EuroCloud
DEUTSCHLAND | eco

CONFIDENTIAL COMPUTING:

Secure and Sovereign in the Cloud

Content

Preliminary Note

Confidential Computing is a relatively new technological approach. It serves to protect sensitive data, even during processing in a cloud or hosting environment. The current text focused on this concept was created by a working group under the umbrella of EuroCloud Deutschland_eco e. V. A list of the authors can be found at the end of the document.

Cologne, January 2025

1. Introduction

Cloud Computing is now firmly established within companies. The majority of organisations today utilise services from one or more cloud providers, whether international hyperscalers or regionally operating providers. Companies aim to benefit from the advantages offered by the delivery model in terms of agility, availability, flexibility and scalability. However, the cloud has created new challenges for companies' security and sovereignty, given that data is processed outside the local IT and thus beyond the perimeter. Cloud Computing concurrently means a loss of control. After all, companies' workloads are carried out on infrastructure that is managed and controlled by a third party (internal or external) and to which the companies themselves have no direct access.

In response to these challenges, the IT industry has developed new security concepts. For example, under the term "Zero Trust", specific mechanisms have been developed to regulate access to applications and data based on the principle of minimal rights. Another concept that is suitable for increasing security in cloud usage is "Confidential Computing". The main idea is to establish confidential execution environments for workloads within the IT infrastructure that cannot be accessed by providers. Isolation and encryption protect data of users from unauthorised access, as well as from cyber threats.

Confidential Computing aims to provide organisations with a sense of trust that their sensitive data is secure, even when it is being processed in the cloud. This white paper provides the reader with an overview of Confidential Computing. From this perspective, the essential technological features are presented, and the basic functioning of the concept is explained. The white paper also highlights the advantages that the concept offers for processing sensitive data. This is particularly beneficial for companies in regulated industries that have high IT security requirements. Confidential Computing can assist them to comply with the applicable regulations. The risks and weaknesses of the concept should not be ignored. Finally, some use cases are outlined.

2. What is Confidential Computing?

Confidential Computing refers to a technology designed to ensure that data is processed in a protected environment. This is primarily achieved by means of encryption. When cloud platforms are used, data is typically encrypted when it is stored - i.e. at rest - and when it is transmitted between users' clients and the cloud providers' servers. However, data is usually decrypted during processing, i.e. when it is being handled by applications or services in the cloud. The purpose of Confidential Computing is to protect data during this critical stage, where it is vulnerable to attacks or errors in system administration. The concept therefore aims to ensure that data remains in a confidential environment, even while applications are running.[1]

2.1. Definition of Terms

In cloud and hosting environments, companies face the challenge of protecting sensitive data during storage, transmission and processing. This means that three states of data play a role: data at rest, data in transit (also referred to as "data in motion") and data in use. While encryption of data in the first two states is already widespread and implemented through technologies such as hard disk or transport encryption, the protection of sensitive data during processing - in other words, data in use - has not yet been widely adopted. This is precisely what Confidential Computing aims to address. If data is encrypted in all three states, this is also referred to as 3D encryption.

In a general definition, Confidential Computing can be described as a concept that is suitable for ensuring the trust of data during execution in untrusted environments. In this context, an untrusted environment refers to any IT environment to which third parties have access for administration and control purposes. Essentially, Confidential Computing enables the secure processing of applications and thus the secure processing of data. Critically, this approach ensures that the environment's operator – such as a hosting or cloud provider - cannot access the applications or data.

Confidential Computing includes 3D encryption technologies along with the necessary accompanying processes and tools that enable the processing of sensitive applications within a protected system environment. In general, this is referred to as a Trusted Execution Environment (TEE). In Confidential Computing, such trusted execution environments are implemented as so-called enclaves. Key technological requirements for this include logical isolation and encryption at the Central Processing Unit (CPU) level.

1 See also: https://next.enclaive.cloud/s/MPpzyDJbfkBBsFc

2.2. Excursus: Homomorphic Cryptography for Encrypting Data in Use

Confidential Computing is a relatively new concept. Before it emerged in recent years and the first marketable solutions were developed, there were already theoretical approaches to encrypting data during processing. This includes homomorphic cryptography. However, this approach has its limitations.

Homomorphic cryptography makes it possible to perform calculations on encrypted data without decrypting it. If data is stored and processed by a third-party provider in the context of Cloud Computing models, homomorphic cryptography is theoretically suitable as a solution that ensures data privacy. This is because the data is only visible to the owner in plain text. However, this technology incurs a significant computing overhead: it requires 100 to 1,000 times more computing power than is needed to process unencrypted workloads. Therefore, this approach is not (yet) commercially viable.

2.3. Key Features of Confidential Computing

Confidential Computing is fundamentally based on providing confidential execution environments in the form of enclaves. These are enabled by hardware-based encryption, which chip manufacturers such as AMD, IBM and Intel have been supporting in their processors for several years. This allows protected areas within a CPU to be isolated for the secure execution of applications.

Furthermore, Confidential Computing offers further security functions. An essential feature here is the ability to verify the integrity and authenticity of data processed within an enclave. This process, known as attestation, uses cryptographic methods to ensure that a workload has not been manipulated or compromised. This strengthens trust in the security of an execution environment.

2.3.1. Encryption as a CPU Feature

The manufacturers Intel and AMD have implemented hardware-based encryption in their processors since 2015 and 2017, respectively. The purpose is to protect workloads in virtual environments from external threats, such as the reading of data from memory via the hypervisor. This complements, but does not replace, important established approaches to protecting availability, confidentiality, resilience and integrity.

The latest chipsets from the vendors are delivered as standard features with technologies such as Intel Software Guard Extensions (SGX), Intel Trust Domain Extensions (TDX) and AMD Secure Encrypted Virtualisation (SEV).

In late 2023, the manufacturer NVIDIA released its H100 GPU series, the first system with a technology that enables the isolated execution of workloads. For ARM processors, it is expected that chips with comparable security features will hit the market at the beginning of 2025, following an announcement by the semiconductor specialist. IBM's System/390 also supports Confidential Computing.

As such, Confidential Computing features are now available for almost all processor architectures that are relevant on the market.

In concrete terms, the technological support for Confidential Computing at the CPU level (as of autumn 2024) is illustrated in the following table:

Technology (linked)	CPU	Codename
AMD Secure Encrypted Virtualization (SEV)	Untervorbehalt: EPYC 2 (SEV ES) EPYC 3 (SEV SNP) or later	Rome Milan
Intel Trusted Domain Extension (TDX)	4th Gen Xeon or related	Sapphire Rapids
ARM Confidential Compute Architecture (CCA)	ARM Cortex A9	Falcon
NVIDIA Confidential Computing	NVIDIA H100 GPU	Hopper
IBM Secure Execution for Linux (SEL)	IBM z15 or later IBM LinuxONE III	

2.3.2. Enclaves

An enclave, as previously outlined, is an isolated environment within a processor. Operations can be carried out securely in this environment, ensuring that the data being processed is protected. An essential feature here is isolation: enclaves are strictly separated from the other parts of the system, which in turn can only communicate with them via specific secure interfaces. Additionally, strong cryptographic technologies are used to encrypt the data within the enclave. This combination of isolation and encryption ensures that data in such an execution environment is protected against unauthorised access.

In this way, enclaves serve the purpose of protecting sensitive data in cloud and hosting scenarios from potential threats such as theft or compromise by malware.[2] The data remains secure even if an attacker penetrates the overall system and the neighbouring system areas are compromised.

Attacks on virtualised systems typically exploit vulnerabilities in the hypervisor or host operating system. This gives attackers the opportunity to gain access to areas of the random-access memory (RAM). In theory, the logical separation of the RAM regions assigned to different virtualised instances can also be breached due to a CPU vulnerability. Examples include the "Meltdown" and "Spectre" vulnerabilities that became known in 2017/2018. These allowed potential attackers to steal sensitive data from the random-access memory.

Enclaves largely counteract such scenarios by encrypting all virtualised storage areas with their own key material. If an attack on virtualised instances occurs via the CPU or the management layer (hypervisor, host operating system), the attacker would at best gain insight into their encrypted data. However, the attacker cannot do anything with this data as long as the key for the respective storage area remains unknown.[3]

2.3.3. Workload Attestation

Attestation, as previously mentioned, is another essential feature of Confidential Computing. It serves to securely verify the integrity and authenticity of workloads running in a Trusted Execution Environment (TEE) in an audit-proof manner. Attestation can be used to verify that the workload in question has not been tampered with, impaired or compromised. This allows the user to trust that an enclave used for processing their data provides the expected protection and does not exhibit any known vulnerabilities. Attestation is a complex procedure. Simply put, it aims to make the flawless security status of an execution environment verifiable.

To determine whether an enclave actually provides the expected high security standard, cryptographic methods are utilised. During workload attestation, the enclave measures the initial configuration of the environment and the application code at the beginning of a computing operation, generating a unique hash from these measurements. This measured value is recorded, along with other relevant information about the enclave's identity and security properties, in a report called the "Attestation Report".

2 However, it must be ensured that the data to be processed has not already been compromised. In theory, malware can run within an enclave if it is overlooked and included during the attestation. This makes a comprehensive attestation all the more important..

3 Research experiments have shown that it is theoretically possible to extract this key with direct access to the system (side-channel attacks)

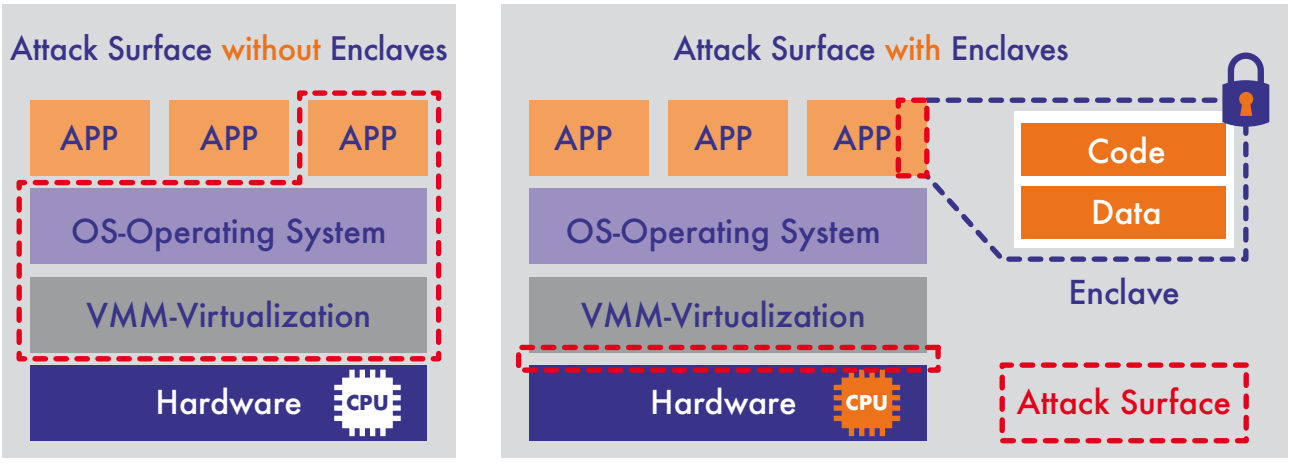

Confidential Computing enables the encrypted processing of application code and data within an enclave. This significantly reduces the attack surface. All of the necessary security functions are implemented in the CPU.

For enclaves in the cloud, the verification of their identity and security status is usually carried out remotely. After all, cloud users do not have direct access to the physical infrastructure on which the enclaves are running. In remote attestation, another entity comes into play: the Attestation Service. As a trusted third party, this service reviews the Attestation Report. It analyses the report, validates the measurement, assesses the workload's integrity and, if successful, confirms the secure state of the enclave with a certificate.

The verification by the Attestation Service is typically based on standard guidelines. Users also have the option to verify the attestation certificate by means of their own user-defined policies. In this way, they can ensure that the enclave meets their individual security requirements. The policies can specify defined measurements, the use of certain software versions or compliance with specific security standards.

What is essential for remote attestation is ensuring that communication between the enclave, the certification service and the external user is consistently protected across the board to maintain the confidentiality and integrity of the transmitted data. Technologies such as encryption and secure channels are used for this purpose.

Certification services can be provided by different types of companies. On the one hand, cloud providers themselves offer such services. For hyperscalers, they already include them in their portfolios. On the other hand, providers of security software or IT service providers can be considered as operators of attestation services. Moreover, large companies with high security requirements can also establish and provide their own internal certification service. To consistently implement the principle, the operator of the environment being attested and the verifying entity should ideally not belong to the same organisation.

Fundamentally, workload attestation is a solid procedure that enables secure interaction between users and enclaves. By verifying the integrity and authenticity of data processed in confidential execution environments in an audit-proof manner, attestation strengthens trust in the use of enclaves and, ultimately, in the use of the cloud.

3. Advantages of Confidential Computing

The advantages of the concept of Confidential Computing stem from its functionality and essential characteristics. By using isolated enclaves whose contents are encrypted and whose security status is regularly checked, companies can effectively protect their sensitive data even when processing it in the cloud. On the one hand, this enhances their security against cyberattacks. On the other hand, the use of confidential execution environments helps them to fulfill legal requirements and internal guidelines.

Although data encryption requires computing power, the additional CPU load is relatively minor compared to the gains in security and compliance. With current enclave implementations, the overhead is usually only two to five per cent of the performance that the workloads would require without Confidential Computing.

It is evident that Confidential Computing is particularly beneficial for organisations that operate in a regulated sector (e.g. finance, healthcare or public) but still want to use the cloud. Despite strict regulatory requirements, they can process data on cloud platforms in a compliant manner.

However, the concept is suitable for companies across all industries. It not only offers advantages for organisations that use IT, but also for IT providers such as software manufacturers or Managed Service Providers (MSPs) who want to provide their customers with secure SaaS applications or secure infrastructure services. In practice, there are use cases in almost all industries and a wide range of application scenarios (examples in Chapter 7).

The advantages of Confidential Computing will be examined in greater detail in the following two sub-sections, divided into operational and security-specific advantages.

3.1. Operational Advantages

Confidential Computing is a technological approach that can help companies implement security and compliance management measures. Ultimately, key principles of the concept, such as the provision of confidential execution environments and end-to-end encryption (3D encryption), ensure that applications and data are protected from unauthorised access. Added to this are further security functions such as access controls and authentication mechanisms, logging of access and operations, and workload attestation, which enable consistent monitoring and documentation of security-related processes.

These features support companies in complying with regulations, protecting intellectual property and fulfilling audit and reporting requirements. Moreover, Confidential Computing can be particularly helpful when migrating workloads from on-premises to the cloud.

3.1.1. Compliance with Data Protection Regulations

Confidential Computing makes it easier for users to comply with legal data protection regulations. As enclaves are strictly separated from the underlying infrastructure, administrative access to their content can be excluded. Companies can therefore ensure that they are compliant with legal acts such as the General Data Protection Regulation (GDPR) when they process personal data in the cloud. This is because the cloud provider has no access to the encrypted data in the enclave, even if legal acts such as the US CLOUD Act were to obligate them to do so (see also Chapter 6).[4]

This advantage applies not only to Cloud Computing, but also to scenarios where companies collaborate with third parties and give them remote access to their systems. Confidential Computing enables them to comprehensively control their partners' access. For instance, the concept supports Multi-Party Computation (MPC), a procedure where multiple parties can jointly analyse data sets, without disclosing sensitive information.

3.1.2. Reduced Reporting Effort

As applications and data are operationally separated from the infrastructure and the data is encrypted during processing, it is not only easier to meet compliance requirements. Rather, Confidential Computing also reduces the effort for regulatory reporting. Because, seen in this light, there is nothing to report. Technically, it is not possible to modify the data within an enclave. Personal data in the sense of the GDPR can neither be shared nor read, since they are encrypted artefacts. Moreover, because administrative access to the content is inherently excluded, there is no reporting requirement in the area of privileged users.

4 See also: https://next.enclaive.cloud/s/Y65T2tZEo8eoWKr)

3.1.3. Protection of Intellectual Property

Through the use of Confidential Computing, companies can allow both customers and partners to access their applications without risking exposure of their source code. In addition, by processing data in a secure environment, akin to a vault, companies can quickly and easily test sensitive digital innovations in the cloud. In doing so, they neither risk data loss nor a compliance violation.

3.1.4. Support for Lift-and-Shift Migrations

When applications are transferred from on-premises to the cloud without modifying their architecture or code, this is referred to as "Lift and Shift". Since cloud infrastructures differ from local infrastructures, this type of migration can cause issues such as performance degradation or outages. Confidential Computing circumvents this issue. This is because an enclave provides the user with a dedicated environment in the cloud that can be configured with the same attributes as an on-premises environment. This allows applications to be migrated without major adjustments.

3.2 Protection Against Attack Vectors

In addition to operational benefits, the use of Confidential Computing also strengthens resilience to cyber threats. This is because the concept can help security teams to eliminate a range of attack vectors, or at least mitigate them in such a way that it is significantly more difficult for attackers to achieve their goal. In the sections below, we will outline some examples of these vectors.

3.2.1. Attacks on Virtual Machines (VM Attacks)

In the cloud, attackers can theoretically access data in Virtual Machines (VMs) by exploiting vulnerabilities in hypervisors or in other VMs. Within confidential execution environments, applications and data are isolated and are therefore protected against such attacks, even if the hypervisor or VMs on the same host are compromised.

3.2.2. Malware and Rootkits

The same principle also defends against malicious code of any kind. Enclaves ensure that the data they process remains secure, even if the hypervisor or operating system is compromised by malware.

3.2.3. Memory Attacks

Confidential Computing safeguards against memory attacks by ensuring workloads are processed in an isolated environment where data is only stored in encrypted form. This prevents attackers from extracting plaintext data from the memory area.

3.2.4. Cold Boot Attacks

Cold boot attacks also aim to extract data from the memory. To do this, attackers quickly reboot a system and capture the RAM contents. Confidential Computing thwarts this attack pattern by encrypting the data in the memory.

3.2.5. Insider Threats

Insider threats, posed by individuals from their own organisation or immediate environment who misuse their access data and internal knowledge for malicious purposes, represent a major security challenge. Confidential Computing counteracts such insider threats by eliminating administrative access to confidential execution environments. Even privileged users such as system administrators or service providers (e.g. cloud providers) do not have access to the data processed in these environments.

3.2.6. Man-in-the-Middle Attacks

A Man-in-the-Middle attack is an assault in which attackers secretly insert themselves into a communication between two parties and intercept or manipulate data. By running workloads within an enclave, the data is protected from threats during processing. By securely distributing cryptographic keys, Confidential Computing can also help to encrypt communications between users and the enclave end-to-end, thereby safeguarding data even during transmission.

4. Risks and Challenges

The already-mentioned security vulnerabilities, "Meltdown" and "Spectre", revealed potential attack vectors targeting the security mechanisms of common processors. If hackers exploit these vulnerabilities, they could theoretically retrieve or manipulate data from memory areas to which they should not have access.

Since then, CPU manufacturers have been consistently providing firmware patches to address known security vulnerabilities. In a worst-case scenario, however, the user may need to switch to a newer version of the processor if the problems cannot be resolved through a firmware update. Maintaining a consistent vulnerability and patch management strategy is critical when operating server systems. Nonetheless, it cannot be ruled out that hacker groups may discover such vulnerabilities before hardware manufacturers become aware of them.

4.1. Vendor Lock-in

Confidential Computing utilises micro-instructions within processors to encrypt data. This makes the hardware of chip manufacturers the foundational technology that users must trust. Leading semiconductor specialists such as Intel, AMD, ARM, IBM and NVIDIA currently offer server processors capable of running applications in trusted environments. The only exception is the RISC-V architecture, which is not yet designed for Confidential Computing. Projects in Europe and China are currently working on evolving processor technology towards Confidential Computing. All manufacturers have now focused exclusively on server systems after Intel initially attempted to bring the concept to desktop PCs.

Since the basic technology for Confidential Computing is imbedded in processors, users are dependent on how chip manufacturers continue to develop their products. Under certain circumstances, specific functions may be discontinued. Furthermore, the procedures used in CPUs cannot currently be neutrally audited by third parties. The question of digital sovereignty thus shifts further down the IT stack to the semiconductor level. While these dependencies may not necessarily be critical, companies should be aware of them.

As is often the case in the IT industry, the solutions offered by the various manufacturers are similar, differing only in implementation details. Currently, dependencies for users arise primarily because hyperscalers and regional cloud providers implement Confidential Computing in different ways. For instance, AWS has developed its own technology under the label "Nitro", which is based on proprietary hardware (plug-in cards). Companies can counter vendor lock-in by using middleware, known as Multi-Cloud Confidential Computing Brokers (MCCCBs), which neutralise these differences. However, this also means that the lock-in tendency is more likely to occur at the Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) level than at the CPU level.

4.2. Potential Vulnerabilities

In recent years, there have been repeated reports of attacks on Confidential Computing technologies. In 2023, researchers from the CISPA Helmholtz Centre for Information Security in Saarbrücken and the Graz University of Technology discovered the vulnerability "CacheWarpe" in AMD SEV. This allowed attackers to gain unrestricted access to enclaves via the cache. AMD closed the gap with a microcode update. Further insights were revealed in 2024 by the "Ahoi Attacks", in which researchers at ETH Zurich demonstrated possible attacks on SEV and TDX environments. These attacks involve using untrusted hypervisors to compromise trusted VMs. Security patches in the Linux kernel can prevent such attacks.

When Intel TDX was still a non-public prototype, Google conducted an in-depth security analysis of the technology. In the course of this analysis, several security vulnerabilities were discovered and subsequently fixed by Intel. The most significant finding was a bug in the Intel Authenticated Code Module (ACM), which is responsible for initialising TDX. This error could have allowed attackers to execute programs in privileged mode and compromise supposedly protected VMs.

To address vulnerabilities and further enhance the security of their enclaves, manufacturers have provided numerous microcode updates and new processor versions over the years. However, further attacks of the described type still seem realistic, especially since the architecture of VMs has been developed thus far on the basis of trusted hypervisors. However, this is not sufficient to ensure the security of the machines. Furthermore, the operating systems within trusted VMs need to be hardened and better protected.

Companies need to rely on experts to identify and assess vulnerabilities. Comprehensive analyses of which potential attacks are relevant for an application and which security measures are required are not provided by manufacturers. Intel has published a "Security Guidance" document for TDX, detailing information on attack vectors and corresponding countermeasures. Similar documentation is still lacking for other Confidential Computing technologies.

Another topic that still needs to be researched is the secure utilisation of trusted hardware from enclaves. It would, for example, be desirable if applications running on such hardware could also access network or graphics cards securely. While architectural concepts already exist for this, such as AMD SEV-TIO and Intel TDX Connect TEE-IO, implementations are yet to be realised. It is also still unclear how the extended environment can be verified by remote parties via attestation.

5. Legal Framework for Confidential Computing

The execution of applications and the processing of data in trusted execution environments requires more than technical integrity protection. Confidential Computing is also subject to legal, regulatory and contractual requirements. Certifications and attestations can guide users and promote trust in the technology.

5.1. CRA and CE Marking for IT Products

The CE mark from the CENELEC standardisation institute indicates that a product meets the fundamental requirements of European directives ("Conformité Européenne").

With the EU's Cyber Resilience Act (CRA), which has come into force in December 2024 and requires manufacturers of digital products to implement security by design, the CENELEC requirements will also increase. Three criticality levels will determine whether a self-declaration by the manufacturer is sufficient for the CE mark or whether an audit by an independent third party must be carried out beforehand. However, for Confidential Computing, a manufacturer's declaration alone will not suffice in any case.

5.2. EUCS and Sovereignty

The European Cybersecurity Certification Scheme for Cloud Services (EUCS) is a framework that regulates the certification of cloud services in terms of cybersecurity. The aim is to establish uniform and high security standards for cloud providers within the EU. The EUCS has not yet been conclusively negotiated and the deadline for this has not yet been set. Therefore, it is still unclear whether the final version will contain any sovereignty requirements and, if so, which ones. However, it seems to be emerging that some countries will have stricter national sovereignty regulations, even if the EUCS does not include any.

Sovereignty can be understood as control: physical control over the infrastructure (data centres, servers, data lines, mobile phone antennas, etc.), but also control over the data, its storage and processing location, as well as control over access permissions and the data lifecycle. With these dimensions of control, Confidential Computing can meet the increased need for protection.

With enclaves, users are provided with dedicated environments that are isolated from the infrastructure. Administrative access to these is excluded, so that even the infrastructure operator, such as a cloud provider, is prevented from unauthorised access into the data, as well as from unauthorised use and disclosure. In this way, Confidential Computing can help companies to achieve their goal of sovereignty, even if they have no control over the infrastructure. This applies regardless of the location of processing, as the data is protected in every dimension when implemented correctly.

5.3. Confidential Computing in the Implementation of Regulations

Confidential Computing is particularly relevant in the context of current data economy regulations. Ultimately, the concept is suitable for ensuring both the security and privacy of data, as well as its integrity. These regulations include, in particular, the General Data Protection Regulation (GDPR), the second edition of the Network and Information Systems Directive (NIS2) and the Digital Operational Resilience Act (DORA), a specific regulation for the financial sector.

The GDPR imposes strict requirements for the protection of personal data. Confidential Computing can help companies comply with the regulation, as the technology ensures the confidentiality and integrity of the data, such as when using the cloud. Since the data remains encrypted even during processing, the risk of data breaches is minimised, which is a central concern of the GDPR.

In the case of the NIS2 Directive, which aims to increase the security of Internet and IT systems, affected companies can also benefit from Confidential Computing. This is because the concept can help to increase the resilience of critical infrastructures against cyberattacks by providing an additional layer of protection during data processing.

Likewise, Confidential Computing can help the financial industry to meet the requirements of DORA. This regulation is intended to better prepare banks and insurance companies for cyberattacks and to increase the resilience of the financial system. Specifically, Article 6 of DORA requires "encryption and cryptographic controls" and "rules for the encryption of data at rest, in transit and, where appropriate, in use, the results of an approved data classification are to be taken into account [...] Where encryption of data in use is not possible, financial entities shall process data in use in a separated and protected environment, or take equivalent measures [...]".

For data classified as "relevant" under DORA, Confidential Computing and enclaves are therefore an essential prerequisite for allowing processing in the cloud. Data is considered relevant if it is essential for the operation of a financial institution and its loss or compromise could lead to significant disruption. Without encryption during use, such data would have to be processed in dedicated systems to meet the high protection requirements.

It is important to note that Confidential Computing can make a significant contribution to simplifying the implementation of technical measures required for compliance with the aforementioned regulations. The same principle applies to adherence to standards, whether these are general security frameworks like ISO 27001ff or sector-specific ones such as ISAX for the automotive sector.

6. Confidential Computing in Hyperscaler Use

When companies operate their IT or parts of it on the platform of a North American hyperscaler, Confidential Computing can serve two key purposes. Firstly, it provides users with the assurance that the hyperscaler cannot access their data in plain text. This also minimises the risk of potential data disclosure by the cloud provider based on regulations such as the CLOUD Act. What the hyperscaler cannot see, it cannot disclose.

Confidential Computing also protects cloud users' data from being accessed by other customers on the same platform. In the public cloud, many users share the same physical infrastructure ("shared infrastructure"), with resources being allocated flexibly and based on demand. This means that a physical piece of hardware can potentially be used simultaneously by several customers. Even if the accounts are logically separated from each other, there is still a residual risk. Confidential Computing ensures that a customer's data cannot be viewed, copied or otherwise processed by other customers at any time.

The hyperscalers are undoubtedly playing a pioneering role when it comes to embracing Confidential Computing. The concept helps them and their customers to meet security and data protection requirements in Europe. However, the hyperscalers' offerings are still under development. Currently, Confidential Computing is not yet available in all regions across Europe. For example, Microsoft primarily offers such options in Ireland and the Netherlands. As of 2024, the cloud provider is also offering enclaves in Germany, but initially only based on AMD SEV. Although AWS's own Nitro system is available throughout Europe, to date the provider only offers Confidential Computing based on AMD SEV in Ireland. Frankfurt is on the roadmap. At Google, confidential environments are available in Frankfurt, but only on the basis of AMD SEV, and not yet on the basis of Intel TDX.

The first regional cloud service providers have also taken up the topic and are beginning to differentiate themselves through specific PaaS offerings. Examples include Adacor Hosting, Open Telekom Cloud (OTC), OVHCloud, STACKIT (Schwarz Gruppe) and vshosting.

The following table provides an overview of the hyperscalers' offerings[5], the supported basic technologies and the range of applications:

Hyperscaler	Supported Technology	Application	Note
Microsoft Azure	Intel SGX Intel TDX AMD SEV-SNP NVIDIA H100*	Virtual Machines Managed Kubernetes (AKS) Confidential AI	*Confidential AI is currently in Private Preview
Google Cloud Platform	AMD SEV-SNP Intel TDX NVIDIA H100*	Virtual Machines Managed Kubernetes (GKS) Confidential AI	*Confidential AI is currently in Private Preview
AWS	AMD SEV-SNP Nitro*	Virtual Machines Managed Kubernetes (EKS)	*Nitro is an AWS-proprietary enclave concept
IBM	Intel SGX IBM SEL	Virtual Machines Container Runtime*	*Part of the IBM Hyper Protect Platform
Oracle Cloud Infrastructure	AMD SEV AMD TSME	Virtual Machines Bare-Metal Machines	
Alibaba	Intel SGX Intel TDX	Virtual Machines Containerised Workloads	

5 See also: https://docs.enclaive.cloud/virtual-hsm/documentation/supported-cloud-configurations

7. Fields of Applications and Use Case Examples

Confidential Computing is particularly suitable for processors of sensitive personal data - as defined, for example, in the GDPR. This includes information about ethnic and cultural origin, political, religious and philosophical beliefs, health, sexuality and trade union membership.

A typical use case is health data, which can only be used to benefit patients if authorised third parties can also access and aggregate this data or exchange the data with other stakeholders as needed. With Confidential Computing, the need to use data does not conflict with the legal obligation to provide this data with special protection.

Since Confidential Computing is a fundamental technological concept, application scenarios can be found across all industries and in companies of all sizes. As described, the benefits of the technology lie in protecting relevant company assets and processes, while also meeting national and international regulatory and standards.

Below you will find a range of examples.

7.1. Confidential Appliances

A software company provides its customers with dedicated analysis environments in the form of highly secure appliances. Confidential Computing ensures that users cannot access the application code and logic. This protects the manufacturer's intellectual property, even when the appliances are deployed outside of the manufacturer's control at the customer site.

7.2. Cloud Platform for Human Resources

A software manufacturer has developed a global, cloud-based platform that enables industry specialists to collaborate across companies and form project-specific teams. The platform, which is operated by a hyperscaler, stores all essential personnel data of experts, including certain health data. Confidential Computing ensures that all personnel data is stored as securely, like in a vault. Customers of the platform provider can only access the sensitive information on the basis of strictly regulated access rights.

7.3. Internet of Things (IoT) as a SaaS Application

An analytics software provider for IoT scenarios has converted its business model from traditional licence sales to Software as a Service (SaaS). Confidential Computing has supported this transformation process. By providing the application in a secure environment, customers have been encouraged to use the software in the cloud instead of installing it locally in their companies. This has enabled the provider to reduce the heterogeneity of its installed base in a manageable period of time.

7.4. Secure Storage of Transaction Data in Retail

Retailers are increasingly required to store transaction data. This involves more than just digitally storing and signing each receipt. Due to the seasonality of the business, the required storage capacity also varies significantly, making the cloud the ideal and economical storage location. One retail company solved this challenge by migrating its existing application redundantly across two clouds. Confidential Computing enables the user to operate the solution compliantly and resiliently. By using a virtual Hardware Security Module (vHSM), the company independently manages its own keys and secrets. This ensures that all customer and payment data remains securely under its control – something that was not the case in the original environment.

7.5. End-to-End Encrypted Communication

A provider of a communication application realised that they could not keep their customer promise of end-to-end encryption on a virtual appliance. The data streams are broken down and decrypted for analysis purposes before the data packets are re-encrypted and forwarded. For this reason, the manufacturer has migrated the central component of its application to a Confidential Computing environment. By doing so, it not only verifiably delivers on its value proposition, but also offers its customers additional flexibility in deploying and securing the solution.

7.6. Flexible Security Environments

A security manufacturer, whose portfolio previously included traditional gateways and appliances, transferred its software to a Confidential Computing environment within two weeks. This environment can now be flexibly deployed for customers in the cloud of their choice or in their own data centre. This has opened up new business models for the provider, such as the situational provision of secure virtual system environments.

7.7 Secure Data Processing in the Healthcare Sector

The healthcare sector is one of the pioneers in the practical use of Confidential Computing. For example, in the implementation of electronic patient files, a solution based on Intel SGX is used: a Trusted Execution Environment (TEE) ensures the trusted environment required by the national agency gematik, thereby protecting sensitive patient data during processing – and also from access by the operators. gematik bears overall responsibility for the telematics infrastructure in Germany's healthcare system.

8. How Can Confidential Computing be Practically Implemented?

In practice, the adoption of Confidential Computing offers a wide range of options for protecting critical processes and sensitive data. Regardless of the approach a company chooses, it is important to ensure that only minimal adjustments are required for application codes, tools and processes. Such an approach makes it easier for companies to quickly test the benefits of Confidential Computing. For instance, in a first step, innovation, data or AI teams can validate a technology within a short period of time. Of course, it is necessary in a second step to involve the infrastructure team, given that operational aspects play a crucial role in successful implementation. However, in the testing phase, they are not yet critical to success.

8.1. Recommendations for User Companies

The introduction of Confidential Computing can vary greatly from company to company, depending on the individual infrastructure and application landscape, the existing data protection measures or the existing knowledge. The following recommendations for action are intended to serve as a generic guide..

8.1.1. Identify Critical Processes and Data

Identifying critical processes and data is the first step towards successfully implementing Confidential Computing. Companies should take a holistic view of their business processes. Which processes are particularly relevant for the company's success? What data is essential for the execution of these processes, and what consequences would data loss or unauthorised use have? In addition, legal and regulatory requirements, such as the GDPR, must be taken into account.

8.1.2. Bring Key Management into Corporate Control

Key management forms the foundation of any cryptographic system. Companies that want to implement Confidential Computing are typically aware of the importance of data protection and therefore encrypt sensitive data. However, key management is often outsourced to service providers or cloud providers' key management services are used. When implementing Confidential Computing, it is advisable to bring all technologies required for key management under internal control. This includes, in particular, Hardware Security Modules (HSMs), which are used to securely manage cryptographic keys. By using their own physical or virtual HSMs (vHSMs), companies gain complete control over all the keys and secrets they need for their applications, regardless of whether these applications are running locally (on-premises), in the cloud or in a hybrid environment.

8.1.3. Identify Initial Pilot Applications

When introducing Confidential Computing, it is recommended to start with pilot projects. A few representative applications should be selected to test the concept in a multi-cloud environment. Ideally, a Multi-Cloud Confidential Computing Broker (MCCCB) should be used for this purpose. These pilot projects serve as a learning phase to understand the specific requirements of the application in a confidential execution environment and to identify potential challenges at an early stage. Based on these insights, specific use cases can then be developed that fully leverage the advantages of Confidential Computing.

8.1.4. Select Suitable Technologies and Service Providers

A range of products and services for Confidential Computing are available on the market. When selecting them, various questions arise - for instance, about the maturity of a solution, references or compatibility. Is the technology easy to integrate into the existing IT landscape and which cloud platforms does it support? In principle, users must themselves decide whether they want to develop a solution themselves or whether they want to implement a largely finished product that they only need to adapt to their specific requirements. In either case, companies should also address the question of the extent to which they can implement the project with their own human resources or the extent they need support from external service providers.

8.1.5. Secure Operating Models and Tools

The adoption of Confidential Computing requires adjustments to existing operational processes and tools. In addition to the implementation of new technologies for key management and the configuration of secure execution environments, continuous monitoring and maintenance is essential. By monitoring the systems, potential security risks can be identified and eliminated at an early stage. In addition, software and hardware should always be kept up to date to close known vulnerabilities and enhance the resilience of the systems.

8.1.6. Evaluate Deployment and Migration Options

The introduction of Confidential Computing requires careful planning of the migration strategy. To minimise risks and reduce the impact on ongoing operations, a step-by-step approach is recommended. Based on pilot projects, companies can expand their knowledge and continuously improve their Confidential Computing scenarios. Close collaboration between the IT department and the business departments is helpful in this regard. Targeted training can raise awareness of the new technology among employees and enable them to fully utilise its advantages. A long-term strategy ensures the sustainable use of Confidential Computing and enables continuous adaptation to evolving requirements.

8.2. Recommended Actions for SaaS Providers

Cloud Computing has given software providers the opportunity to provide applications as a service. The SaaS model relieves customers of the need to procure server infrastructure and offers them flexible usage options. By operating their SaaS applications in enclaves, providers enhance the security of data processing, as the data is protected from unauthorised third-party access. This, in turn, strengthens trust in their offering.

Thanks to the availability of Multi-Cloud Confidential Computing Brokers (MCCCB) and open-source tools, SaaS providers can test the concept in a matter of days. This allows them to evaluate how their applications function in confidential execution environments and whether Confidential Computing represents a viable business case for them. For the relevant technologies, the usual OEM models on a subscription basis apply.

8.3. Recommended Actions for Service Providers

Confidential Computing also holds potential for service providers, enabling them to enhance the security of their infrastructure services. A three-phase approach is recommended.[6] In the first step, the foundations should be established, i.e. identifying the server hardware (or procuring new hardware if necessary) that supports Confidential Computing, defining clear guidelines for access control, key management and logging, and training their teams in the new technology.

The second phase is dedicated to building the Confidential Computing environment. In this phase, service providers should develop a virtualisation concept that can be used to manage confidential VMs and, ideally, also confidential containers. Moreover, they should test the environment intensively for security. In the service deployment phase, the aim is to offer confidential cloud services that address the needs of specific industries, build a partner ecosystem, acquire specific certificates and actively market the advantages of Confidential Computing.

At this point, detailed recommendations for both SaaS providers and service providers would go beyond the scope of this white paper, which primarily covers the fundamental functionality and advantages of Confidential Computing. Such recommendations for action could be the subject of another, more practice-oriented document.

6 See also: https://next.enclaive.cloud/s/rssJ5c54RJfLCp5J

9. Conclusion and Outlook

Confidential Computing is not a fundamentally new concept. However, with the technologies available today - and thanks to the support of cloud and managed infrastructure providers - it is no longer just a theoretical possibility, but can be practically realised. In fact, the feasibility has now been simplified to such an extent that users and software providers can now leverage the concept's possibilities without in-depth cryptographic knowledge.

By combining hardware-based security mechanisms, independent key management and the ability to validate the security status of a used environment, Confidential Computing offers a robust solution for protecting data from unauthorised access. With the help of this technology, companies can securely process their sensitive data in the cloud and thus use the hyperscaler platforms in a compliant manner.

Together with a consistent Zero Trust model, companies have two powerful approaches available to help them ensure generic protection gaols such as confidentiality and integrity, while constantly improving. Especially in industries such as finance or healthcare, where the protection of sensitive data is of the utmost importance, Confidential Computing opens up new opportunities for innovative business models and strengthens the trust of customers and business partners.

However, the introduction of Confidential Computing is not a trivial matter. In addition to technical aspects such as performance and costs, organisational aspects such as process adjustments and training must also be taken into account. These challenges can nevertheless be overcome through careful planning and continuous monitoring.

The future of Confidential Computing is promising. With further development of hardware and software, as well as integration with other technologies, new application possibilities will emerge. Companies that adopt the concept on an early stage can gain a competitive advantage and future-proof their business models.

Authors of the Text

Achim Astel
noris network AG

Anna Fischer
secunet Security Networks AG

Prof. Dr. Sebastian Gajek
Flensburg University of Applied Sciences

Nicolas Maeding
IBM Deutschland
Research & Development GmbH

Prof. Dr. Norbert Pohlmann
Institute for Internet Security –
if(is) at the Westphalian University

Andreas Walbrodt
enclaive GmbH

EuroCloud
DEUTSCHLAND | eco

EuroCloud Deutschland_eco e.V.
Lichtstr. 43h
50825 Cologne, Germany

Phone: +49 (0) 221 / 70 00 48 – 0
Fax: +49 (0) 221 / 70 00 48 – 111

Email: info@eurocloud.de
Web: https://www.eurocloud.de/