. We therefore present different peripheral access modes and discuss their trade-offs in terms of security, functionality, and implementation effort.

We implement TEEtime on a real smartphone, with a Signal client running in isolation from the main operating system. We also prototype TEEtime on Arm's Fixed Virtual Platform simulator to demonstrate different peripheral access modes. Our evaluation shows that TEEtime requires only small additions to the trusted computing base while maintaining usable performance for the demonstrated smartphone app.

By building on existing Arm hardware, TEEtime provides a foundation for running sensitive workloads alongside

everyday smartphone software. It gives users and organizations greater flexibility in choosing which software they trust with their data and peripherals.

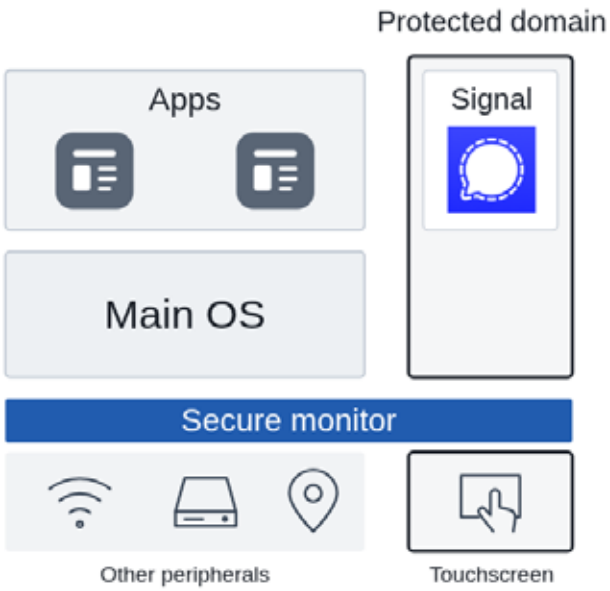
Further information

Publication:
Friederike Groschupp, Mark Kuhne, Moritz Schneider, Ivan Puddu, Shweta Shinde, and Srdjan Capkun. "It's TEEtime: Secure Interrupt Isolation for Normal-world Enclaves on Arm." IACR Transactions on Cryptographic Hardware and Embedded Systems, 2026(1)

Project website
sovereign-smartphone.github.io/projects/teetime

Researchers

Friederike Groschupp, Mark Kuhne, Dr. Moritz Schneider, Dr. Ivan Puddu, Prof. Shweta Shinde, Prof. Srdjan Capkun (ETH Zurich)



TEEtime runs Signal in a protected domain isolated from the main operating system, with protected access to the touchscreen. The secure monitor (trusted firmware) and hardware enforce memory, peripheral, and interrupt isolation between domains.

Research Highlights 2026

Breaking and securing DDR5 DRAM



Dr. Ali Hajiabadi

Memory corruption vulnerabilities such as Rowhammer have demonstrated serious security threats in DDR3 and DDR4 DRAM devices. Rowhammer allows an attacker to induce memory bit flips by repeatedly activating an aggressor row, causing neighboring rows to lose charge faster and eventually experience bit flips. Such bit flips have been exploited for privilege escalation in various settings, including local code execution, browsers, and remote attacks over the network. However, since DDR5 devices were introduced in 2021, a common belief has been that DDR5 is finally secure and deploying stronger mitigations, reinforced by the absence of demonstrated end-to-end Rowhammer exploits on DDR5 systems.

Phoenix: Rowhammer Rises Again in DDR5

Phoenix (S&P '26) challenges this belief and demonstrates that all tested DDR5 DRAM devices from SK Hynix, currently the largest DRAM manufacturer, still rely on security-by-obscurity assumptions and remain vulnerable, with the main difference being that more sophisticated attack strategies are required to bypass their mitigations than the known attack patterns for DDR4 devices.

Mitigations in DDR5

DDR5 devices deploy proprietary Target Row Refresh (TRR) mechanisms that sample row activations and refresh potential victim rows. The analyzed SK Hynix designs use longer and more complex sampling schedules than those assumed by previous attacks, making known DDR4 hammering patterns ineffective. DDR5 also integrates on-die error-correcting codes (OD-ECC), which can correct a limited number of errors but do not prevent the underlying disturbance.

Bypassing the mitigations

Phoenix reverse engineers TRR's sampling behavior and identifies intervals in which aggressor activations can escape detection. It exploits these blind spots to construct two hammering patterns spanning 128 and 2,608 refresh intervals. To execute such long patterns reliably, Phoenix introduces self-correcting refresh synchronization, which detects missed refresh operations and realigns the hammering sequence. Prolonged hammering also allows errors to accumulate beyond ODECC's correction capability.

Practical exploitation

Phoenix triggers bit flips on all 15 tested SK Hynix DDR5 DIMMs, manufactured between 2021 and 2024. Its page-table exploit obtains root privileges on a commodity system under default settings in as little as 109 seconds, with an average of 5 minutes and 19 seconds. Tripling the refresh rate prevents Phoenix from triggering bit flips on the tested systems but incurs an 8.4% performance overhead on SPEC CPU2017, highlighting the need for principled Rowhammer defenses.

Phoenix reception

Due to its impact, Phoenix received media attention from several outlets such as BleepingComputer, Cyber Security News, The Hacker News, and Security Week. Additionally, Phoenix was honored with a Distinguished Paper Award at the IEEE S&P 2026.

Industry solution for principled DDR5 mitigation

The most recent JEDEC DDR5 specifications introduce Per Row Activation Counting (PRAC) as a principled Rowhammer mitigation, extending DRAM arrays with per-row counters to precisely track activations, alongside an ALERT Back-off (ABO) protocol that allows DRAM to request additional time from the CPU for mitigation (i.e., by receiving Refresh Management commands) whenever a row reaches a critical activation count. Such a solution represents an aggressive industry effort to solve Rowhammer once and for all. However, the PRAC specifications leave security-critical design choices undefined, such as what exactly the per-row counters must count, when they should reset, and how critical rows requiring mitigation should be queued and prioritized. Our work, VeriPRAC (CCS'26), fills the specification gaps in PRAC and its implementations to ensure that such implementations can finally put Rowhammer to rest.

Formal PRAC contracts

We introduce PRAC contracts that formalize the key elements of a PRAC mitigation, such as counting, resetting,

and queuing mechanisms together with the JEDEC back-off protocol. Each contract specifies the expected behavior of these elements and a security bound, A_{max}, representing the maximum actual activation count that an optimal attacker can achieve on an aggressor row. This provides an explicit security guarantee against which a concrete PRAC implementation can be verified.

EchoTrail: deriving optimal security bounds

EchoTrail provides a provably optimal attack structure for the analyzed PRAC contracts, covering a range of counting, resetting, and queuing mechanisms proposed in prior work. It systematically identifies silent survival strategies that exploit imprecisions in these elements; for example, basic counting tracks explicit activation commands but ignores implicit activations induced by refreshes. EchoTrail combines silent survival strategies, such as exploiting these uncounted activations and priming decoys before the victim's refresh window, with active decoy hammering. By determining which strategies apply and their optimal ordering and duration, EchoTrail derives optimal security bounds for each analyzed contract.

VeriPRAC: verifying concrete implementations

VeriPRAC automates contract verification by executing EchoTrail instances against a concrete PRAC implementation, tracking both explicit and implicit activations, and comparing the strongest observed attack with the intended contract. It identifies contract violations in both MOAT and QPRAC, as state-of-the-art PRAC implementations. At an ALERT threshold of 64, MOAT permits 7,483 activations despite an intended bound of 99, because its queuing mechanism fails to properly account for refresh-induced activations. For QPRAC, accounting for decoy priming before the victim's refresh window raises the bound from 542 to 567 at an ALERT threshold of 512. These findings demonstrate how formal contracts and automated verification can expose implementation flaws and incorrect security assumptions before deployment.

VeriPRAC reusability

VeriPRAC has been awarded the Distinguished Artifact Award at CCS 2026, alongside an Evaluated---Reusable badge. These recognitions demonstrate the reusability of VeriPRAC as a security analysis tool that significantly exceeds standard evaluation requirements and facilitates reuse by the community. More broadly, as the first formal analysis tool for Rowhammer mitigations, VeriPRAC initiates and supports rigorous and automated security-analysis practices for such mitigations within the community.

In summary, Phoenix demonstrates why DDR5 needs principled Rowhammer defenses; VeriPRAC establishes the formal foundation for designing and verifying secure PRAC implementations as a principled defense.

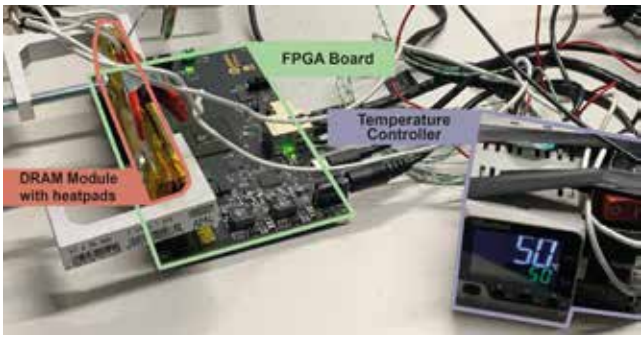
Further information

Phoenix: https://comsec.ethz.ch/phoenix
VeriPRAC: https://comsec.ethz.ch/veriprac

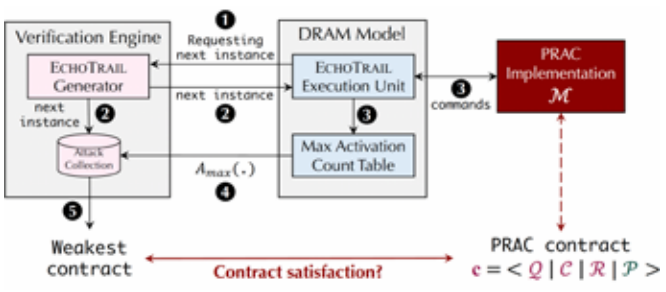
Publications

* Meyer, D., Jattke, P., Marazzi, M., Qazi, S., Moghimi, D. and Razavi, K., "Phoenix: Rowhammer attacks on ddr5 with self-correcting synchronization". In IEEE S&P 2026.

* Hajiabadi, A., Niederer, S. and Razavi, K., "VeriPRAC: Verifiable Formal Contracts for Secure Per Row Activation Counting". In ACM CCS 2026.



Phoenix FPGA testing setup with heating infrastructure



VeriPRAC design to verify contract satisfaction

Research Highlights 2026

Why Johnny Should Not Delegate Email Encryption to Gateways

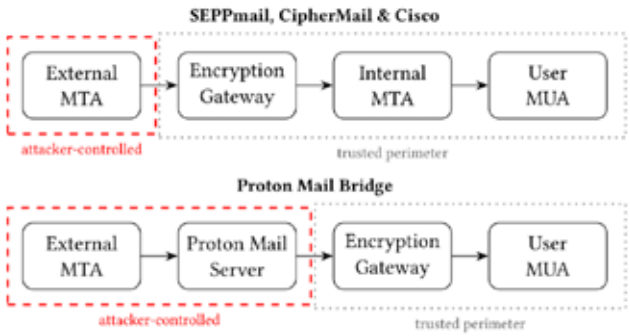


Figure 1



Dr. Matteo Scarlata

Introduction

Despite the rise of instant-messaging platforms, email remains the dominant medium for business communication. A broad array of regulatory frameworks, such as the GDPR and HIPAA, require that sensitive messages be encrypted in transit. Even in the absence of explicit regulation, many organisations adopt email encryption as a matter of policy, driven by contractual obligations or internal risk assessments.

The standards for securing email communication, S/MIME and PGP, have existed for decades. Yet deploying them end-to-end is notoriously hard: every user needs keys, certificates, and a compatible email client.

Email Encryption Gateways

In practice, many organisations therefore centralise cryptographic operations in a server-side encryption gateway: an appliance that sits on the mail path, transparently encrypting and signing outbound messages, and decrypting and verifying inbound ones on behalf of end users. Gateways transform the problem of training every user to configuring one appliance.

Products such as SEPPmail, CipherMail and Cisco Email Security Appliance follow this model. SEPPmail alone counts thousands of installations across banks, insurance companies, government offices and healthcare providers in Europe, particularly in the German-speaking (DACH) region. CipherMail is an open-source alternative in the DACH region and Benelux, whose customers include banks, law firms, pension funds, government agencies, and healthcare providers. Cisco integrates its Email Security Appliance service into a broader email security platform deployed at enterprise scale for hundreds of large business customers, mostly in the United States. Gateways also appear in consumer settings: the Proton Mail Bridge lets users read their end-to-end encrypted Proton mailbox from a conventional email client. It is effectively a gateway, but one that runs locally.

However, gateways introduce a fundamental tension: they perform exactly the operations that end-to-end encryption is designed to keep away from intermediaries. After (automated) decryption, the plaintext travels through internal mail servers, content scanners, and email clients that were never designed to handle it, and that have no idea it was ever encrypted. The gateway also needs to automate signature checks, and to communicate the authentication status of incoming email to the end user. We set out to understand how this impacts the security of S/MIME and PGP when used in a gateway setting. Figure 1 shows the trust assumption for the different gateway products we looked at.

Breaking Email Encryption Gateways

We analyzed four products: SEPPmail, CipherMail, Cisco ESA, and the Proton Mail Bridge. Using the standards, vendor documentation, source code where available, and reverse engineering where not, we modeled how each product receives, decrypts, verifies, rewrites, and forwards mail. We then looked for weaknesses where cryptography meets ordinary mail-processing features: HTML rendering, MIME assembly, bounce generation, and subject-line tagging.

In total, we found 29 attacks (11 against SEPPmail, 12 against CipherMail, 4 against Cisco, and 2 against the Proton Mail Bridge) that allow an attacker to fully decrypt messages or forge their authenticity. We validated every attack with a proof-of-concept against an isolated lab deployment. The attacks fall into four classes.

Reviving EFAIL

In 2018, the EFAIL attacks showed that an attacker who intercepts an encrypted email can trick the recipient's client into leaking its plaintext, for instance by splicing the ciphertext between fragments of an HTML image tag pointing to an attacker's server. Email clients deployed

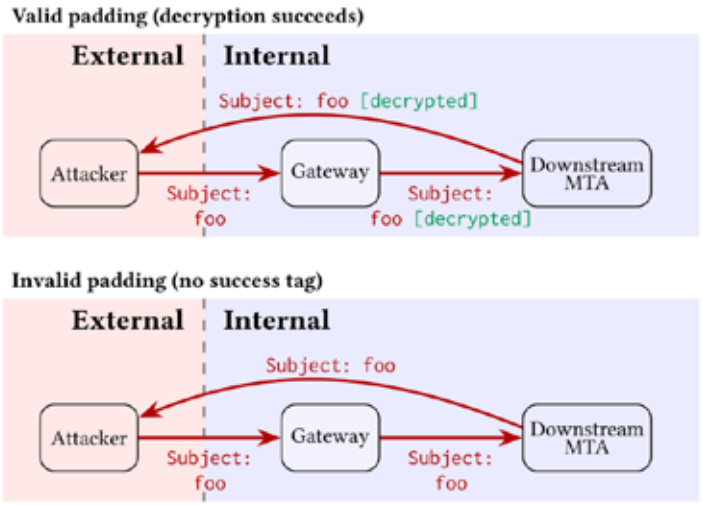


Figure 3

countermeasures, but some of them only activate when the client knows that part of the message came out of an encrypted wrapper. Gateways decrypt emails on a server, and then forward the plaintext content: this means that EFAIL mitigations won't trigger, and it is therefore possible to resuscitate all classic EFAIL attacks.

Mixed Content

A gateway must also put messages back together after decrypting them, and it can get confused when a message mixes encrypted, plaintext, and signed fragments.

We discovered that some gateways don't correctly handle emails with a mix of signed and non-signed parts: an attacker can compose an email out of invisible parts with valid signatures (e.g. by copying them from a real signed message), and an arbitrary part that is visible but not signed. This email will be marked as signed, lending legitimacy to spearfishing attacks for instance (see Figure 2).

We also discovered a new attack vector, which we dub MIME boundary injection. Multipart emails use boundary strings to separate their parts. SEPPmail and CipherMail reuse the attacker's boundaries when re-assembling decrypted mail, while the Proton Mail Bridge derives them deterministically from a server-chosen message ID. Either way, an attacker can hide boundary strings inside encrypted parts so that, after decryption, the message restructures itself and places the victim's secret inside an attacker-controlled image URL.

Bounces as Decryption Oracles

When a mail server cannot deliver a message, it sends a bounce back to the sender. If the message bounces only after having been processed by the gateway, the message being bounced is the decrypted one. In the simplest case, the bounce quotes it in full: an attacker resends a captured ciphertext, provokes a bounce (a full mailbox, or a message the gateway accepts but a stricter internal server rejects), and receives the plaintext in the bounce. CipherMail and Cisco ESA include the full message in their own bounces by default.

Even bounces that quote only headers are enough. Gateways mark successful decryption with a subject tag such as [secure] or by changing the Content-Type header, and either of these headers will show up in the bounce. This one-bit signal is all that is needed for a classic CBC padding oracle attack, since S/MIME still mandates support for unauthenticated AES-CBC (see Figure 3). In the end-to-end setting, padding oracles are impractical: they need thousands of queries, each requiring a human to open an email. A gateway, by contrast, decrypts everything automatically. In our lab, recovering one plaintext byte took between 4 seconds (Cisco) and under 3 minutes (SEPPmail), enough to steal a 16-character code within an hour on all three server-side gateways. Using backwards-compatibility tricks, the attacks even extend to messages encrypted with authenticated AES-GCM.

Spoofing Security Indicators

Since the email client never sees the cryptography, how does a user know a message was signed? SEPPmail and CipherMail append tags such as [signed OK] to the subject line, which the attacker otherwise controls, so the gateways must strip forged tags.

We explored three ways of bypassing or misusing these tags: by composing forged tags in the subject line using Unicode lookalikes (a Cyrillic "е" instead of a Latin "e"), by bamboozling the sanitization code ("[sig[signed]ned]" becomes "[signed]" after one pass), and by padding the subject with hundreds of spaces to push genuine warning tags off-screen in the victim's email client. The lookalike variant works in Apple Mail, Outlook, and Thunderbird alike (see Figure 4). Cisco ESA and the Proton Mail Bridge avoid these attacks only because they give users no cryptographic indicator at all.

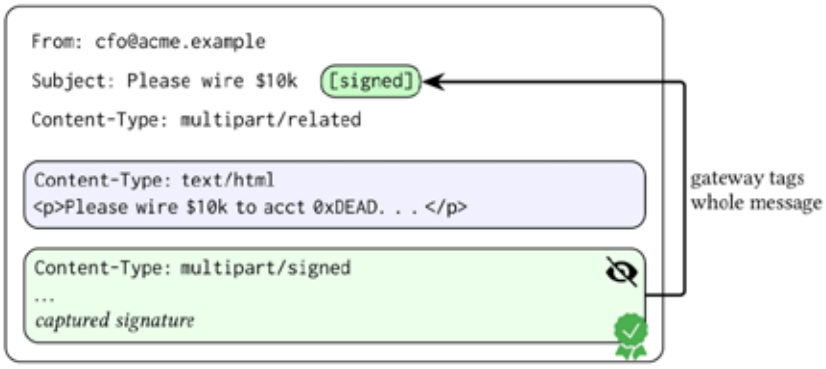


Figure 2

Why Johnny Should Not Delegate Email Encryption to Gateways

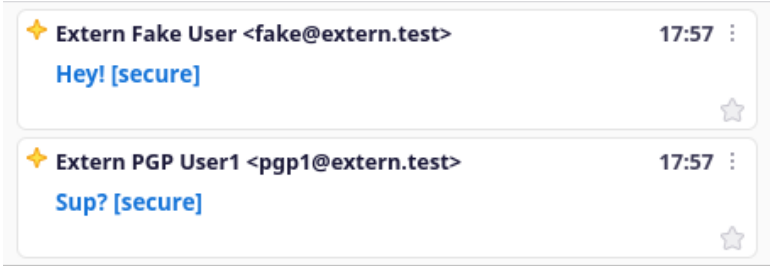


Figure 4

Disclosure and Impact

We disclosed all findings to the vendors, starting with SEPPmail in October 2025. SEPPmail has since shipped fixes across versions 15.0.1 to 15.0.5, CipherMail 6.3.3 addresses all but one of our findings, and Proton fixed MIME boundary injection in Bridge 3.25.0 and awarded a bounty. Cisco published an advisory with two CVEs in September 2026. Our proof-of-concept code is now public.

Lessons Learned

Many of these vulnerabilities are easy to patch individually, but they are not mere software bugs: they stem from the gateway architecture itself. We identify three root causes.

First, the loss of cryptographic signal. Once decryption moves to the gateway, the email client can neither apply its own defenses nor show the user a trustworthy indicator of a message's security status.

Second, reconstruction hazards. Re-assembling mail after decryption forces the gateway into security-critical decisions (which boundaries to use, what counts as a signed message) that are easy to get wrong.

Third, a different threat model. S/MIME and PGP were designed for a human recipient who decrypts one message at a time. Gateways turn the mail infrastructure into an automated decryption oracle, exposing weaknesses that legacy cryptography could get away with until now.

Centralization also creates a single point of failure: a single flaw in the gateway affects every user in the organization at once. Our mitigations make exploitation harder, but many are stopgaps. If email encryption is to deliver meaningful security, the S/MIME standard needs a cryptographic refresh and the backwards-compatibility features of PGP must be discarded. More broadly, our work shows once again that cryptography cannot be analyzed in isolation from the environment in which it is deployed.

Further information

[1] Suter-Dörig, A., Scarlata, M., Paterson, K. G. "Why Johnny Should Not Delegate Email Encryption to Gateways". To appear at ACM CCS 2026. Full version on the Cryptology ePrint Archive: https://eprint.iacr.org/2026/1911

Proof-of-concept code: https://github.com/AppliedCrypto-ETHZ/email-encryption-gateways-pocs

Researchers

Andris Suter-Dörig, Matteo Scarlata, Kenny Paterson
Applied Cryptography Group, ETH Zurich

Research Highlights 2026

Secure mesh messaging for the ICRC: successes and challenges



Dr. Lenka Mareková

Prof. Paterson's group. Our collaboration with the ICRC on "Wireless Mesh Networking for Humanitarian Communication" has entered its second year. In the 2025 ZISC Annual Report, we described the requirements gathering and the initial design phase of the project. Now, we can report on the results as well as on the challenges we have faced on the way.

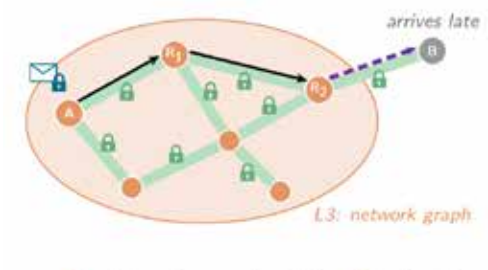
Recall that in this project, we develop a secure mesh messaging application for humanitarian organisations operating in remote locations. The application uses point-to-point Wi-Fi Aware connections to create a local mesh network. It is designed for use in locations such as local ICRC mission headquarters or medical facilities. These are scenarios where ICRC delegates have no Internet connection, cannot bring with them additional hardware (beyond mobile devices), and in some areas, cannot even use encryption; this is why our application has two modes: open mode and secure mode. We designed the system as a modular, 5-layer network stack (see figure).

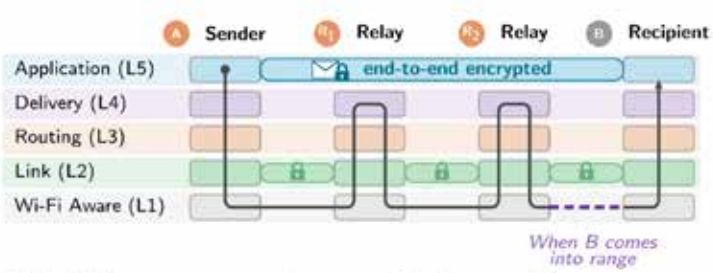
A lot has changed on the project in its second year. We have gained new contributors to the project from interested students, some of whom joined us for a second time. We have focused our work on making our initial prototype application much more robust and performant by improving the implementation of our lowest networking layer (L2), in particular on Android. We also worked on integrating all of the security measures that were designed in the earlier stage of the project, enabled the Android application to function as a background process, and implemented a myriad of other small improvements.

This implementation work uncovered several engineering challenges, which forced us to extend the duration of the project until Autumn 2026. While we were initially hopeful about a newly released implementation of Wi-Fi Aware on iOS, our extensive testing demonstrated that it was not a feasible cross-platform solution in our setting as for example iOS requires manual pairing of devices.

Further, undocumented hardware differences between device models on Android (even among successive models of devices from a single manufacturer) required us to rearchitect our routing (L3) and delay tolerant layers (L4) to enable devices to frequently teardown and reestablish connections with new devices. This allowed us to overcome unforeseen connection limits (which are quite severe on certain device models) and improve network connectivity.

In this way, our work has helped to highlight limitations that are faced by any developers wishing to make use of mesh networking technology on consumer smartphones.





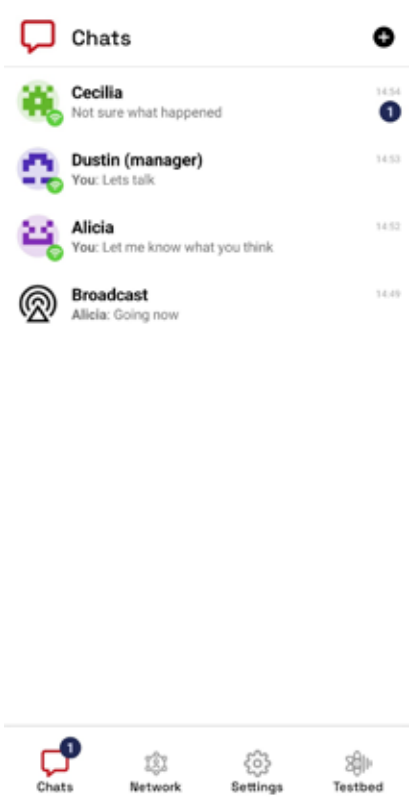
ICRC

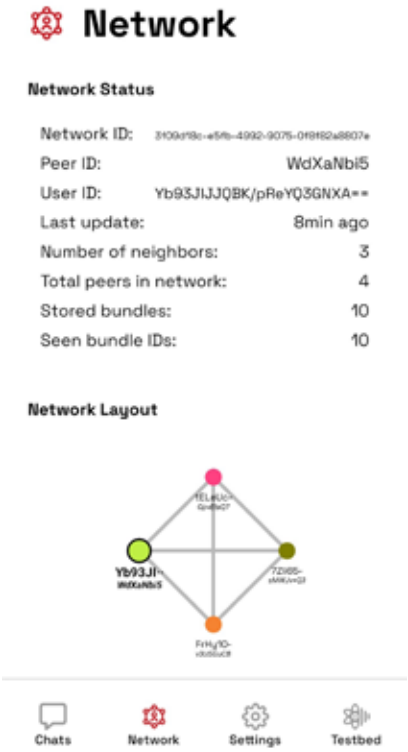
We performed internal testing of the Android application with up to a dozen participants to confirm functionality and measure range, latency, and throughput. In open mode with devices 1-100m apart, we measured an average latency of message round trip times to be within 40-50ms, dropping to around 70ms at 180m distance, and average throughput at around 600 messages per second. We were also able to demonstrate maximum line of sight range of 215m outdoors, and measure the effect of indoor walls within the spaces of ETH Zurich. Additional tests confirmed that packets are delivered to all destinations among respective shortest paths (L3) and that a single device can successfully transport messages between two partitioned groups (L4).

Our work also brings some lessons for researchers wishing to have real-world impact. We advocate for building tools that are founded on concrete user requirements, rather than building out from the state-of-the-art academic work. This can lead to cleaner and simpler designs that have practical deployment potential. For example, our E2EE messaging design (L5) diverges from much of the state-of-the-art work on secure messaging (largely built upon the Signal protocol), because our security and functionality requirements do not include properties such as post-compromise security and deniability. Moreover, some properties, like dynamic groups, could be desirable, but are not compatible with a network that can be partitioned for long periods of time. That is, the results of our requirements gathering determined that we could opt for simpler protocol design choices, which are simultaneously compatible with real-world network limitations.

Future work will involve field testing with ICRC delegates at a larger scale and productisation by ICRC or a third party partner. We are still actively searching for possible partners in this effort.







Research Highlights 2026

Collaborations with the International Committee of Red Cross (ICRC)



Prof. David Basin

Prof. Basin's group. In times of armed conflict, the emblems of the red cross, red crescent, and red crystal are used to mark physical infrastructure. They signal that the assets bearing them are protected under international humanitarian law (IHL), so that military units can avoid attacking them. Since 2020, in collaboration with the International Committee of the Red Cross (ICRC) and as part of our work at the Centre for Cyber Trust, funded by the Werner Siemens-Stiftung, we have been working on extending such protection to digital infrastructure through a digital emblem. In 2021, we finalized a proposal for a digital emblem that we call ADEM: An Authentic Digital Emblem.

ADEM provides a unique combination of security properties: authentication, accountability, and a property that we call covert inspection. Covert inspection means that those wishing to authenticate assets as protected must be able to do so without revealing that they may attack unprotected entities. Moreover, ADEM was designed to fit the context of international diplomacy, as its decentralized design allows states to make sovereign decisions. Using ADEM, one can now mark network-connected infrastructure as protected under IHL. Using Oblivious Digital Tokens (ODTs), which we designed in 2025, one can also signal the protected status of personal devices.

Since its beginnings in 2020, the Digital Emblem project has made tremendous progress. ADEM is now the ICRC's designated prototype for a digital emblem and has thus transitioned from academia into practice. The project has also reached several exciting milestones.

In 2024, all members of the 34th International Conference of the Red Cross and Red Crescent, which include over 190 states, encouraged the further development of a digital emblem. This laid an important diplomatic foundation for the project. In 2025, a working group dedicated to digital emblems was formed at the Internet Engineering Task Force (IETF), one of the world's most influential standards organizations. That same year, ADEM was tested with militaries of NATO member states during a military exercise. In 2026, the ICRC announced Phase 2 of the Digital Emblem project, focusing on further testing and standardization of a digital emblem. These developments illustrate the remarkable impact that the Digital Project has, and we are excited to witness the project's future development.

Education 2026

Center for Computer Science Education

The ZISC collaborates with ETH's **Center for Computer Science Education**[1] (ABZ, "Ausbildungs- und Beratungszentrum für Informatikunterricht") to promote K-12 security education, starting in primary school.

The ABZ was established around 20 years ago with the goal of promoting sustainable computer science education in Swiss schools, with security competences being one of its pillars. For instance, the ABZ conducts training for pre- and in-service teachers to equip them with the necessary skills and tools.

One cornerstone are so-called school projects, where experts conduct exemplary lessons, on security-related or other topics, in schools all around Switzerland to implement on-the-job training for computer science teachers.

In another project, a cybersecurity platform called **CyberQuest**[2] is developed that simulates typical applications with which students interact in their everyday lives. The tool contains a lesson center and so-called target applications which mimic real-world applications. Currently, the following target applications are implemented: a social network, a chat application, a database interface, and an email client.

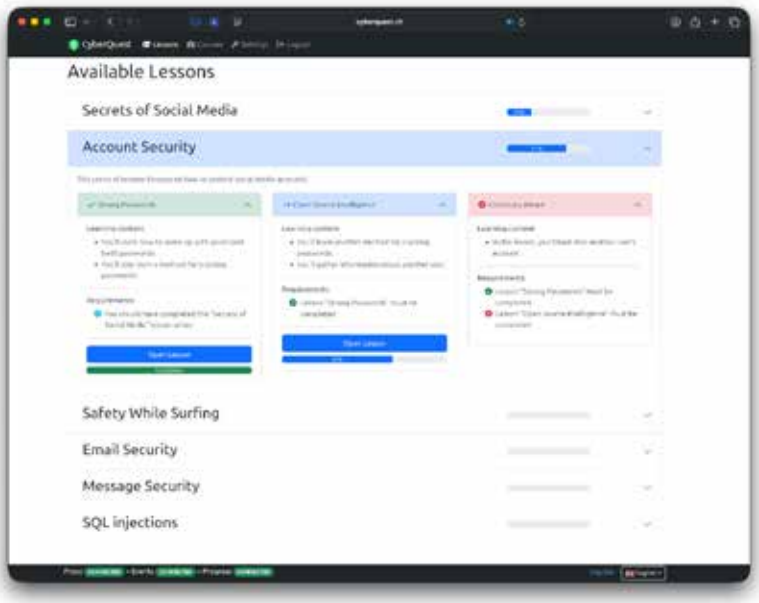
Additionally, students have different helper tools available that allow them to analyze what is going on behind the scenes if they interact with a target application. For instance, a student might be able to inspect the HTTP requests, the cookie data, or certain SQL queries when using the social network. The helper tools are only available to the students if the current lesson requires them, to minimize cognitive load.

The tool is particularly suitable for use in the classroom, as students can work through the lessons at their own pace while still interacting with their classmates. It contains user management capabilities that enable the teachers to manage their students. There are currently over 700 accounts on the platform, which is free to use. Users can also work through the lessons individually and follow a more autodidactic approach.

The following paragraphs give a more detailed overview of the changes since last year's ZISC report.

Lesson Rework

Throughout the last year, we have started reworking the available lessons. With the help of a high school teacher, we are improving the didactic value of the lessons to create a better learning experience for the students. The existing lesson series on "The Secrets of Social Media" and "Social Media and the Internet" have been reworked from the ground up and now provide a more fitting experience for the usage in the classroom, while still keeping the same interactive elements and perspective changes.



1 https://abz.inf.ethz.ch
2 https://www.cyberquest.ch

Account Security Series

As part of the rework, we have also introduced a brand-new lesson series about "Account Security" which consists of three individual lessons, as well as one challenge lesson.

In the first lesson, students learn how to keep their personal social network accounts secure by becoming able to spot and generate strong and weak passwords. Additionally, students are taught about brute-force attacks in terms of their strengths and weaknesses.

In the second lesson, students learn about Open Source Intelligence (OSINT). By examining carefully crafted accounts on the mock social network, students gather sensitive information on their "victim," taking the perspective of a potential attacker.

In the third lesson, students utilize the gathered information of the OSINT part to simulate a dictionary attack to try to guess the password of the victim account. To facilitate this mock attack, we developed a special helper tool that uses common methods of password generation (such as replacing characters by similar-looking numbers or adding two tail numbers) to automatically generate and test passwords based on the provided dictionary.

In the challenge lesson, students go even further by running a simulated SIM swap attack. In a simulated chat environment, students interact with a "support agent" of a telecommunications company. They need to convince this support agent that they are in fact the victim, again utilizing the gathered information of the previous OSINT part. If successful, the students gain access to the SIM card of the victim, making it possible to bypass SMS-based 2FA.

As with the other lessons on CyberQuest, the environment is fully contained to the classroom. No content generated by the students is visible to any user outside of their class. Even if certain vulnerabilities are contained in the system on purpose, additional security measures are implemented to keep a tightly secured platform.

Study: Teachers wanted!

We are currently planning a large-scale study to evaluate CyberQuest in the upcoming spring semester. If you are a high school teacher interested in joining our study, please get in touch with us by sending an email to sven.gruebel@inf.ethz.ch.

Main Research Areas

Sovereign Smartphone

Prof. S. Shinde

The majority of smartphones either run iOS or Android operating systems. This has created two distinct ecosystems largely controlled by Apple and Google—they dictate which applications can run, how they run, and what kind of phone resources they can access. Barring some exceptions in Android where different phone manufacturers may have influence, users, developers, and governments are left with little control. Specifically, users need to entrust their security and privacy to OS vendors and accept the functionality constraints they impose. Given the wide use of Android and iOS, immediately leaving these ecosystems is not practical, except in niche application areas.

We are building a new smartphone architecture that securely transfers the control over the smartphone back to the users while maintaining compatibility with the existing smartphone ecosystems. Our architecture, named TEEtime, implements novel TEE-based resource and interrupt isolation mechanisms which allow the users to flexibly choose which resources (including peripherals) to dedicate to different isolated domains, namely, to legacy OSs and to user's proprietary software. We have shown the feasibility of TEEtime design via a prototype on ARM platform and are working towards building a fully functional phone.

Foundations of Cryptography

Prof. D. Hofheinz

Cryptographic building blocks (such as encryption schemes or zero-knowledge protocols) ensure the secrecy and integrity of information, and help to protect the privacy of users. Still, most actually deployed cryptographic schemes are not known to have any rigorously proven security guarantees.

Our goal is to provide practical cryptographic building blocks that come with rigorously proven security guarantees. These building blocks should be efficient enough for the use in large-scale modern information systems, and their security should be defined and formally analyzed in a mathematically rigorous manner. Specifically, we are interested in the foundations of theoretical cryptography, and in general ways to derive constructions and security guarantees in a modular fashion.

Future Internet Architecture SCION

Prof. A. Perrig

SCION is the first clean-slate Internet architecture designed to provide route control, failure isolation, and explicit trust information for end-to-end communications. SCION organizes existing ASes into groups of independent routing subplanes, called isolation domains, which then interconnect to form complete routes. Isolation domains provide natural isolation of routing failures and human misconfiguration, give endpoints strong control for both inbound and outbound traffic, provide meaningful and enforceable trust, and enable scalable routing updates with high path freshness.

As a result, the SCION architecture provides strong resilience and security properties as an intrinsic consequence of good design principles, avoiding piecemeal add-on protocols as security patches.



Secure Positioning and Localization

Prof. S. Capkun

With the development of new location-based services and the expected deployment of cyber-physical systems (e.g., autonomous cars and drones) the reliance on location and time information in critical applications will only increase. Today's positioning systems are vulnerable to location spoofing by which devices can cheat on their own positions or can manipulate the measured positions of other devices. Proximity-based access control systems are insecure against man-in-the-middle relay attacks (e.g., passive keyless entry system in automobiles, contactless access and payment cards). This problem cannot be fixed by a simple upgrade - existing positioning systems rely on legacy distance measurement techniques and protocols that were designed without security considerations or with security as an after-thought. For more than a decade, we have been designing secure technologies to address these issues.

Trusted Execution Beyond CPUs

Prof. S. Shinde

Modern data centers have grown beyond CPU nodes to provide domain-specific accelerators such as GPUs and FPGAs to their customers. From a security standpoint, cloud customers want to protect their data. They are willing to pay additional costs for trusted execution environments such as enclaves provided by Intel SGX and AMD SEV. Unfortunately, the customers have to make a critical choice—either use domain-specific accelerators for speed or use CPU-based confidential computing solutions.

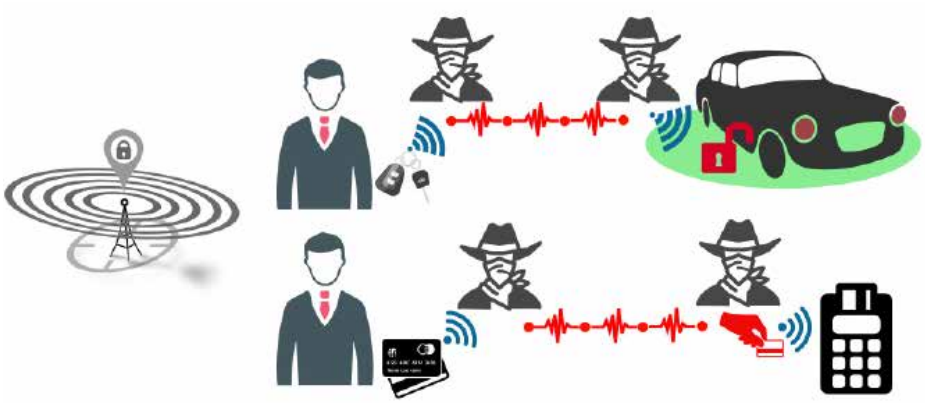
To bridge this gap, we are building data-center scale confidential computing that expands across CPUs and accelerators. Having wide-scale TEE-support for accelerators presents a technically easier solution, but is far away from being a reality. Instead, we aim to provide enclaved execution guarantees for computation distributed over multiple CPU nodes and devices with/without TEE support, which presents security, scalability, and performance challenges.

Machine Learning Security

Prof. F. Tramèr

Machine learning systems are becoming critical components in various industries, yet they face clear security and privacy challenges. Attacks on a machine learning models data can destroy the integrity of the entire system; deployed models can memorize and leak sensitive training data; and models themselves can be copied and stolen.

In our research, we study the behavior of machine learning systems in adversarial settings, to better understand the current limitations and risks of this nascent and booming technology. We then draw on this knowledge to propose new defense mechanisms to safeguard machine learning applications and their users.



Main Research Areas

Access Control

Prof. D. Basin

Access control has wide range of applications, from physical access control, e.g., to buildings, over to access control in single applications, to enterprise-wide access control solutions for an entire company's data management.

Our work covers multiple facets of the challenges in effective access control. We work on languages for efficiently analyzable yet expressive access control policies, techniques for ensuring that no security-critical access control queries are omitted, mining access control policies, and modern systems for access control in physical systems.

Trusted Computing

Prof. S. Capkun

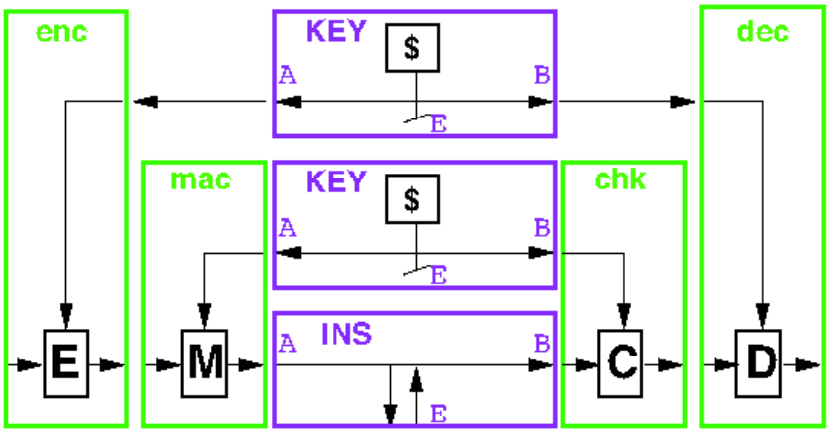
Trusted Computing is a security approach that uses special hardware features to make sure computers run software as expected and protect sensitive data. One key technology in this field is Trusted Execution Environments (TEEs), which create isolated execution environments inside a processor where sensitive operations can be executed safely, even if the rest of the system is compromised. We are interested in all aspects of Trusted Computing: from discovering new attacks on modern processors, analyzing side-channels of security-critical software, to designing better TEEs.

Applied Cryptography

Prof. K. Paterson

Cryptography provides a fundamental set of techniques that underpin secure systems. It includes basic techniques to enable services such as confidentiality and integrity of data in secure communication systems, as well as much more advanced methods such as cryptographic schemes that enable searches over encrypted data.

It draws broadly from theoretical computer science (algorithms, complexity theory), mathematics (number theory, probability) and engineering (both electronic- and software-engineering). Our research in Applied Cryptography brings all of these strands together to produce impactful research that improves the security of today's and tomorrow's cryptographic systems.



Security Protocol Verification

Prof. D. Basin

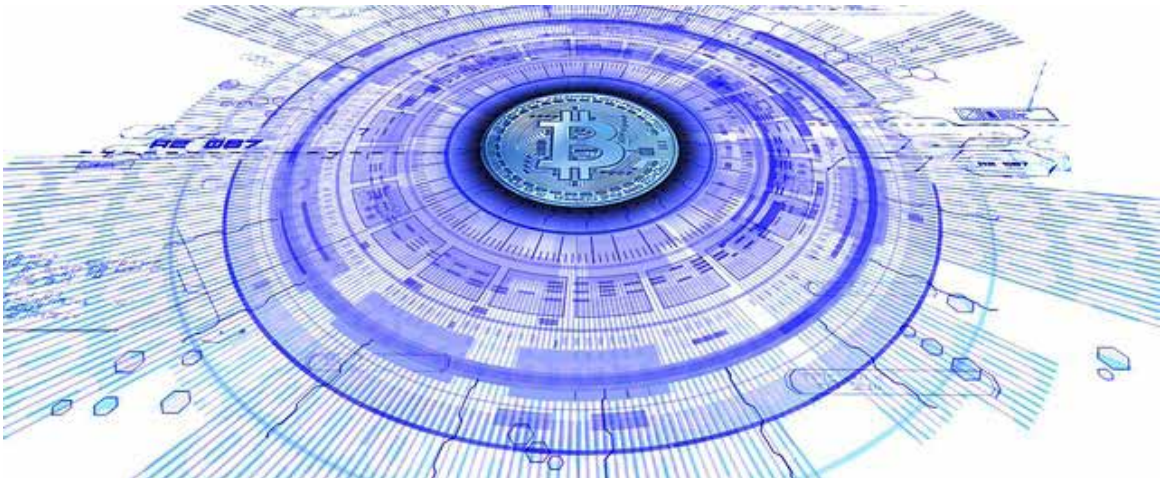
We develop tools for security protocol analysis in the symbolic model, particularly the Tamarin-prover. Security protocols are well-known to be error-prone, thus having a formal analysis enhances trust in the protocol's correctness. Our tools have theoretic foundations guaranteeing their conceptual correctness.

The symbolic model has a high level of abstraction compared to bitstrings over the wire, but provides automation and the ability to consider all modes of operation of a protocol at once. Practical results and impact has been seen in the standardization of TLS v1.3 and the analysis and discovery of serious vulnerabilities in the EMV protocol used for worldwide electronic payments.

Decentralized Systems

Prof. S. Capkun

Decentralized systems can reduce trust requirements and reliance on third parties, thereby enabling greater user autonomy. As part of our research, we explore different applications of decentralized systems, such as blockchains and self-sovereign identity (SSI) systems, and related topics like electronic identities (e-ID) and Central Bank Digital Currencies (CBDCs). Our goal is to develop innovative techniques, build systems that provide new functionalities and to analyze and enhance the security, privacy, and usability properties of existing systems.



Research Projects

Highly Available Communication for Financial Networks

Communication, in particular for critical infrastructures, requires a high level of availability that remains available despite earthquakes, power outages, misconfigurations, or network attackers. One example is the financial industry, which has high requirements on availability to ensure that up-to-date trading information is accessible, that financial transactions are executed within short time windows, and that end customers can execute banking applications online.

The financial industry is generally a prime target for network attackers, mainly because of the importance of availability for banking applications. The importance of availability has led to several extortion attacks in the past, where banks would sometimes pay for attack termination in the short term, rather than protecting their networks for the long term.

To provide high availability and thus to make the financial industry robust against the aforementioned attacks and calamities, we investigate the deployment of multi-path connectivity between branches of one of our ZISC banking partners. In the context of the future Internet architecture SCION, designed and developed at ZISC, we focus on connecting branches over multiple existing links at the same time. We are deploying a multi-path communication system that automatically selects multiple independent, high-quality paths to avoid outages even if some of the independent paths fail.
We are deploying a multi-path communication system that automatically selects multiple independent, high-quality paths to avoid outages even if some of the independent paths fail.

To further increase the resilience against attacks, in particular in the context of DDoS defense and IoT security, our architecture offers the option to hide paths from the public and thus to prevent attackers from flooding such invisible paths. Moreover, we have developed a scalable bandwidth-reservation scheme that protects inter-domain communication by establishing fine-grained resource allocations to ensure no links between networks are saturated.

Further information

Grace Cimaszewski, Francesco Da Dalt, Thomas Moser, and Adrian Perrig,
SCION and cross-border payments: Enhancing security and compliance in distributed ledger networks.
In SNB Working Papers 2025.

C. Krähenbühl, S. Tabaeiaghdaei, C. Gloor, J. Kwon, D. Hausheer, A. Perrig, and D. Roos,
Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure.
In Proceedings of the ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2021.

B. Rothenberger, D. Roos, M. Legner, and A. Perrig, PISKES: Pragmatic Internet-Scale Key-Establishment System,
In Proceedings of the ACM Asia Conference on Computer and Communications Security (ASIACCS) 2020.

M. Legner, T. Klenze, M. Wyss, C. Sprenger, and A. Perrig,
EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet,
In Proceedings of the USENIX Security Symposium 2020.

Researchers

Various members of the Network Security Group.

Industry partner

Zürcher Kantonalbank

SIX

Scalable and Privacy-Preserving Financial Fraud Detection

Financial institutions are legally obligated to screen customers and transactions for fraud under anti-money-laundering regulations. Fraud markers identified by one bank are valuable to others, since illicit funds are typically funneled through accounts at multiple institutions, but banks rarely share these markers with their peers. Banks would be willing to share fraud markers under strict data minimization properties: a bank should learn only which of its own events match, without learning the broader marker set, and without transmitting any of its transaction data to third parties. Private set intersection protocols meet the core matching requirement but are inherently interactive, requiring the receiving bank to send an encoding of its transactions to the marker holder and thereby leaking how often and when it screens against shared markers. Non-interactive alternatives, such as distributing hashed markers, avoid this leakage but cannot currently bound how many times a marker may be matched or let issuing banks revoke markers after distribution.

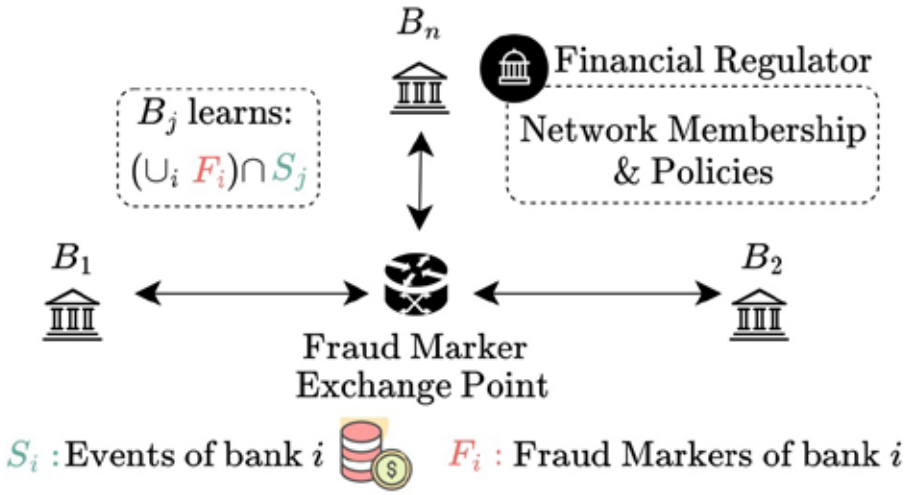
This project designs and analyzes non-interactive protocols for cross-institutional fraud-marker sharing. It aims to construct mechanisms that let a bank locally match its transactions against externally issued fraud markers, with no exchange of transaction data and no interaction-based leakage, while enforcing a bound on how many times each marker can be matched and allowing issuers to revoke markers after the fact. The resulting privacy-preserving protocols should meet the latency and scale requirements of real-world transaction volumes.

Researchers

Prof. Dr. Srdjan Capkun
Dr. Kari Kostiainen
Nicolas Dutly

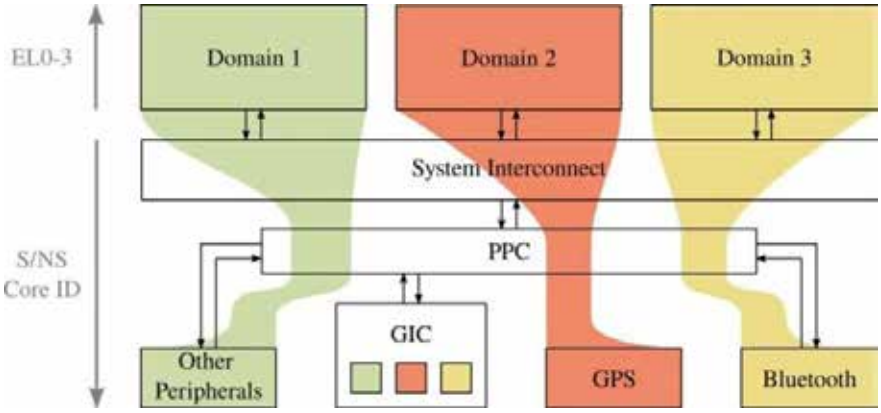
Industry partner

Swissquote
THE SWISS LEADER IN ONLINE BANKING



Research Projects

Secure and Sovereign Smartphone Platform



The majority of smartphones either run iOS or Android operating systems. This has created two distinct ecosystems largely controlled by Apple and Google. While they generally provide phones which are rich in functionality, they also constraint what can be done with the devices: the two aforementioned companies dictate which applications can run, how they run, and what kind of phone resources they can access. End-users have no choice but to fully entrust their security and privacy to OS vendors and accept the functionality constraints they impose. When handling high sensitive information, like with our industry partners, this is far from desirable. Rather, organizations need full control over all the code (including the OS) that can access sensitive data (e.g., messaging apps). However, given the wide use of Android and iOS, immediately leaving these ecosystems is not practical, except in niche application areas.

In this project, we explore the development of a new smartphone architecture that securely transfers the control back to users, developers, and organizations while maintaining compatibility with the rich existing smartphone ecosystems (to maintain a high user experience). Specifically, our solution allows mobile users to install and run security-sensitive apps in isolation, thus protecting them from other applications as well as the operating system. The platform will maintain compatibility with legacy operating systems (e.g, Android) and legacy apps (e.g., as provisioned by Android App Store). Thus, the user can run sensitive and non-sensitive applications on the same device.

Our key design choice is to use a hardware-based mechanism called trusted execution environments (TEEs). Using a TEE, we introduce a small security monitor (SM) that executes at the highest privilege on the phone. The SM uses the underlying hardware features to ensure that the security-sensitive applications are isolated from the legacy OS and legacy applications. The SM is not omnipotent in our approach: It can manage apps and the legacy OS, but it cannot inspect their memory or interfere with their execution. We call this the management-without-inspection primitive. Further, the SM ensures that secure apps can access system resources (e.g., Bluetooth) without the legacy OS being able to inspect or interfere.

Further information

Website: https://sovereign-smartphone.github.io/

It's TEEtime: Secure Interrupt Isolation for Normal-world Enclaves on Arm
Friederike Groschupp, Mark Kuhne, Moritz Schneider, Ivan Puddu, Shweta Shinde, Srdjan Capkun
IACR Transactions on Cryptographic Hardware and Embedded Systems (CHES 2026)

Researchers

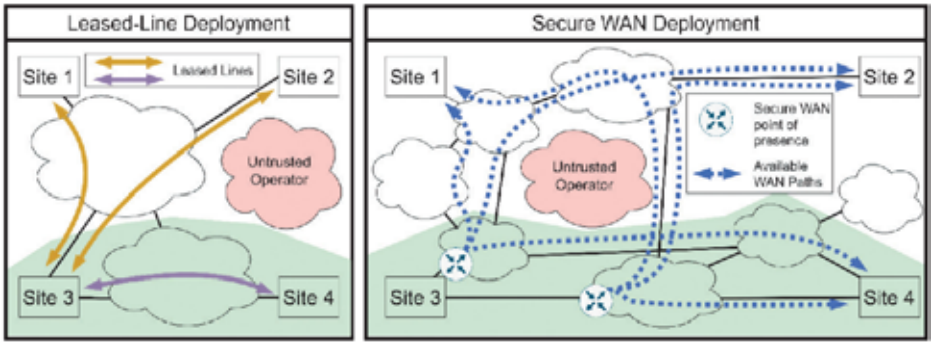
Friederike Groschupp
Mark Kuhne
Dr. Moritz Schneider
Dr. Ivan Puddu
Prof. Shweta Shinde
Prof. Srdjan Capkun

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

Fine-grained Internet Path Control and DDoS Defense



The Network Security group collaborates with armasuisse to build a highly secure WAN infrastructure based on the SCION future Internet architecture, particularly focusing on achieving fine-grained path control and increased communication availability.

In a recent collaboration with armasuisse titled "On Building Secure Wide Area Networks over Public Internet Service Providers," we investigate the challenges associated with constructing a secure WAN over public Internet infrastructure. We found that some of the most difficult challenges include (i) ensuring communication availability between WAN sites, especially during network congestion induced by DDoS attacks, and (ii) achieving fine-grained path transparency and control. We identified several systems that help mitigate such attacks and assessed their technology readiness level (TRL). This year, our focus will be on enhancing the TRL of these systems while also conducting further research to strengthen WAN communication availability.

Fine-Grained Path Control.
We aim to extend SCION's capabilities of path transparency and control to the intra-domain router level. The idea is to enable network operators to communicate information about their internal router topologies in the form of router policies, which are then made accessible to applications along with a set of selectable forwarding paths. This allows each application to choose suitable router policies and encode the chosen policies within its data packets. In our past collaboration with armasuisse, we designed and analyzed such a system (called FABRID), and subsequently worked on a prototype implementation. This year, we will extend its operational maturity to TRL 6 / 7 by rolling out and deploying a prototype on the SCIONLab infrastructure. We will also investigate incentive mechanisms for ISPs, essentially devising business models that would incentivize ISPs to deploy the system in their infrastructure.

DDoS Defense.
We are collaborating on a next-generation bandwidth reservation system. The system's modular and flexible design allows for novel control planes, such as a control plane implementing a marketplace for bandwidth reservations, and introduces features such as the option to request reservations ahead of time. This year, we will extend the operational maturity of this reservation system to TRL 6/7 by rolling out and deploying a prototype on the SCIONLab infrastructure.

Furthermore, we will conduct research to ensure communication reliability for short-lived intermediate-rate traffic, including DNS communication, command and control system traffic, and access to websites. These types of traffic are challenging to protect using bandwidth reservations.

Since our proposed in-network DDoS defense systems have requirements, for example, related to time synchronization and key distribution, we will also design alternative systems to minimize such requirements while still achieving high communication availability.

Other Directions.
In the 5G/6G testbed area, we will continue supporting armasuisse in applying SCION for wide-area experiments. Additionally, in DNS vulnerability analysis and defenses, we will maintain our collaboration through joint supervision of thesis projects. The Network Security Group will also continue supporting armasuisse's security analysis and pentesting efforts. In this task, we will also assist with coordinating with ISPs and vendors, as well as supporting analyses of testbed deployments.

Researchers
Prof. Dr. Adrian Perrig, Dr. Marc Wyss, Jelte van Bommel, Juan A. García-Pardo

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

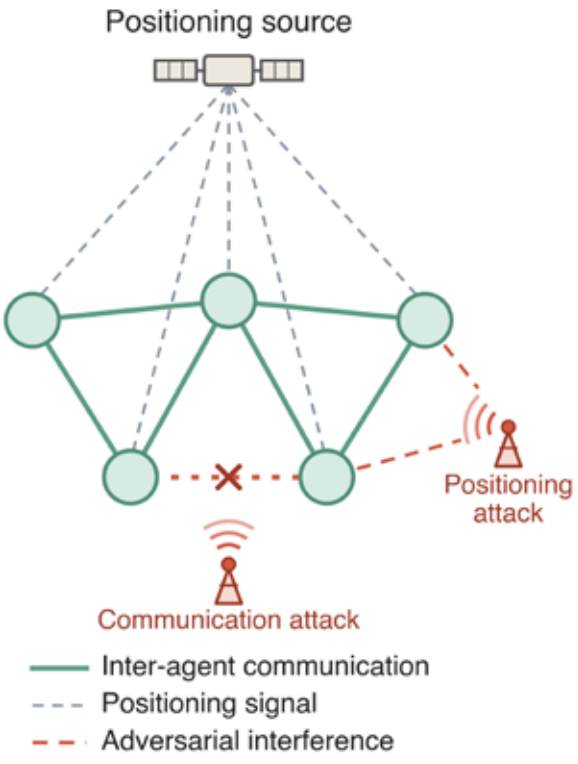
Research Projects

Adversarial Robustness of Autonomous Multi-Agent System

Researchers
Prof. Srdjan Capkun
Prof. Aanjhan Ranganathan

Autonomous multi-agent systems, such as coordinated groups of unmanned aerial vehicles, increasingly rely on decentralized algorithms to achieve collective behaviors without central control. These systems depend on two essential inputs: each agent's ability to determine its own position, typically through satellite-based navigation, and reliable communication between agents. Both inputs can be disrupted or manipulated, raising open questions about how such systems behave when the information they rely on becomes unreliable.

This project investigates how different classes of coordination algorithms respond to disruptions in localization and communication. Using a simulation environment, we systematically characterize how various forms of interference affect collective behavior, with the aim of understanding which information flows are most critical and why certain algorithms degrade under adverse conditions. The resulting analysis is intended to inform the design of more resilient multi-agent systems in future work.



Enhanced Security and Efficiency in Sorting Centers through Computational Robotics

Swiss Post processes large volumes of parcels every day. It is desirable to increase the efficiency of the system, to ensure that parcels are where they are supposed to be at the right time, and to accurately measure the utilization of the system in order to allocate resources correctly. At the same time, it must be ensured that parcels are not tampered with by malicious actors.

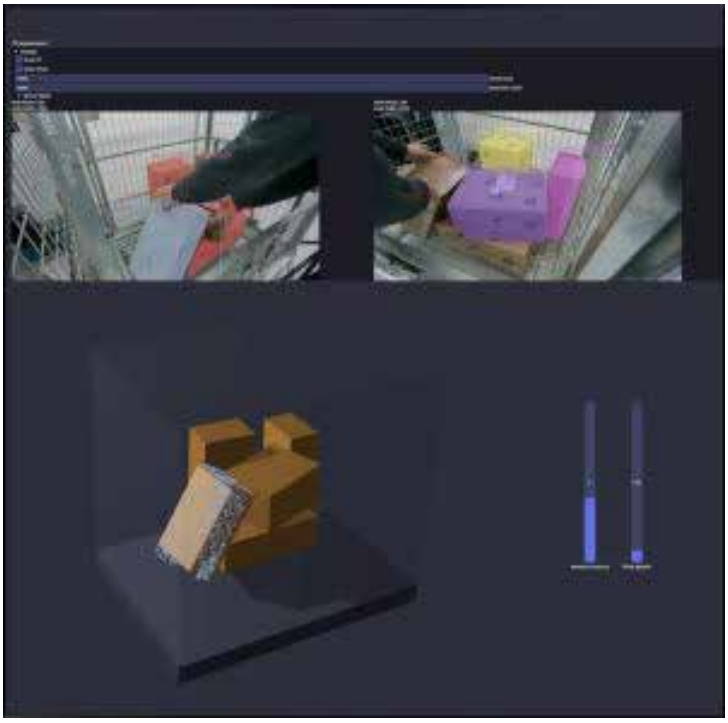
The project aims to enable the tracking and reconstruction of parcels throughout a parcel facility. This enables dynamic routing for improved efficiency. At the same time, such a tracking system can act as a safeguard, i.e., enable anti-fraud and anti-tampering measures by ensuring that a parcel does not deviate from its intended path, and is not removed from where it is supposed to be. Further, the system can accurately measure and report the types of parcels along with the volume for recording transport utilization.

To achieve these goals, we are working on computer vision methods to reconstruct and track parcels, in addition to robotics manipulation in order to handle parcels and orient them in such a way that markers are visible and parcels do not obstruct the views of other parcels.

These robotics problems include simple rearrangement problems that can be handled by single robots efficiently, and multi-robot collaboration problems where multiple robots need to work together to handle large or heavy parcels, in addition to multi-robot coordination problems where multiple robots work together to handle streams of parcels efficiently.

Researchers
Simon Huber (ETH)
Valentin N. Hartmann (ETH)
Stelian Coros (ETH)
Steffen Ochsenreither (Swiss Post)

Industry partner



Research Projects

Cryptography for Secure, Efficient, and Private Digital Identity

Digital identity infrastructure is evolving toward portable, cryptographically-protected digital credentials that users can carry and present to both private and public services. Think of your ID living in your phone, and of showing only what a service really needs to know, and nothing else. Many countries are rolling out electronic identification (eID) schemes. Notably, the Swiss Federal Council is developing a Swiss eID system called swiyu, which will be rolled out in 2026 and aligns with the European Union's effort to introduce an interoperable digital identity framework called the European Digital Identity (EUDI) Wallet.

In these schemes, privacy-preserving properties (unlinkability) are crucial and mandated by legislation, which leads to solutions that involve complex cryptographic protocols. However, the current solutions proposed by all European countries involve batch issuance of one-time credentials, a design that potentially satisfies the desired security properties but fails to achieve unlinkability.

Looking closely at the systems being built

Our first goal is to look hard at the e-ID systems that exist or are being proposed. We are carrying out a rigorous, formal analysis as follows: find out whether there are vulnerabilities and design weaknesses in the current protocols, tell the designers what we find, and eventually prove the security of the patched protocols in a suitable threat model. In this process, we have found and disclosed a recurrent design pattern in the designs of the EUDI Wallet and the Swiss eID that weakens the device-binding (non-transferability) property, enabling credential lending attacks. In other words, a credential that should belong to one person and one device could end up being lent to someone else. We are also disclosing several other weaknesses in European wallets in deployment, such as the Danish eID. In parallel, we are working on developing suitable formal game-based security models to prove the patched constructions secure.

Making the cryptography faster and more flexible

One of the main technical requirements of e-ID systems is to use secure hardware to guarantee non-transferability (device binding) of credentials, but this imposes severe restrictions on the cryptography that can be used, and in particular prevents the deployment of many families of anonymous credentials. A prominent solution is to leverage zero-knowledge proofs, which let you prove that a statement is true without revealing anything beyond that fact. Efforts to bring zero-knowledge proofs into this world, such as the recent Longfellow-ZK library by Google, are encouraging, but they don't extend to primitives other than ECDSA (e.g. ML-DSA), they can't easily be updated, and proving times are still slower than we'd like. We want to find alternative solutions and test the soundness of the current approaches.

On the constructive side, we have developed two novel, efficient proof systems for batch relations, which are particularly well-suited for eIDs. Both works are currently under review. On the cryptanalytic side, we started supervising an MSc thesis that aims to implement and improve existing attacks on the Fiat–Shamir transform, a core technique in many proof systems, which was recently shown to be open to surprisingly practical attacks.

Summary

So far, we have found weaknesses in several eID designs, started the work to fix them and to prove the fixes sound, and laid the first foundations for faster, more flexible zero-knowledge proofs. There is plenty of road ahead, and we look forward to sharing more as the disclosures conclude and the papers start to take shape.

Researchers

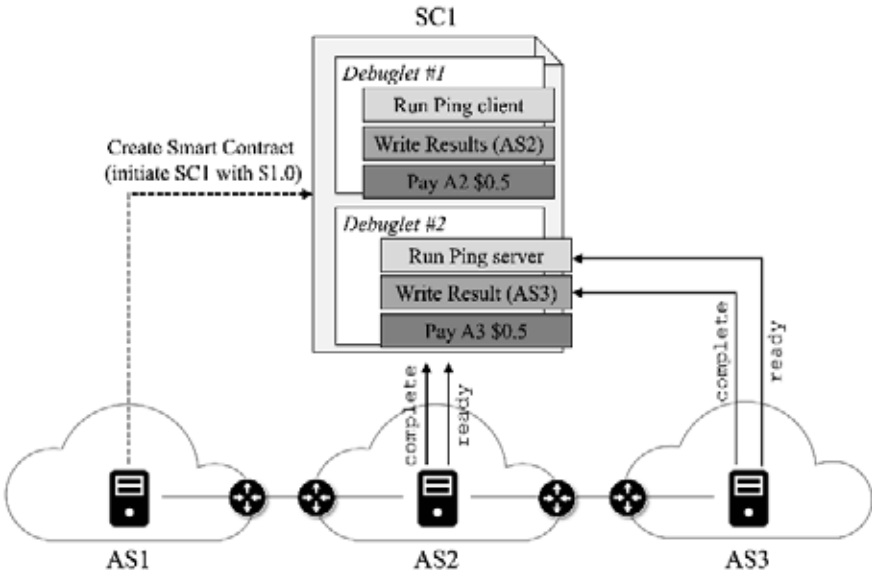
Prof. Kenneth G. Paterson
David Balbás Gutiérrez
Applied Cryptography Group, ETH Zurich

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

Debuglet: Programmable Network Debugging Infrastructure



Debuglet is an advanced distributed network debugging infrastructure designed to enhance end-user debugging on the Internet. Today's end-user debugging is limited to primitive tools like ping and traceroute, leaving users with insufficient data for isolating network faults and no means of external result validation. Debuglet offers a deployable, incentivized architecture that enables near-path network debugging using real data packets and user-defined code for precise and adaptable network performance measurements.

In this system, Autonomous Systems (ASes) deploy small-scale cloud services for network measurement applications. The Debuglet executor, distributed at the edge of ASes, provides a policy-constrained remote code execution environment. Users can perform real-world data-plane operations by distributing and executing network measurement applications at different vantage points.

Users can pay to run their network debugging apps in these environments, and the results can be certified by the deploying AS for third-party verification. This infrastructure significantly improves end-user debugging, accelerates network issue identification, and introduces innovative business models for ASes, with even a partial deployment proving beneficial for users.

Further information

Yi-Min Lin, Mark Csurgay, Jonghoon Kwon, and Adrian Perrig,
TRACER: Flexible and Responsible Distributed Active Network Measurements,
In Proceedings of the IEEE International Conference on Network Protocols (ICNP) 2026.

S. Tabaeiaghdaei, F. Costa, J. Kwon, Pl. Bamert, Y. Hu, and A. Perrig,
Debuglet: Programmable and Verifiable Inter-domain Network Telemetry,
In Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS) 2024.

Researchers

Prof. Dr. Adrian Perrig (ETH)
Dr. Jonghoon Kwon (ETH)
Yi-Min Lin (ETH)
Patrick Bamert (ZKB)

Industry partner

Zürcher Kantonalbank

Research Projects



Sovereignty for Routing, DNS, and Time Synchronization

The Network Security (NetSec) group collaborates with armasuisse to advance sovereignty for critical distributed infrastructures. They propose to pursue the overarching theme of system dependency scoping for availability, where the dependencies of networked systems are transparent and under explicit control. In today's distributed systems, identifying and controlling dependencies is challenging, as reflected in Leslie Lamport's well-known observation: "A distributed system is one in which the failure of a computer that you didn't even know existed can render your own computer unusable." To enable sovereign operation of critical infrastructures, we aim to establish a foundation that provides sovereignty in routing, DNS, and time synchronization.

The SCION Internet architecture provides a strong foundation for sovereign routing and packet forwarding over shared, multi-provider infrastructure. Fine-grained control over packet forwarding is enabled through the FABRID extension, developed in collaboration with armasuisse. Routing sovereignty is achieved through fault isolation domains (ISDs), which allow routing policies to be defined and enforced independently of external entities. However, the creation of ISDs introduces management complexity and additional routing overhead, potentially limiting scalability. We therefore aim to develop light-weight ISDs that substantially reduce the operational cost of ISD creation while preserving SCION's scalability and security guarantees. By enabling dependency scoping for routing, light-weight ISDs provide a foundation for higher-layer services, including DNS and time synchronization.

DNS is a critical dependency for most networked applications, and DNS failures can directly affect their availability and correct operation. As a globally distributed system spanning multiple administrative domains, DNS often relies on entities outside an operator's control. It is therefore essential to restrict dependencies to trusted parties and to protect critical DNS infrastructure from external denial-of-service attacks. We aim to design and implement a DNS architecture that guarantees request processing while eliminating unnecessary external dependencies. The resulting system will ensure continued name resolution for critical infrastructure systems.

Time synchronization is another essential service for distributed systems, where disruptions to availability or correctness can lead to system failures. NetSec aims to design and implement a time synchronization architecture that maintains close system-wide synchronization even when external reference time sources are unavailable or cannot be trusted. Building on their previous work on robust and highly available time synchronization, they will extend these mechanisms to provide sovereign operation and reduce reliance on external entities.

Further information

Marco Pioppini, Marc Frei, Sven Meier, Kevin Schärer, Thomas Schaub, Marc Wyss, and Adrian Perrig,
Resilient Clock Synchronization over Path-Aware Wide-Area Networks,
In Proceedings of the IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS) 2026.

Researchers

Marco Pioppini (ETH)
Eduarda Assunção (ETH)
Marc Frei (ETH)
Dr. Tilmann Zäschke (ETH)
Dr. Marc Wyss (ETH)
Prof. Dr. Adrian Perrig (ETH)

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

Wireless Mesh Networking for Humanitarian Communication

Effective communication is vital for coordinating humanitarian operations. However, regions in which humanitarian workers operate often suffer from frequent internet disruptions or lack connectivity altogether. Internet outages leave smartphones stranded, despite their widespread availability. To address this issue, the project will analyze a novel communication solution for times when cellular networks and non-terrestrial (satellite) networks are unreachable. We will explore the combined use of wireless mesh networks and Delay Tolerant Networks (DTNs).

Researchers

Prof. Dr. Kenny Paterson (ETH)
Lenka Mareková (ETH)
Kien Tuon Truong (ETH)
Shannon Veitch (ETH)
Stefan Mangold (former ETH)

Industry partner

COMITE INTERNATIONAL GENEVE
ICRC



Research Projects

Money Mule Detection

Fraud is a significant problem for users and financial institutions, resulting in high losses. To combat fraud, most banks today leverage machine learning (ML) models. However, the industry's existing techniques are unable to detect fraud sufficiently. The goal of our research is to develop improved ML solutions with stronger fraud detection capabilities, thereby reducing losses for users and financial institutions.

The particular focus of this project is on detection of money mules which have become an essential part of organized crime. Criminal organizations use attack vectors such as phishing campaigns to take over financial accounts. Using compromised accounts, financial assets are sent to money mules, that is, individuals who have been hired to assist in the transfer and laundering of stolen funds. After several such payments hops, the stolen funds typically find their way to cryptocurrency exchanges, offshore accounts and similar locations.

The goal of our research is to detect such malicious financial activities. While recent developments in machine learning provide a good starting point for our research, there are also significant challenges, such the scarcity of labeled training data, short active lifetime of money mule accounts, and constantly changing tactics of criminal organizations. In this research, we tackle the following questions: Professional criminal organizations are likely to adapt their fraud patterns as new detection techniques are introduced - what is the best fraud strategy that remains for fraudsters when an improved fraud detection ML system is implemented? How exactly do the strategies of fraudsters change over time when a fraud detection system is put in place? Can money mule detection systems be adjusted to increase the economic costs for fraudsters? Are there ways to make certain types of fraud economically unattractive?

Researchers
Prof. Dr. Srdjan Capkun (ETH)
Dr. Kari Kostiainen (ETH)
Dr. Daniele Lain (ETH)
Sheila Zingg (ETH)
Carolin Beer (ETH)
Dr. Lucrezia Bruni (Swissquote)

Industry partner

Swissquote
THE SWISS LEADER IN ONLINE BANKING



Fraud Prevention

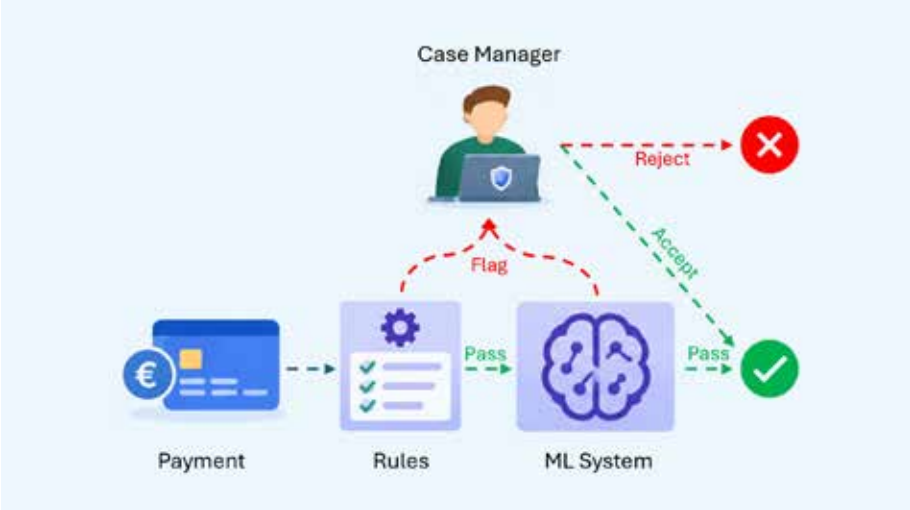
Fraud is a significant problem for users and financial institutions alike, resulting in increasingly high losses. A big challenge for efficient and practical fraud detection is that fraud patterns are very diverse and often difficult to distinguish from benign transactions. For instance, there are scams like romance scam and investment scams that trick a victim users into making a payment to an attacker. As such payments are triggered by the rightful account owner, distinguishing them from benign payments is challenging. In order to detect fraudulent transactions, most banks and financial institutions implement a combination of rule-based fraud detection, ML models for fraud detection, and human case managers to investigate suspicious cases. However, the significant and constantly increasing amount of fraud demonstrates that the existing industry practices are clearly insufficient.

In this project, we study fraud detection with a focus on scenarios where the rightful owner of the bank account (victim) has been tricked to make a malicious payment to the adversary. In particular, we are interested to understand how recent developments in machine learning could be used to address this problem. For example, we investigate how relatively simple tree-based learning models compare to more sophisticated deep learning methods in a use case where lots of training data might be available, but the training data is typically highly unbalance since vast majority of financial transactions are not malicious. We also investigate how state-of-the-art ML techniques could help the day-to-day work of case managers who analyze suspicious transactions.

Researchers
Prof. Dr. Srdjan Capkun (ETH)
Dr. Kari Kostiainen (ETH)
Dr. Daniele Lain (ETH)
Sheila Zingg (ETH)
Carolin Beer (ETH)
Dr. Daniel Donatsch (ZKB)

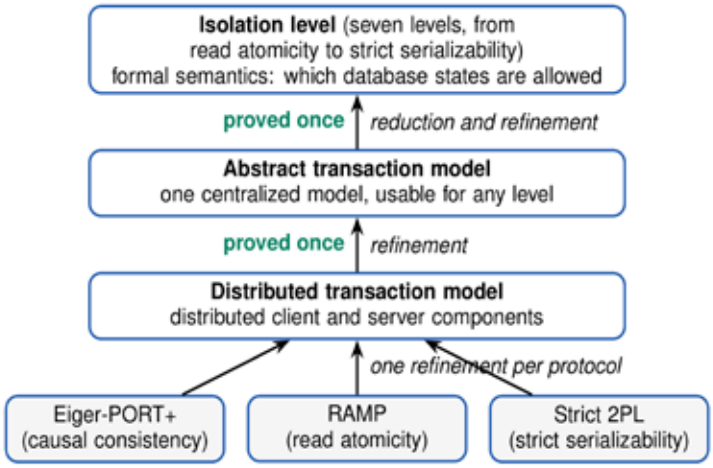
Industry partner

Zürcher Kantonalbank



Research Projects

Verified Isolation Guarantees for Distributed Databases



Databases are the backbone of banking, e-commerce, health care, and nearly every online service. To be fast at scale, distributed databases run transactions concurrently and promise an *isolation level*, a contract that states which effects of other concurrent transactions a client may observe. These contracts are subtle, and the protocols implementing them are highly complex. Isolation bugs keep surfacing in carefully designed and heavily tested production databases, with consequences ranging from lost updates to exploitable inconsistencies, such as a transfer that has arrived on one account but not yet left the other. We therefore need a rigorous way to verify that a protocol satisfies the isolation level it claims.

Our Verlso framework (VLDB 2025) makes isolation guarantees formally verifiable in the proof assistant Isabelle/HOL. It mechanizes the semantics of seven common isolation levels as one abstract model, usable for any level, in which the guarantees of that level hold by construction. A protocol is then verified by proving that a concrete model of the protocol *refines* the abstract model of its claimed level, that is, that every protocol behavior is also a behavior of the abstract model. With Verlso, we verified the strict two-phase locking protocol and uncovered a violation of atomic visibility in the published TAPIR protocol (SOSP 2015).

For weaker isolation levels, which allow more concurrency, refinement alone is sometimes insufficient. The performanceoptimal protocol Eiger-PORT+ (TACAS 2025), for example, commits transactions out of timestamp order to reduce latency, which makes a direct refinement proof infeasible. We therefore combined refinement with a *reduction* argument, showing that executions can be rearranged into an order the abstract model accepts without changing the database states that can arise. The proof succeeded, but the reduction has to be redone for every new protocol, at a cost of thousands of lines.

Our current work removes this cost. A hierarchy of abstract models between the isolation-level semantics and the protocols (see figure) captures the two hardest parts of such proofs once and for all: the reduction needed by the weak isolation levels, and the refinement step from a centralized model to distributed client and server components. Verifying a new protocol then requires a single, simpler refinement proof against the distributed model.

Our verification of Eiger- PORT+ in this setup yielded a proof that is 56% smaller and checks five times faster, and we verified the RAMP protocol (TODS 2016) with a single refinement, where Verlso alone would again have required an ad hoc reduction proof.

Our result provides a reusable foundation for trustworthy databases, where the isolation level a database advertises can be formally verified by a machine-checked proof.

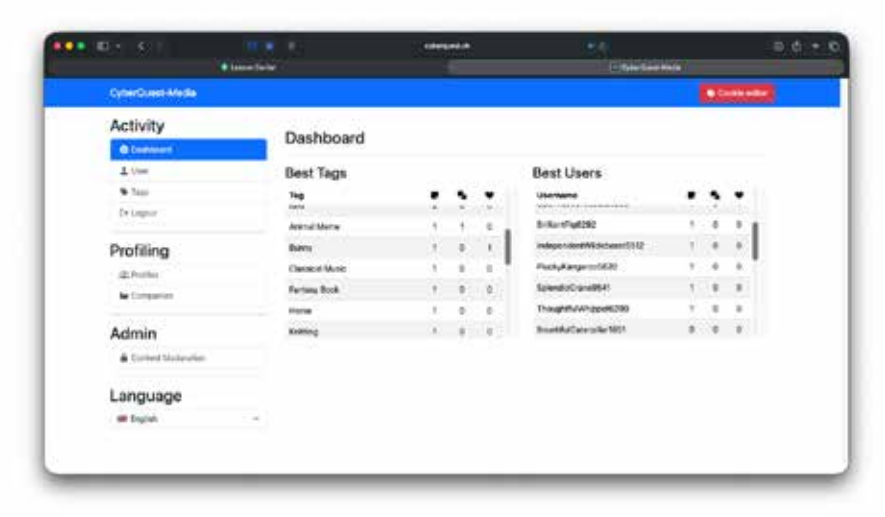
Further information

S. Ghasemirad, S. Liu, C. Sprenger, L. Multazzu, and D. Basin, Verlso: Verifiable Isolation Guarantees for Database Transactions, In Proceedings of the VLDB Endowment, 18(5), 2025.

S. Ghasemirad, C. Sprenger, S. Liu, L. Multazzu, and D. Basin, Pushing the Limit: Verified Performance-Optimal Causally-Consistent Database Transactions, In Proceedings of Tools and Algorithms for the Construction and Analysis of Systems (TACAS) 2025.

Researchers

Prof. Dr. David Basin (ETH)
Dr. Christoph Sprenger (ETH)
Shabnam Ghasemirad (ETH)
Dr. Si Liu (CUHK)



CyberQuest

In today's online world, the importance of cybersecurity is ever-increasing. It feels like a company is hacked every other week. To combat this issue, education in cybersecurity has drastically increased over the past years—on a technical front, as well as on an awareness front.

Various interactive platforms exist which teach different topics of cybersecurity to their users. However, most of these platforms are either focusing on furthering the education of existing industry professionals, or they are used in a leisure setting to host various cybersecurity-related challenges ("capture the flag" competitions, say). The small number of platforms which are actually targeting a younger and/or less experienced audience mostly still have a rather high entry barrier, as they require the setup of additional software and potentially depend on existing programming knowledge.

CyberQuest (https://cyberquest.ch) introduces a novel highly interactive learning platform focussing on teaching cybersecurity topics to high school students. The platform is fully browser-based and does not require any prior programming experience, thus removing the entry barrier of existing platforms.

It consists of three basic components: a lesson center, target applications, and helper tools.

The lesson center holds various lessons which can be worked through by the students, while the target applications serve as playgrounds for the students to apply their newly learned knowledge and experiment with their own ideas. More specifically, the idea behind the target applications is to mimic tools that the students interact with regularly.

In turn, the helper tools provide a means to take a look "behind the scenes" when interacting with the target applications and raising awareness for certain aspects of these interactions.

Currently implemented target applications are a mock social media network, an e-mail client, a chat application, and a frontend for an SQL database.

As a concrete example, the social media network can be set up for high school classes, where teachers receive a privileged account. Students can then interact with each other, create posts, like posts of their friends, etc. At the same time, an HTTP request inspector serves (among others) as helper tool.

Through the interactive lessons presented in the lesson center, the students are guided towards investigating what kind of data is stored about users, or even exploiting the vulnerabilities of the social media network.

The students need to answer questions and solve tasks based on their interactions with the social media network, while learning about cookies and authentication mechanisms, say. They also discover that "deleting an image" may not necessarily mean that the image is in fact deleted from the server. Another goal is to "hack" the platform and hijack a target account.

Students therefore take the role of an "adversary" in order to sensitize them for the pitfalls of social media.

Experiments with teachers and students are ongoing to investigate effective ways to promote cybersecurity knowledge and awareness with CyberQuest being a central element of the lessons.

Researchers
Prof. Dr. Dennis Komm
Prof. Dr. Srdjan Capkun
Dr. Daniele Lain
Sven Grübel

Startup Companies

The companies founded by ZISC researchers are listed here*













*Infineon acquired 3db Access AG (3db) in Fall 2023









Seminar Talks 2026

February 27, 2026
Towards a New Generation of Cryptographic Software
Prof. Peter Schwabe

March 05, 2026
Optimistic MEV in Ethereum Layer 2s: Why Blockspace Is Always in Demand
Dr. Lioba Heimbach

March 24, 2026
The DNS as National Critical Infrastructure in an Era of Geopolitical Tensions
Michael Hausding

April 24, 2026
The Confidential Computing Frontier: A New Paradigm for Compute and AI
Dr. Daniel Moghimi

May 28, 2026
Reconciling Trusted Computing Services and Real-Time Embedded Systems
Prof. Ivan De Oliveira Nunes

June 25, 2026
Why you can't trust your system software, and how you might
Prof. Timothy Roscoe

September 24, 2026
Behavioural Stealth in Autonomous Red Teaming
Daniel López Gala

October 08, 2026
Forging 1024-bit RSA signatures in nearly SNFS time
Prof. Nadia Heninger (UCSD)

October 19, 2026
Sketching the Boundaries: Rethinking Intelligent Analysis of System Provenance
Prof. Zhenkai Liang (NUS)

November 12, 2026
Secure AI Computations at Scale on the Alps Supercomputer
Dr. James Chiang (ETH)

November 26, 2026
CryptanalysisBench: Can LLMs do Cryptanalysis?
Lukas Fluri (ETH)

Further Information

For more information:
https://zisc.ethz.ch/

How to find us:

Postal address

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Universitätstrasse 6
CAB/CNB F
8092 Zurich

Physical address

Entrance to CNB building

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Unversitätstrasse 6
Buildings CNB and CAB, floor F (ZISC OpenLab F100.9)
8006 Zurich
Schweiz

phone +41 (0)44 632 86 89



Contact

https://zisc.ethz.ch/

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Universitätstrasse 6
CAB/CNB F
8092 Zurich
Schweiz

© ETH Zurich, ZISC, December 2026