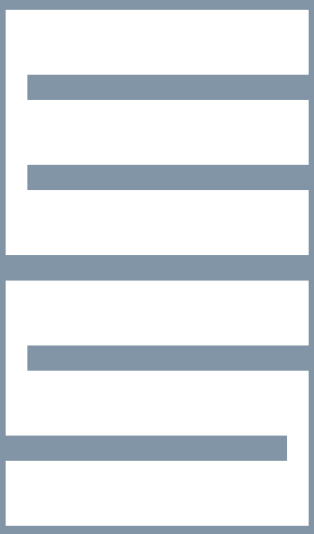

EDGELESS
SYSTEMS

WHITEPAPER

CONFIDENTIAL COMPUTING 101

https://**edgeless.systems**

May 2021, v1.0

Imprint

Edgeless Systems GmbH
Castroper Str. 12
44791 Bochum, Germany

1 CONTENTS

2 INTRODUCTION

It is common knowledge that data is getting increasingly important for our modern economy.

Organizations are progressively embracing data-heavy topics such as artificial intelligence while cloud computing has hence become part of the day-to-day business. Information is constantly generated, consumed, shared, and stored including very sensitive details such as credit card data, health-related data, personal data, firewall configurations, and geolocation data.

When it comes to data protection, one must keep in mind its three existing states, which are data at rest meaning data that is in storage e.g., in a database, data in transit traversing the network and data in use which is being processed at that time.

While techniques to protect data at rest and data in transit are well known and frequently deployed, the protection of data in use as well as its verifiability while being processed are not yet established.

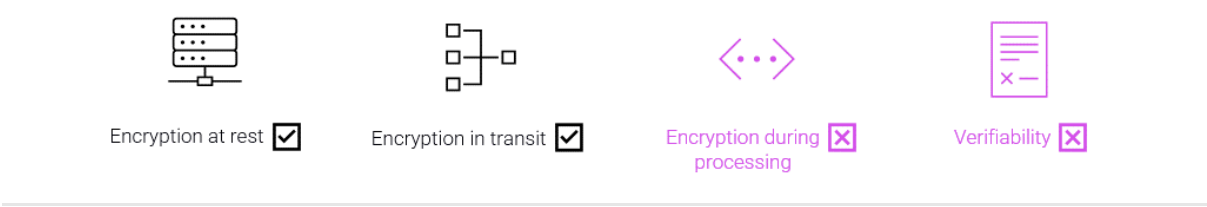

Figure 1: Previous data-security solutions only provide encryption at rest and encryption in transit.

Confidential computing is a new paradigm that enables encryption of data at runtime through the use of hardware based Trusted Execution Environments (TEE). It can be mathematically ensured that no unauthorized access or modification by a potential malicious actor or even the system administrator is possible.

3 WHY DO WE NEED CONFIDENTIAL COMPUTING?

Security and privacy expectations are increasing among regulators, companies and consumers. Straight in mind comes in this context the GDPR (General Data Protection Regulation, EU 2016) and the CCPA (California Consumer Privacy Act of 2018) as well as recent large-scale data breaches.

Since current technology is not properly addressing these issues, many companies are missing out on creating value from their data. For example, today companies rarely share sensitive and valuable data, because once shared, there is no going back. As a result, many companies are unable or unwilling to use the cloud for certain types of data processing, thus using their IT-infrastructure inefficiently.

Furthermore, this security risk has gotten the attention of attackers who now have been targeting data in use on several occasions, including high-profile memory scraping, such as the Target breach in 2013, and CPU-side-channel attacks. In addition, the triton attack in 2017 and the Ukraine power grid attack in 2015 are only two of several high-profile attacks on data in use involving malware injection.

The protection of data and applications during execution is increasingly important and must be part of the overall defense strategy.

This is why we need Confidential computing.

4 WHAT IS CONFIDENTIAL COMPUTING?

As stated above Confidential computing involves protecting data and code within hardware-based secure enclaves, also referred to as Trusted Execution Environment (TEE). The most prominent enclave implementation to date is Intel SGX. In a nutshell, enclaves enable isolated and auditable processing of data on untrusted computing systems – this can be the user's own computer or a machine in the cloud. With Intel SGX, the contents of an enclave even remain encrypted in memory at runtime. The following chapter will dive deeper into the topic.

Confidential computing not only takes general security to a new level, but also enables new types of data-driven applications. The verification aspect of Confidential computing is key here: remote parties can verify exactly how data is processed, who provides the input, and who gets access to the results.

This enables, for example, the secure and rule-based exchange of data between potentially distrustful parties while enabling high performance and confidentiality. Similarly, companies can process their customers' sensitive data while proving that no one, including their own analysts and administrators, can see the raw data at any time.

Confidential computing is therefore believed to be a game changer. Forbes Magazine has named it one of the Top 10 Digital Transformation Trends for 2021 and it is listed in Gartner's Hype Cycle for Compute Infrastructure as well as its Hype Cycle for Cloud Security of 2020.

Many American and Chinese big-tech companies and a few start-ups, including Edgeless Systems, have joined forces in the Confidential computing Consortium to drive adoption and define standards.

However, the whole topic is still quite nascent. There clearly is a lack of software in general and in particular of devops/dev tools. For example, tooling for the orchestration of workloads in clusters of secure enclaves, e.g., on Kubernetes in the cloud is a key piece that has been missing so far. The goal is that the development and delivery of confidential computing apps will be as smooth as that of normal cloud-native apps.

As stated above, the main element of Confidential computing is the so-called Trusted Execution Environment (TEE). It is commonly known as an isolated processing environment in which applications can be securely executed irrespective of the rest of the system. They provide a level of assurance of data and code integrity, meaning that the data is kept confidential and cannot be replaced or modified by unauthorized entities.

In these isolated enclave data and code are always encrypted in memory, even during processing. Not even administrators or the operating system can access an enclave. Likewise, physical attackers only see ciphertext. A good picture for an enclave is therefore a super secure black-box.

The enclave has a number of benefits. First of all, it virtually eliminates the risk of data breach, since no one outside the enclave can see or access the data inside. Furthermore, it may lead to securing all data and applications by default, counteracting the current practice of securing only the most sensitive and valuable data, e.g., long-term signing keys, in specialized hardware. In addition, enclaves offer the same protection wherever they run, including the public cloud or untrusted locations in general. Data can therefore be protected everywhere it is used, be it at rest, in transit or even in use. For companies, this benefit could be the most significant argument to fully embrace cloud computing as a part of their IT infrastructure.

Another huge benefit is that attacks from insiders can be prevented, because physical access still does not enable access to the enclave.

With an origin in the 2010s, Confidential computing as a discipline is still fairly young. Thus, it is unsurprising that there are only few established standards yet. In addition, the current open-source landscape can be challenging to navigate because the available technology has yet to distinguish itself. While this may discourage the less technically inclined, the current state of Confidential computing presents itself as a great opportunity for experienced developers.

Also, one of the main benefits is the assurance of customer data privacy through the use of enclaves. Companies can verify, that their employees or any other party will not have access to the data and can use that as a significant competitive advantage.

Companies that currently consider the exchange of data with partners, e.g. in order to jointly train AI models, will find the security features of Confidential computing to be especially beneficial. By utilizing enclaves, they can grant TEE software access to their data without anyone else being able to intervene. As enclave memory can only be accessed by enclave code, the software can be seen as having means to protect itself. Additionally, with Remote Attestation, any of the involved parties can be assured of the integrity of data and code.

Lastly, enclaves are eliminating costly security software layers, by offering a security solution on the lowest layer possible. Combining physical, network, system, endpoint, and application security protections and installing and integrating them into an enterprise IT organization is no longer necessary. This not only has beneficial financial aspects, but additionally very much reduces the attack surface as shown in the illustration below.

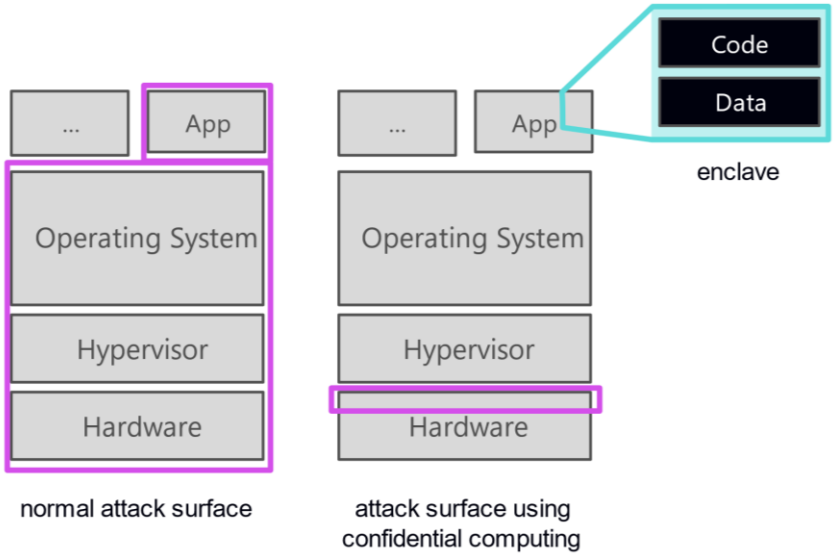

Figure 2: Comparison of attack surfaces.

4.1 WHICH HARDWARE IS READY FOR CONFIDENTIAL COMPUTING?

Confidential computing hardware, i.e., hardware that is able to provide enclaves, is surprisingly prevalent. Many recent client and server processors from Intel are equipped with a feature called SGX – which is short for Software Guard Extensions. Without doubt, Intel SGX is the most widely spread Confidential computing hardware provider. It is available in several cloud offerings including Microsoft Azure and Alibaba Cloud.

Conceptually, Intel SGX enclaves can be used to create confidential functions or confidential apps. Confidential functions denotes concept of splitting an application, e.g., a database, into a trusted and an untrusted part. The trusted part runs inside the enclave and has access to secrets, while the untrusted part remains outside. Since the task of splitting an existing application can be cumbersome or even close to impossible, the confidential apps concept is gaining popularity. In this approach, entire applications are run inside enclaves. The EGo framework makes it for example easy to lift and shift applications written in the Go programming language to enclaves.

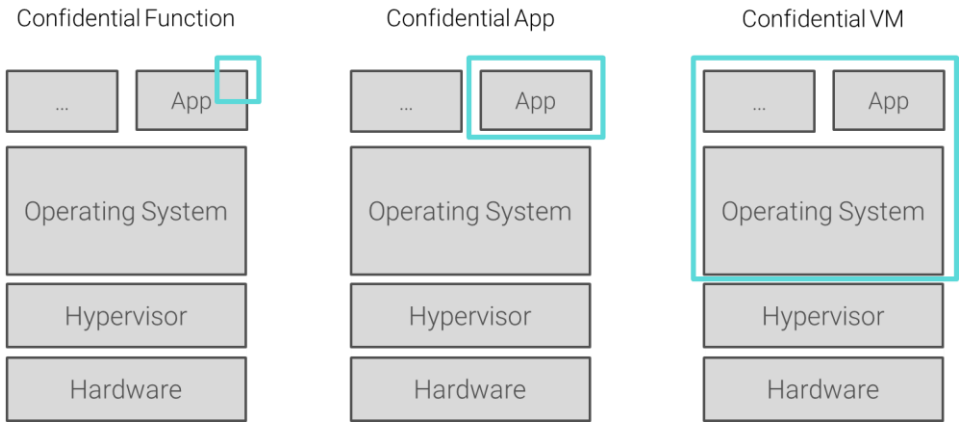

Figure 3: The enclave concepts of confidential functions, confidential apps, and confidential VMs

Besides Intel SGX, three other Confidential computing hardware providers exist or have been announced: AMD SEV, Intel TDX and Arm Realms. These three have in common that they follow the concept of confidential VMs: they support running entire virtual machines inside enclaves. This approach makes it even easier to lift and shift existing software. However, it also comes with the concern of an enlarged attack surface, since a VM typically comprises billions of lines of code, which all become part of the so-called trusted computing base (TCB).

AMD SEV hardware is already available in several clouds including Microsoft Azure and Google Cloud. Intel TDX and Arm Realms have been announced but are not yet commercially available.

Figure 4 depicts the security and usability trade-offs between confidential functions, confidential apps, and confidential VMs.

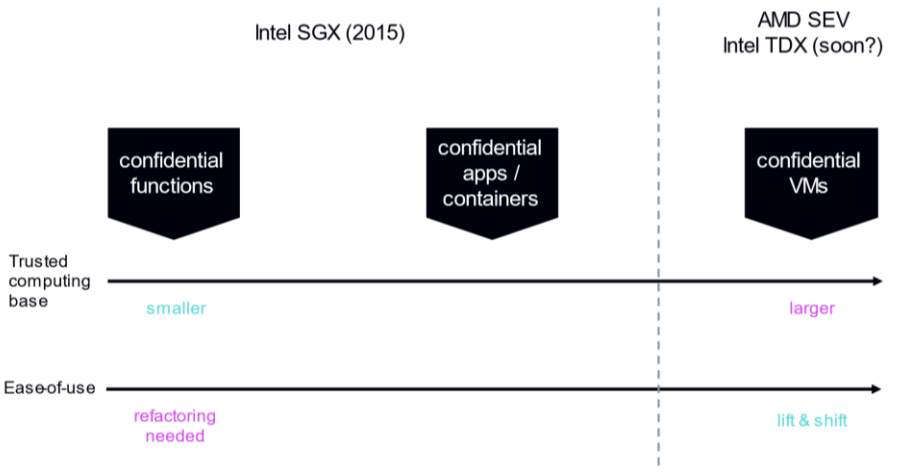

Figure 4: Trade-offs between confidential functions, confidential apps, and confidential VMs.

5 WHAT ARE INDUSTRY USE CASES FOR CONFIDENTIAL COMPUTING

There are a number of different industries that will benefit from Confidential computing. To give some examples several use cases will be described.

Confidential computing in Mobility

With Confidential computing, sensor data from networked vehicles can be aggregated and processed in an end-to-end encrypted and end-to-end verifiable way. If implemented correctly, even the vehicle manufacturer and the application operator only get access to the aggregated and filtered output data. It can be mathematically ensured that no relevant conclusions can be drawn from the output data about individual drivers. Using the verifiability properties of Confidential computing, this can even be proven to customers, partners, and legislators. Thus, Confidential computing-based applications can dramatically increase customer acceptance for the use of their data and help with compliance.

We at Edgeless Systems recently engaged in a related project with the German automotive and tech giant Bosch. We recommend you take a look at Bosch's insightful blog post on the project.

Confidential computing in mechanical engineering and manufacturing

"Industry 4.0" is a much-hyped term in mechanical engineering and manufacturing – in particular in Germany. In a nutshell, Industry 4.0 stands for the approach to increase productivity through massive use of sensors and corresponding data analytics. However, it is not uncommon for industrial data to be of a sensitive nature, as it often contains business secrets and special know-how. Companies are therefore often unwilling to share this data or process it in the cloud. Confidential computing can address these concerns comprehensively. The principle of "sharing data without sharing it", which is enabled by Confidential computing, will create lots of value in the future regarding predictive maintenance, digital twins and other data-driven industry applications. For instance, BCG estimates in a recent blog post that $100 billion of value could be unlocked through the exchange of data in manufacturing.

Confidential computing in healthcare and medical research

When it comes to patient data in healthcare, one is clearly dealing with highly sensitive and regulated data. In this context, Confidential computing can enable secure multi-party training of AI for different purposes. For example, multiple hospitals can combine their data to train AI for detecting diseases, say, given pictures from CT scans. The patients' data remains confidential during each step of the process. This way, the patients' privacy is protected and hospitals or other data owners remain in control of their valuable data.

One of the biggest healthcare applications based on Confidential computing, which is currently in production, will be the "E-Rezept" (electronic prescription). Starting in 2021, the E-Rezept infrastructure will handle drug prescriptions within the German national healthcare system. With the help of Confidential computing, patient data will be strongly protected throughout the lifecycle of a prescription – from the doctor's practice to the pharmacist's counter.

Confidential computing in Finance

Another example where the secure combination of data between two parties can unlock substantial value is finance. Through Confidential computing, a retailer and a credit card company can cross-check their customer and transaction data for potential fraud while neither of them gets access to the original data. The privacy of their customers sensitive data is ensured by Confidential computing alongside the whole process.

6 THE OPEN-SOURCE LANDSCAPE FOR CONFIDENTIAL COMPUTING IN 2021

Confidential computing and open-source have a special relationship. In the following the reasons for that as well as an outline of the current available open-source tools will be given.

As discussed earlier, confidential computing is based on three key concepts: isolation, runtime encryption, and verifiability. It might be counterintuitive, but verifiability is at least as important as the other two. In fact, without verifiability, the whole concept of secure enclaves doesn't make much sense.

If one can't verify that the workloads are actually running in secure enclaves, any malicious actor (e.g., a hacker in the cloud provider's systems) could pretend that they are and access the data once it is being sent. They would just simulate to be an Intel SGX or AMD SEV processor and run everything in plain.

In confidential computing, one can verify the following: a certain software running in a genuine secure enclave has produced some output. Such a statement is only insightful, if it is known what that piece of software precisely does. And this is where open source comes into play: if the software running in a secure enclave is open source, then it is much easier to establish trust in it.

Figure 3 sketches the relations between relevant confidential computing open-source projects on GitHub. It is intended to show the whole picture as of today. It shows the relationships of projects and names their main corporate sponsors.

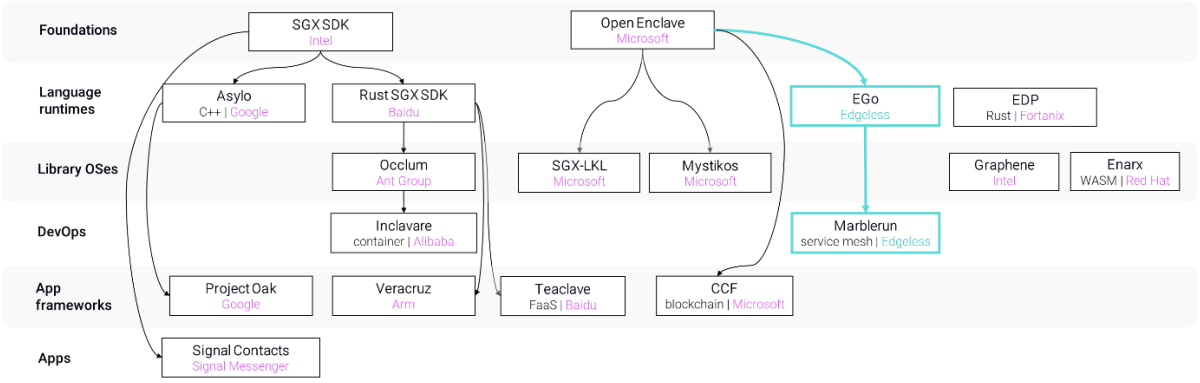

Figure 5: The open-source landscape in confidential computing.

On the lowest layer ("foundations"), one finds the two well-known projects SGX SDK and Open Enclave. They act low-level and for example handle in-enclave exceptions and make enclaves programmable in C and subsets of C++. Except for Graphene and Enarx, all other projects are directly or indirectly based on one of the two. They can roughly be categorized into language runtimes, library OSes, DevOps tools, app frameworks and actual apps.

Here, Red Hat's Enarx is again an outlier as it doesn't really fit into this taxonomy. The circumstance that it runs WebAssembly (WASM) and that it has an enclave-specific implementation of the WASI makes it somewhat a hybrid between a language runtime and a library OS. It also has features that are normally found in app frameworks.

What's interesting is that the landscape actually grew from the far ends towards the middle. Applications where build as soon as basic tooling was available. Signal was built on SGX SDK directly, CCF is based on Open Enclave and Project Oak evolved from the groundwork of Asylo. Meanwhile, tooling improved and support for more modern/high-level programming languages like Rust or Go was added by projects like the Rust SGX SDK or Edgeless RT.

In parallel, the breed of library OSes appeared, which promise to run most code regardless of programming language at the cost of a more complicated trusted computing base (TCB). In essence, library OSes try to mimic a Linux-like environment inside enclaves. The SGX-LKL project even goes as far as actually running a full Linux kernel towards this end.

Finally, now the first DevOps tools are starting to appear, namely Inclavare and Marblerun. Both were released fairly recent. It can be seen as a great sign of maturity for confidential computing that more and more blank spots on the open-source map are closing. With Marblerun, for instance, it is now possible to have secure and scalable cloud-native confidential-computing apps running on vanilla Kubernetes.

What is striking is that there are only very few apps built on the latest and greatest tools and frameworks. The prediction for 2021 is that the now available tooling will set loose a wave of innovation in apps and that we'll continue to see substantial improvements on the tooling side.

7 CONCLUSION

Many experts agree that Confidential computing has the potential to be a gamechanger in cloud security. For example, Gartner's 2020 Hype Cycle for "Cloud Security" lists the topic and Forbes Magazine ranked Confidential computing as a top 10 trend for 2021.

In addition, numerous large American and Chinese tech companies, such as Microsoft, Google or Alibaba, and some start-ups, such as Edgeless Systems from Bochum, have joined forces in the Confidential computing Consortium to drive the adoption of confidential computing and define standards.

The landscape of tools is rapidly evolving and provides business and end-users with the opportunity to protect data when it is most vulnerable: while being processed.

The goal of ensuring the confidentiality of the sensitive, business critical information and workloads is being tackled with various approaches and gets projected on numerous use case scenarios.

However, it is also clear that the evolution of the entire area is still at its very beginning. For example, there is still a distinct lack of software - especially in the devops/dev tools in this area.

But the hardware foundations for confidential computing have already been laid by providers such as Intel. In addition, all the major cloud providers are increasingly jumping on the bandwagon and offering TEEs in their data centers. However, for the possibilities of confidential computing to be used sustainably, it is crucial that a software ecosystem is formed in this area. In the future, we will therefore see many more projects that enable developer6s to deploy Confidential computing applications.

8 SOURCES

Bosch, "Trustworthy computing – data sovereignty while connected," [Online]. Available: https://www.bosch.com/research/know-how/success-stories/trustworthy-computing-data-sovereignty-while-connected/. [Accessed 04 05 2021].

Boston Consulting Group, "How Manufacturers Can Unlock Value from Data Sharing," 21 01 2020. [Online]. Available: https://www.bcg.com/publications/2020/manufacturers-unlock-value-from-data-sharing/. [Accessed 04 05 2021].

Forbes, "Top 10 Digital Transformation Trends For 2021," 21 09 2020. [Online]. Available: https://www.forbes.com/sites/danielnewman/2020/09/21/top-10-digital-transformation-trends-for-2021/?sh=6313f4e4c6f4/. [Accessed 04 05 2021].

Gartner, "Top Actions From Gartner Hype Cycle for Cloud Security, 2020," 27 08 2020. [Online]. Available: https://www.gartner.com/smarterwithgartner/top-actions-from-gartner-hype-cycle-for-cloud-security-2020/. [Accessed 04 05 2021].

Feel free to contact us directly or follow our socials to keep up to date with Confidential Computing.

https://www.edgeless.systems/

contact@edgeless.systems

https://github.com/edgelesssys

https://twitter.com/edgelesssystems

https://www.linkedin.com/company/edgeless-systems

https://discord.com/invite/rH8QTH56JN