Scientific Papers within the ReCRED Project

ReCRED Consortium

University of Piraeus

Cyprus University of Technology

cnit

institute imdea networks

Telefónica
Telefónica I+D

verizon

certSIGN

Baker & McKenzie

upcom
Bringing business and IT together

The Productizers

EXUS

wedia

ReCRED

Everything about Research & Innovation in ReCRED

Within the lifespan of the ReCRED project several scientific papers were published in prestigious peer-reviewed journals and conferences. This issue is dedicated to all these publications that aimed at raising awareness for all the scientific work done in the ReCRED project. The ReCRED academic partners provided cutting-edge research and insights on what is important to the scientific community in topics such privacy and anonymity in ICT applications.

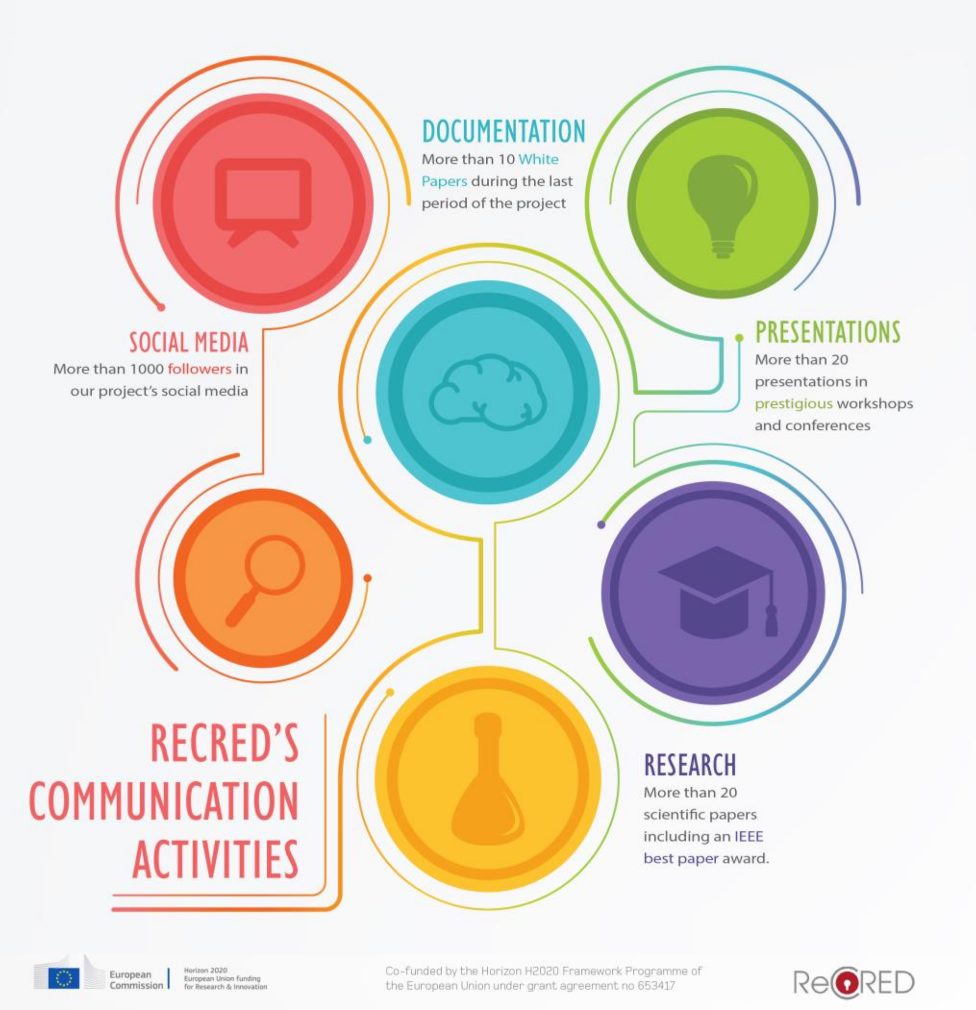
DOCUMENTATION
More than 10 White Papers during the last period of the project
SOCIAL MEDIA
More than 1000 followers in our project's social media
PRESENTATIONS
More than 20 presentations in prestigious workshops and conferences
RECRED'S COMMUNICATION ACTIVITIES
RESEARCH
More than 20 scientific papers including an IEEE best paper award.
European Commission
Horizon 2020 European Union funding for Research & Innovation
ReCRED

Co-funded by the Horizon H2020 Framework Programme of the European Union under grant agreement no 653417

In the next pages, a short summary of each publication is presented but the full-length papers can be found in our website https://goo.gl/1jSxxZ and in the OpenAIRE repository under the ReCRED project https://goo.gl/iNr32p.

More details in
https://goo.gl/TrAcxX

Resource location based on precomputed partial random walks in dynamic networks

Víctor M. López Millán, Vicent Cholvi, Antonio Fernández Anta, Luis López

This paper proposes a mechanism to locate a desired resource in randomly built networks with dynamic resource behaviour (resource instances can appear and disappear) or dynamic node behaviour (nodes can join and leave the network). The mechanism is based on building a total walk with partial walks that are precomputed as random walks and available at each network node. When precomputing each partial walk, information on the resources held by its nodes is stored and associated to it. This information is used by the searches, so that they can jump over partial walks in which the desired resource is not located. Two versions of the mechanism have been described. In the choose-first version, one of the partial walks at the current node is randomly chosen, and then checked for the desired resource. In the choose-first version, all the partial walks of the node are checked for the resource, and then one is randomly chosen among those in which the resource was found. We have presented an analytical model that predicts the expected search length achieved by the two versions of the mechanism. Simulation experiments have been used to validate the model and to assess the effect of resource and node dynamics. We have found that the choose-first version achieves large reductions of the average search length in relation to searches based on simple random walks. These reductions remain significant even in the face of high volatility of resources or nodes. The choose-first version produces larger reductions as we increase the number of partial walks precomputed at each node (at the corresponding extra cost). Results have been found to be very similar for networks with different degree distributions (k-regular, Erdös-Rényi and scale-free). Finally, we have analysed the cost of the PW-RW mechanisms, concluding that the choice of the length of the PW precomputation interval does not have a significant impact on the cost in a wide range of interval lengths.

More details in
https://goo.gl/9vKUzk

Web Identity Translator: Behavioral Advertising and Identity Privacy with WIT

Fotios Papaodyssefs, Costas Iordanou, Jeremy Blackburn, Nikolaos Laoutaris, Konstantina Papagiannaki

Impact on Advertisers: WIT is designed to balance user privacy and utility to advertisers. This means WIT’s interventions should not stop the efeffective targeting of users with relevant ads. We have already shown that with a rather limited number of interventions (10-15%) WIT protects the majority of users which is a sign that damage to ad profiles is small: advertisers still get to see most of a users’ original requests.

Next we show preliminary validation that WIT does not significantly impact advertisers. We manually played back both the original and WIT intervened histories for three users that WIT appeared to have significant impact on. We then checked the interest tags of the histories via the BlueKai registry, which allows users to see their tags collected by the BlueKai tracking service. Although not a perfect measure, if a user’s original and modified browsing histories have similar tags, it indicates a minimal effect on advertising.

More details in https://goo.gl/91QirM

CoVer-ability: Consistent Versioning in Asynchronous, Fail-Prone, Message-Passing Environments

Nicolas Nicolaou, Antonio Fernandez Anta and Chryssis Georgiou

In this paper we have introduced versioned registers and a new property for concurrent versioned registers, we call coverability. A versioned register associates a version with its value, and with each operation that wants to modify its value. An operation may modify the value and the version of the register, or it may just retrieve its value-version pair. Coverability defines the exact guarantees that a versioned register provides when it is accessed concurrently by multiple processes with respect to the evolution of its versions, over a total order of its operations. We combine coverability with atomicity to obtain coverable atomic registers. The successful writes on the register follow the total order of atomicity, while preserving the properties required by coverability. We note that a different total ordering could be used with coverability to obtain other types of “coverable objects”. In fact, we believe it would be interesting to investigate further the use of coverable objects for the introduction of distributed algorithms for various applications. The fact that each operation is enhanced by the version of the object provides the flexibility to manipulate the effect of an operation under some conditions on the version of the object with respect to the version of the operation.

More details in https://goo.gl/TB4sNM

Internet Computing: Using Reputation to Select Workers from a Pool

Evgenia Christoforou, Antonio Fernandez Anta, Chryssis Georgiou, and Miguel A. Mosteiro

This paper presents a reputation-based mechanism of an Internet-based Master-Worker task computing approach. Traditionally, a master process sends tasks, across the Internet, to worker processors. Workers execute, and report back a result. Unfortunately, the disadvantage of this approach is the unreliable nature of the worker processes. Through different studies, workers have been categorized as either malicious (always report an incorrect result), altruistic (always report a correct result), or rational (report whatever result maximizes their benefit). The reputation-based mechanism guarantees that, eventually, the master will always be receiving the correct task result. We model the behaviour of the rational workers through reinforcement learning, and we present three different reputation types to choose, for each computational round, the most reputable from a pool of workers. As workers are not always available, we enhance our reputation scheme to select the most responsive workers. We prove sufficient conditions for eventual correctness under the different reputation types. Our analysis is complemented by simulations exploring various scenarios. Our simulation results expose interesting trade-offs among the different reputation types, workers availability, and cost.

More details in https://goo.gl/s86Ptb

I Always Feel like Somebody's Watching Me. Measuring Online Behavioural Advertising

Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli and Nikolaos Laoutaris

This paper presents a methodology to identify and quantify the presence of OBA in online advertising. We have implemented the methodology into a scalable system and run experiments covering a large part of the entire spectrum of definitions, metrics, sources, filters, etc. that allows us to derive conclusions whose generality is guaranteed. In particular, our results re-veal that OBA is a technique commonly used in online advertising. Moreover, our analysis using more than 50 trained personas suggests that the volume of OBA ads received by a user varies depending on the economical value associated to the behaviour/interests of the user. More importantly, our experiments reveal that the on-line advertising market targets behavioural traits associated to sensitive topics (health, politics or sexuality) despite the existing legislation against it, for instance, in Europe. Finally, our analysis indicates that there is no significant geographical bias in the application of OBA and that do-not-track seems to not be enforced by publishers and aggregators and thus it does not affect OBA. These essential findings pave a solid ground to continue the research in this area and improve our still vague knowledge on the intrinsic aspects of the online advertising ecosystem.

More details in https://goo.gl/Qn3jox

Understanding the Detection of View Fraud in Video Content Portals

Miriam Marciely, Ruben Cuevas, Albert Banchsz, Roberto Gonzalez, Stefano Traverso, Mohamed Ahmedy and Arturo Azcorraz

This work is the first one to pro-pose a set of tools to monitor and audit the view audit systems of online video portals, and enable independent and external parties to measure their performance. The application of the tools and methodology to the view counting behaviour of five different video portals has highlighted some interesting observations. We find that only YouTube deploys a sufficiently discriminative view audit systems for the public view counter. All the other portals studied are susceptible to very naïve view inflation attacks. Clearly, this raises a problem for users with regard to the accuracy of the numbers that are reported by these portals.

The analysis in this paper reinforce the call by industry for (i) consistent and independently measurable principles on how [Supply sources (SSPs/exchanges, ad networks, and publishers)] should identify and expunge fraudulent traffic and (ii) more efficient antifraud mechanisms. In future work, we intend to refine and better scale the tools, and methods developed here, and explore how to make them available to the wider community.

More details in
https://goo.gl/838hLy

Quantifying the Economic and Cultural Biases of Social Media through Trending Topics

Juan Miguel Carrascosa, Ruben Cuevas, Roberto Gonzalez, Arturo Azcorra and David Garcia

In this paper, we show how news in social media manifest through Local TTs in Twitter, analyzing two alternative large-scale datasets of TTs in different countries. We validated our analysis of the leader-follower relationships between countries testing the hypothesis of priority processes in queue systems, finding a power-law distribution of delay times between the appearance of TTs. This finding conveys knowledge about the dynamics of how TTs travel across countries, in an analogous manner as how power-law degree distributions reveal dynamics of preferential attachment or edge copying mechanisms. Applying the statistical physics of priority processes has potential applications to the analysis of communication dynamics in other online communities, from dialogues to collective reactions.

This work shows how content in social media can break international borders through Twitter TTs, revealing that Twitter is used as an alternative communication channel with respect to mass media. On the other hand, we found significant biases with respect to economic, demographic, and cultural factors. This portrays Twitter as a mixed and multipurpose community, in which content can flow without constraints, but also in which mass media have a strong influence and in which economic and cultural factors bias the flow of content.

More details in
https://goo.gl/re5W95

Independent Auditing of Online Display Advertising Campaigns

Patricia Callejo, Ruben Cuevas, Angel Cuevas and Mikko Kotila

This paper illustrates the lack of transparency and accurate information that advertisers are suffering from in the current online advertising ecosystem. This avoids advertisers from accurately assessing the efficiency and quality of their online campaigns. As a result they lack the required information to take decisions and actions to protect, for instance, their Brand Safety.

These results should encourage advertisers to request the Ad Tech industry to standardize the use of independent measurements methodologies, as the one presented in this work. Doing so would allow advertisers to independently assess the quality of their online advertising campaigns as well as auditing the reporting practices of various vendors such as Ad Networks and DSPs.

More details in
https://goo.gl/2xZYPr

Ensuring the Authenticity and Fidelity of Captured Photos Using Trusted Execution and Mobile Application Licensing Capabilities

Kostantinos Papadamou, Riginos Samaras and Michael Sirivianos

This paper describes a security framework that preserves the authenticity of a captured photo and ensures that it remains intact while transferred to a remote server. The key insight is to use a background service that is tied to the photo-capturing application and uses secure key storing and cryptographic computation capabilities offered by the Trusted Execution Environment (TEE) of commodity Android devices. At the same time, this paper leverages Playstore's Licensing Verification Library (LVL) to remotely attest the authenticity of the photo-capturing application at registration time. The evaluation of our prototype implementation demonstrates the efficacy of the proposed framework in terms of performance overhead and usability.

In the last decade, Android devices have become the most popular gadget used in the daily life. According to YAHOO! Tech in 2014, 52% of U.S. smartphones' owners used a handset that runs an Android operating system. Except from an entertainment product, an Android device is also a channel and a storage of sensitive information. At the same time, mobile devices are equipped with a variety of sensors and have become the eyes and ears of a lot of applications, including Internet of Things (IoT) applications, by providing their sensing information. Therefore, the authenticity and fidelity of such information must be ensured while staying in memory or when sent to a remote client or server. Data authenticity and fidelity is crucial for user identification techniques that employ acquisition of physical documentation over the Web. Such techniques request from users to capture a photo of their identity documentation through the camera of their mobile device. The authenticity and fidelity of such photos must be preserved and the service performing the identity verification must be able to determine whether the received photo is authentic. This is because a malicious user may want to upload a modified identity documentation photo pretending to be someone else or use it for various other purposes. Any digital data can be assumed to be authentic if we can prove that it has not been corrupted or modified after their creation. Especially in the area of user identification, where a strict sense of data authenticity is applied, any processing means corruption. The data is considered authentic only if is the outcome of the acquisition process of a real-world document.

Our primary goal is to empower the remote service with the ability to determine about the authenticity of the received photo and to decide whether to accept it or not. Performing cryptographic operations, such as RSA signature in our case, means that you have the private key and you share a public key with the verifier of the signature. In our approach, an application is considered licensed only if it has been downloaded from the Google Play Store and it is a compliant implementation of our framework. We assume that the user does not have root privileges on her device and no other application that runs on the same device can access

the internal storage of our application or use the private keys stored by our application in the TEE. The application may also run securely on rooted devices under some conditions that we will describe in the discussion section.
We have implemented a prototype of our framework for Android on a Nexus 5X, which is powered by a Qualcomm processor with the ARM TrustZone Technology. Our custom registration and photo capture and submission processes have a reasonable execution time without affecting the usability of the application. A short evaluation shows that our approach is feasible and can ensure authenticity and fidelity of captured photos with an acceptable performance overhead.

More details in
https://goo.gl/LQNxBm

On the feasibility of attribute-based encryption for WLAN access control

Claudio Pisa, Tooska Dargahi, Alberto Caponi, Giuseppe Bianchi, Nicola Blefari-Melazzi

User authentication at Wi-Fi Access Points (APs) is becoming an important issue. Wi-Fi APs are indeed ubiquitous, but existing authentication methods such as WPA/WPA2 static pre-shared secret key (PSK), or 802.1X-based online authentication services (e.g., RADIUS servers/proxies) have their theoretical or practical limitations. In a previous work, we proposed WIFAB, a new authentication mechanism which neither requires online backend access control infrastructure, nor relies on a static pre-shared secret key. In this paper, we extend WI-FAB by removing the need for having a central authority for user authentication and credential issuing.
Our main contribution is twofold: (i) adopting decentralized multi-authority CP-ABE, we support the users who have authentication/authorization credentials from multiple authorities. We decouple the user credentials issuing from the management of the WPA2-PSK, so that neither the credential issuing authority can track the users, nor the AP can access the real identity of the users. Considering an extensive attack model, we show that the proposed approach is secure and preserves the privacy of the users. (ii) We provide a real-world implementation of the proposed approach on off-the-shelf embedded hardware to demonstrate its feasibility and efficiency.

More details in
https://goo.gl/Pa4KHx

WI-FAB: attribute-based WLAN access control, without pre-shared keys and backend infrastructures

Claudio Pisa, Alberto Caponi, Tooska Dargahi, Giuseppe Bianchi, Nicola Blefari-Melazzi

Two mainstream techniques are traditionally used to authorize access to a WiFi network. Small scale networks usually rely on the offline distribution of a WPA/WPA2 static pre-shared secret key (PSK); security hence relies on the fact that this PSK is not leaked by end user, and is not disclosed via dictionary or brute-force attacks. On the other side, Enterprise and

large scale networks typically employ online authorization using an 802.1X-based authentication service leveraging a backend online infrastructure (e.g. Radius servers/proxies). In this work, we propose a new mechanism which does not require neither online operation nor backend access control infrastructure, but which does not force us to rely on a static pre-shared secret key.
The idea is very simple, yet effective: directly broadcast in the WLAN beacons an encrypted version of the secret key required to access the WLAN network, so that only the users which possess suitable authorization credentials can decrypt and use it. This proposed approach clearly decouples the management of authorization credentials, issued offline to the authorized end users, from the actual secret key used in the WLAN network, which can thus be in principle changed at each new user's access. The solution described in the paper relies on attribute-based encryption, and is designed to be compatible with WPA2 and deployable within standard 802.11 management frames. Since no user identification is required (access control is based on attributes rather than on the user identity), the proposed approach further improves privacy. We demonstrate the feasibility of the proposed solution via a concrete implementation in Linux-based devices and via relevant testing in a real-world experimental setup.

More details in
https://goo.gl/VHHRc8

On the Feasibility of Attribute-Based Encryption on Internet of Things Devices

Moreno Ambrosin, Arman Anzanpour, Mauro Conti, Tooska Dargahi, Sanaz Rahimi Moosavi, Amir M. Rahmani, Pasi Liljeberg

The Internet of Things (IoT) is emerging with the pace of technology evolution, connecting people and things through the Internet. IoT devices enable large-scale data collection and sharing for a wide range of applications. However, it is challenging to securely manage interconnected IoT devices because the collected data could contain sensitive personal information. The authors believe that attribute-based encryption (ABE) could be an effective cryptographic tool for secure management of IoT devices. However, little research has addressed ABE's actual feasibility in the IoT thus far. This article investigates such feasibility considering well-known IoT platforms--specifically, Intel Galileo Gen 2, Intel Edison, Raspberry Pi 1 Model B, and Raspberry Pi Zero. A thorough evaluation confirms that adopting ABE in the IoT is indeed feasible.

More details in
https://goo.gl/qgJMYt

FEBA: An Action-Based Feature Extraction Framework for Behavioural Identification and Authentication

Luigi Stammati, Claudio Pisa, Tooska Dargahi, Alberto Caponi, Giuseppe Bianchi

While the usage of behavioral features for authentication purposes is gaining more and more consensus in the community, there is less consensus on which specific behavioral traits may be useful in eventually different settings. This calls for flexible tools which the application developer can leverage to automate the extraction and management of behavioral features for identification and authentication. This paper specifically describes a framework called FEBA (Feature Extraction Based on Action), which to the best of our knowledge is the first open-source framework that provides the developer with simple and flexible means to: i) define application-specific actions, ii) recognize actions based on the received raw data, and iii) finally extract the action-specific features. We have built a complete implementation of FEBA, and made it available online to facilitate future research in such context. To prove the performance of FEBA, we provide an experimental evaluation of a use case scenario, i.e., mouse movements feature extraction and pattern recognition. We believe that FEBA will help researchers and developers to design and implement novel behavioral authentication mechanisms.

More details in https://goo.gl/LCwsbj

Protecting sensitive information in the volatile memory from disclosure attacks

Stefanos Malliaros, Christoforos Ntantogian, Christos Xenakis

The protection of the volatile memory data is an issue of crucial importance, since authentication credentials and cryptographic keys remain in the volatile memory. For this reason, the volatile memory has become a prime target for memory scrapers, which specifically target the volatile memory, in order to steal sensitive information, such as credit card numbers. This paper investigates security measures, to protect sensitive information in the volatile memory from disclosure attacks. Experimental analysis is performed to investigate whether the operating systems (Windows or Linux) perform data zeroization in the volatile memory. Results show that Windows kernel zeroize data after a process termination, while the Linux kernel does not. Next, we examine functions and software techniques in C/C++ programming language that can be used by developers to modify at process runtime the contents of the allocated blocks in the volatile memory. We have identified that only the Windows operating system provide a specific function named SecureZeroMemory that can reliably zeroize data. Finally, driven by the fact that malware scrapers primarily target web browsers, we examine whether it is feasible to extract authentication credentials from the volatile memory allocated by web browsers. The presented results show that in most cases we can successfully recover user authentication credentials from all the web browsers except when the user has closed the tab that used to access the website.

More details in https://goo.gl/qxrt2D

Analyzing, Quantifying, and Detecting the Blackhole attack in Infrastructure-less Networks

Christoforos Panos, Christoforos Ntantogian, Stefanos Malliaros, Christos Xenakis

The blackhole attack is one of the simplest yet effective attacks that target the AODV protocol. Blackhole attackers exploit AODV parameters in order to win route requests, and thus, attract traffic, which they subsequently capture and drop. However, the first part of the attack is often neglected in present literature, while the majority of attempts in detection focus only on the second part of the attack (i.e.,packet drop). This paper provides a comprehensive analysis of the blackhole attack, focusing not only on the effects of the attack, but also on the exploitation of the route discovery process. As a result, a new critical attack parameter is identified (i.e., blackhole intensity), which quantifies the relation between AODV's sequence number parameter and the performance of blackhole attacks. In addition, a novel blackhole detection mechanism is also proposed. This mechanism utilizes a dynamic threshold cumulative sum (CUSUM) test in order to detect abrupt changes in the normal behavior of AODV's sequence number parameter. A key advantage of the proposed mechanism is its ability to accurately detect blackhole attacks with a minimal rate of false positives, even if the malicious node selectively drops packets.

More details in
https://goo.gl/fuUvUS

A Security Evaluation of FIDO's UAF Protocol in Mobile and Embedded Devices

Christoforos Panos, Stefanos Malliaros, Christoforos Ntantogian, Angeliki Panou, Christos Xenakis

The FIDO (Fast Identity Online) Universal Authentication Framework is a new authentication mechanism that replaces passwords, simplifying the process of user authentication. To this end, FIDO transfers user verification tasks from the authentication server to the user's personal device. Therefore, the overall assurance level of user authentication is highly dependent on the security and integrity of the user's device involved. This paper analyses the functionality of FIDO's UAF protocol and identifies a list of critical vulnerabilities that may compromise the authenticity, privacy, availability, and integrity of the UAF protocol, allowing an attacker to launch a number of attacks, such as, capturing the data exchanged between a user and an online service, impersonating a user at any UAF compatible online service, impersonating online services to the user, and presenting fake information to the user's screen during a transaction.

More details in
https://goo.gl/dvwzoP

(U)SimMonitor: a mobile application for security evaluation of cellular networks

Christos Xenakis, Christoforos Ntantogian, Orestis Panos

The lack of precise directives in 3GPP specifications allows mobile operators to configure and deploy security mechanisms at their sole discretion. This may lead to the adoption of bad security practices and insecure configurations. Based on this observation, this paper presents the design and implementation of a novel mobile application named (U)SimMonitor that captures and analyses the security policy that a cellular operator enforces i.e., the invocation and employment of the

specified security measures to protect its users. (U)SimMonitor achieve this by executing AT commands to extract network related parameters including encryption keys, identities, and location of users. Using (U)SimMonitor as our basic analysis tool, we have conducted a set of experiments for three mobile operators in Greece in a time period of 9 months. The obtained results allow us to quantify, compare and evaluate their applied security as well as pinpoint a set of generic critical observations. Numerical results and security measurements show that mobile networks have poor security configurations and practices, exposing subscribers to several attacks.

More details in
https://goo.gl/xae79s

Evaluation of Cryptography Usage in Android Applications

Alexia Chatzikonstantinou, Christoforos Ntantogian, Georgios Karopoulos, Christos Xenakis

Mobile application developers are using cryptography in their products to protect sensitive data like passwords, short messages, documents etc. In this paper, we study whether cryptography and related techniques are employed in a proper way, in order to protect these private data. To this end, we downloaded 49 Android applications from the Google Play marketplace and performed static and dynamic analysis in an attempt to detect possible cryptographic misuses. The results showed that 87.8% of the applications present some kind of misuse, while for the rest of them no cryptography usage was detected during the analysis. Finally, we suggest countermeasures, mainly intended for developers, to alleviate the issues identified by the analysis.

More details in
https://goo.gl/pjWFc8

RiSKi: A Framework for Modeling Cyber Threats to Estimate Risk for Data Breach Insurance

Angeliki Panou, Christoforos Ntantogian, Christos Xenakis

Historically, the financial benefits of cyber security investments have not been calculated with the same financial discipline used to evaluate other material investments. This was mainly due to a lack of readily available data on cyber incidents impacts and systematic methodology to support the efficacy of cyber investments. In this paper we propose an innovative, cyber investment management framework named RiSKi that incorporates detection and continuous monitoring of insiders societal behavior, to the extent permitted by the law, to proactively address implied anomalies and threats and their potential business impact and risks. Moreover, it provides access to published security incidents data to enable businesses to advance their understanding of cybersecurity and awareness of the threats and consequences related to cyber breaches, and, eventually, enable faster recovery from an event. RiSKI armed with the above information, employs a methodology, and develops a supporting scenario-based cyber investment tool, for quantifying the benefits of cybersecurity investments against the many ways that potential cyber risks can affect the operation of a business.

More details in
https://goo.gl/yBfvs4

Commix: Automating Evaluation and Exploitation of Command Injection Vulnerabilities in Web Applications

Anastasios Stasinopoulos, Christoforos Ntantogian, Christos Xenakis

Despite the prevalence and the high impact of command injection attacks, little attention has been given by the research community to this type of code injections. Although there are many software tools to detect and exploit other types of code injections, such as SQL injections or Cross Site Scripting, there is no dedicated and specialized software that detects and exploits, automatically, command injection vulnerabilities. This paper proposes an open source tool that automates the process of detecting and exploiting command injection flaws on web applications, named as COMMand Injection eXploiter (Commix). We present and elaborate on the software architecture and detection engine of Commix as well its extra functionalities that greatly facilitate penetration testers and security researchers in the detection and exploitation of command injection vulnerabilities. Moreover, based on the knowledge and the practical experience gained from the development of Commix, we propose and analyze new identified techniques that perform side-channel exploitation for command injections allowing an attacker to indirectly deduce the output of the executed command (i.e., also known as blind command injections). Furthermore, we evaluate the detection capabilities of Commix, by performing experiments against various applications. The experimental results show that Commix presents high detection accuracy, while at the same time false positives are eliminated. Finally and more importantly, we analyze several 0-day command injection vulnerabilities that Commix detected in real-world applications. Despite its short release time, Commix has been embraced by the security community and comes preinstalled in many security-oriented Operating Systems (OS) including the well-known Kali Linux.

Co-funded by the Horizon H2020
Framework Programme of the European
Union under grant agreement no 653417