Ann Arbor, Michigan
info@agitalabs.com

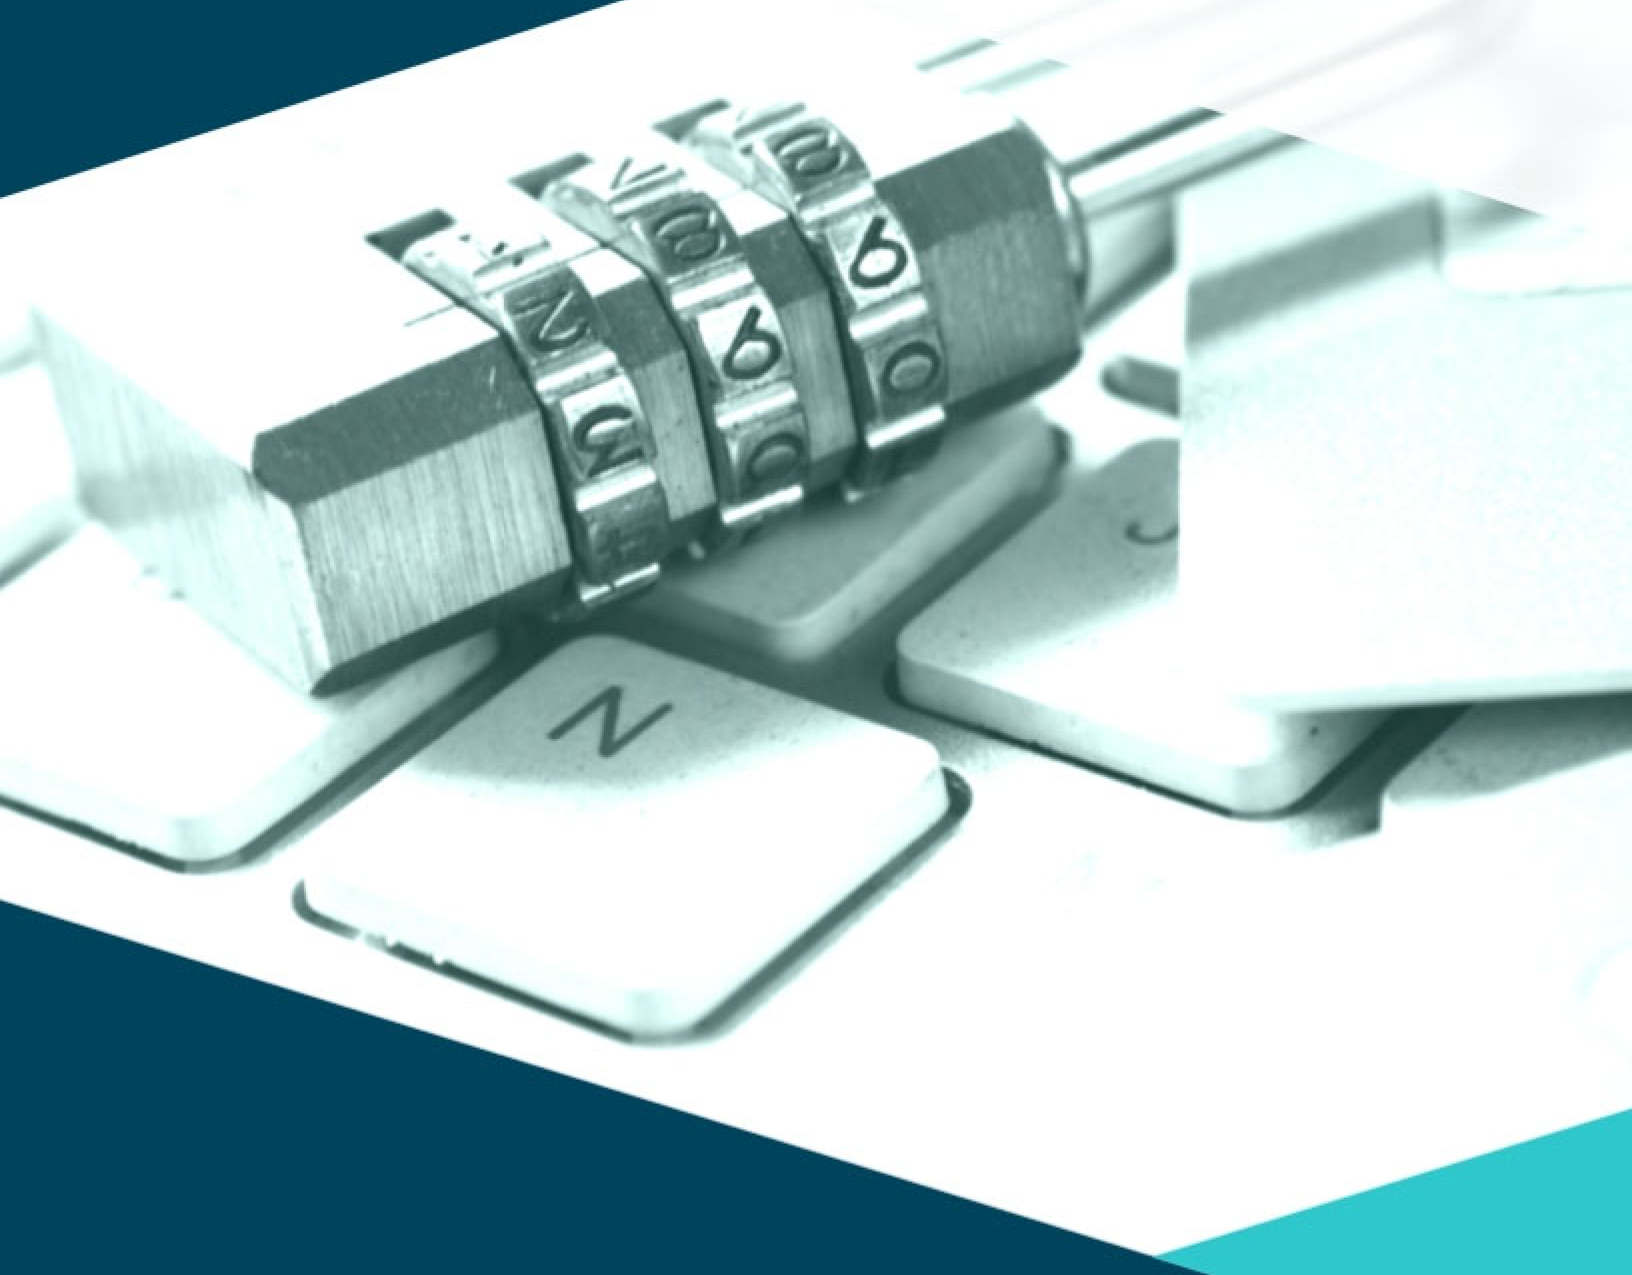

ØZone Privacy Platform

DURABLE DATA SECURITY FOR AN UNTRUSTWORTHY WORLD

Table of contents

EXECUTIVE SUMMARY
Data Security

Data security, or rather lack of it, is one of the key challenges for information technology today. IBM estimates that in any year, 1 in 4 sites on the web will be hacked and their data stolen, and the cost of each incident will average about $4M. And, the problem is getting worse each year, with new advanced security attacks making systems increasingly vulnerable to attackers.

Why is data always at risk? Two primary reasons:

- First, modern systems use software to protect data, and software is always eventually hackable.
- Second, sensitive data resides in resources shared with untrusted parts of the system, leading to side-channel attacks that can leak sensitive data.

Agita Labs' ØZone Privacy Platform utilizes novel sequestered encryption (SE), which eliminates all trust in software and any sharing of critical resources.

The technology is being deployed into the Amazon AWS and Microsoft Azure clouds where it can be used to protect sensitive data and the programs that process that data.

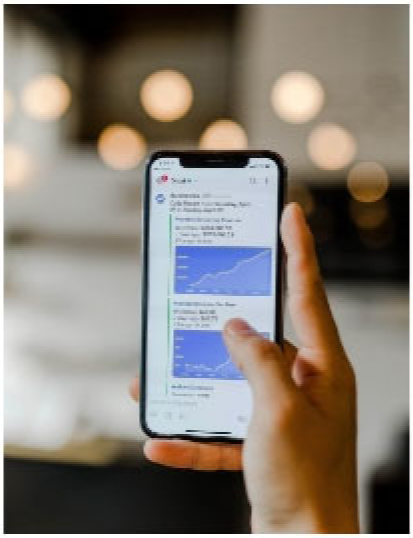

TODAY'S SECURITY ISN'T VERY SECURE

With the growth of cloud computing and IoT, data security has never been more important. With cloud computing, we hand over our personal and private information to cloud providers and their customers, and we can only hope that they steward our data well. For IoT devices, we install them everywhere in our home, cars, and workplace, and then we trust these devices to not spy on us. We are all extending much trust to these computing systems today, and in many cases they are letting us down.

The world of cloud computing is replete with examples of data breaches and poor stewardship of sensitive data. Perhaps the most impacting of these stories for Americans was the Equifax breach in 2017 [1]. In this sophisticated data breach, the financial credentials (e.g., social security numbers and home addresses) of nearly one-half of all Americans were stolen by as-yet unidentified attackers. The attackers used advanced code injection techniques to breach systems owned by Equifax, from which they exfiltrated a myriad of sensitive data. Many similar stories can be found in the world of IoT. IoT devices are notorious for having little or no security. As an example, a North American casino was famously attacked through the fish tank in its lobby [2]. To protect its sensitive financial data, they maintained a high-security intranet. On this intranet, however was a wireless fish tank thermometer that attackers hacked into via the wireless interface. From that base of operations, they were then able to exfiltrate data from systems on the intranet.

The cost of poor data security to business operations is enormous. IBM estimates that in any year, 1 in 4 businesses on the web will be hacked and their data stolen, and the cost of each incident will average about $4M [3]. Similarly, Amazon encourages companies to budget 1/4 of their cloud operating expenses for data breach recovery expenses. To counter this trend, the data security market has exploded (to $160B per year) with many service-oriented technologies that promise to identify breaches when they happen. Still, even sophisticated protections (like those that were used by Equifax) can be breached by talented and persistent attackers.

WHY IS DATA SECURITY SO WEAK?

Hackable Software is Protecting Data

To understand why it is today impossible to stop data breaches, we need to first understand that in today's computing environments, software is protecting data, and software can always be hacked. To be unhackable, a program's security defenses cannot be overcome by any known or unknown attack. It's the "unknown" qualification that makes this declaration unreasonable, since an unknown attack could be a zero-day attack that is known only by some clever attacker, but not by the security community. Even more challenging, an "unknown" attack could be an attack invented in the future.

Worrying about attacks invented in the future may seem like we are grandstanding for effect, but future attacks are a clear and present danger. As an example, consider the Spectre and Meltdown attacks that were disclosed in January 2018 [4]. These attacks used microarchitectural priming techniques to craft custom mispeculation execution code sequences that accessed protected data and converted their values to signals in the cache. These clever attacks allow most memory protections (e.g., virtual memory and Intel SGX) to be side-stepped without any trace that the data breach has occurred. After the attacks were announced, researchers demonstrated that these same attacks could side-step the memory protections of systems that had been built more than 20 years earlier. Thus, an unhackable program written in the mid-90's would have had to anticipate the invention of Spectre and Meltdown in 2018! Consequently, when software is protecting data, that data is always at risk.

Of course, the fundamental deficiencies of program-based data security have been known for quite some time, and the security community has been developing powerful technologies that have partially addressed the problem. Two of the most impactful technologies are Trusted Execution Environments (TEE) and Homomorphic Encryption (HE). Both of these technologies reduce the amount of software that needs to be trusted to ensure data security, which ultimately improves the overall state of data security.

Trusted Execution Environments (TEEs):

To reduce the risk of a program getting hacked, TEEs wall off the execution of trusted software from the rest of the system (including the operating system and drivers) through encryption and selective isolation. Examples of these technologies include Intel's SGX and ARM's TrustZone. Using TEEs, developers should only have to worry about the durability of software running inside the TEE, ignoring the rest of the system's vulnerabilities. **The approach falls short in two important ways:**

- Software remains inside of the TEE, which can be hacked.
- Trusted data resides in the same resources utilized by the main core's software, leading to resource sharing that enables side-channel attacks that can exfiltrate secrets out of TEEs.

As such, TEEs provide better security, but by no means solve all data security challenges, as they are today plagued by a wide variety of effective attack scenarios.

Homomorphic Encryption (HE):

- In a stunning achievement, IBM researchers introduced in 2009 the concept of fully homomorphic encryption, which was the first fully capable data security technology that eliminated all trust in software. HE works by performing mathematics directly on unknown encrypted values. A HE ‘add’ operation, for example, performs addition on two encrypted values to produce an encrypted addition result. The user that originally encrypted the source operands can later decrypt the result of the HE ‘add’ operation to safely and securely learn the value of the privacy-enhanced addition. Consequently, it becomes possible to compute on secret data without ever knowing its decrypted values.

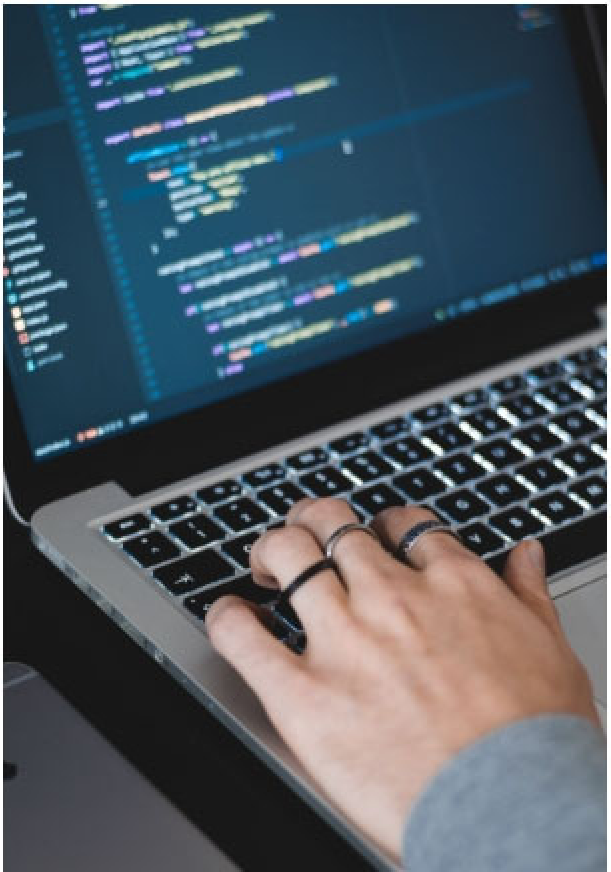

As such, HE based systems have no trust in the system's software or hardware, and therefore, no amount of hacking can penetrate the always-encrypted secret data. If attackers penetrate a system executing HE programs, all there is to steal is encrypted data! This unique property has led to significant interest in HE technologies and research to address the **primary drawbacks of HE:**

- HE computation is very slow compared to native computation, with HE operations typically requiring on the order of 10,000x - 100,000x more computation. This has led to much work on accelerating HE with more capable hardware.

- HE systems only work with pure mathematics, and as of this writing, the primary operations supported are addition, subtraction and multiplication. If one's algorithm can be fully expressed using only these operators, then today's HE will suffice. However, if one's algorithm requires relational operations (e.g., <, <=, ==) or decision-making (e.g., if-then-else), these algorithms are notably challenging to express in HE. For example, to perform a "<" operation, most HE frameworks will have to utilize a linear Taylor-series expansion of the non-linear less-than function!

- It is also notable that fully capable HE-based systems are based on ciphers that are at most a decade old, and they have not yet received the cryptanalysis and side-channel attack analysis that widely used symmetric and asymmetric ciphers have received. While it is not a focus of research today, we expect that soon these algorithms will be exposed to significant attack analysis, which may find additional vulnerabilities.

Despite these drawbacks, **the promise of zero software trust is incredibly exciting.** For example, if one were to build a voting machine with HE, it would preserve the confidentiality of the current vote totals despite being penetrated by attackers. Quite simply, no other data security technology would even dare to boast such a property, since all other voting machines would have software protecting the secret votes, and software can always be hacked!

ØZONE SECURITY IS A REMEDY

At Agita Labs, our goal is to advance the state-of-the-art in data security technology, with data security as strong as HE, while retaining the high performance and programmability that current HE technologies lack. To accomplish this goal, we are deploying the ØZone technology, which is an advanced hardware TEE, that eliminates the problems that plague current TEE technology.

ØZone technology addresses the deficiencies of existing TEEs: namely,

- ØZone's implementation contains no trusted software, thus, software hacking cannot penetrate ØZone data security, and

- ØZone technology is built to eliminate all side-channel attacks.

We achieve this goal with a novel technology called sequestered encryption (SE). This technology possesses powerful properties that give SE data security the strength of homomorphic encryption, but at the same time significantly better performance and programmability. It achieves these goals by using well-known cryptographic protocols (e.g., RSA and Simon) to gain a zero-trust stance with respect to software, by placing the data security trust solely in a hardware-based trusted execution environment.

Our initial deployment of the technology will provide support to protect sensitive data, e.g., medical and genomic data. In addition, the technology will protect program pointer values, which significantly hardens a program against penetration attacks. An earlier version of the ØZone technology (from The University of Michigan Morpheus project [5]) took part in a recent DARPA-sponsored red-teaming effort (FETT). A RISC-V system with sequestered pointers was subjected to attacks from more than 500 attackers over a three-month period. The system was never penetrated by any of the attacks, representing a new milestone in durable security defenses. The ØZone Privacy platform builds on this success.

ØZONE TECHNOLOGY
How Secure Is It?

ØZone security is designed to be a as strong as the cryptography that is it implemented upon, which is currently the asymmetric cipher RSA and the symmetric cipher Simon. This strong security property is derived from ØZone's approach to data security: all ØZone operations produce high-entropy encrypted results, without any measurable side channels. As a result, hacking into an ØZone program will only reveal high entropy ciphertext to an attacker!

To demonstrate the strength of ØZone security, we have validated its ciphertext indistinguishability properties. Ciphertext indistinguishability is a property where samples of ØZone ciphertext cannot be distinguished from pure random data. This property represents one of the strongest security stances from the field of cryptography, dubbed IND-CPA.

References

1. A. Ng. 2018. How the Equifax hack happened, and what still needs to be done. CNET (Sep 2018). https://www.cnet.com/news/equifaxs-hack-one-year-later-a-look-back-at-how-it-happened-and-whats-changed.
2. A. Schiffer. 2017. How a fish tank helped hack a casino. Washington Post (Jul 2017). https://www.wapo.com/news/innovations/wp/2017/07/21/how-a-fish-tank helped-hack-a-casino
3. IBM Corporation, Cost of a Data Breach Report 2020. (Jun 2020). https://www.ibm.com/security/data-breach
4. Spectre and Meltdown, (Jan 2018), https://meltdownattack.com M. Gallagher, L. Biernacki, S. Chen, Z. B. Aweke, S. F. Yitbarek, M. T. Aga, A. Harris, Z. Xu, B. Kasikci, V. Bertacco, S. Malik, M. Tiwari, and T. Austin, "Morpheus: A Vulnerability-Tolerant Secure Architecture Based on Ensembles of Moving Target Defenses with Churn", ASPLOS, 2019