GlobalLogic®

Confidential Computing: Third Pillar of Data Encryption

Vikas Mangla, Director of Engineering, GlobalLogic

Introduction

*Gartner lists "**privacy-enhancing computation**" as one of the Top Tech Trends for 2021*[10]

In order to achieve their next level of growth, Cloud service providers are continuously working, among other things, to address the concerns of privacy, security and confidentiality of customer data.

Security is an integral non-functional requirement for digital data management in the current political, sociological and economic landscape. A range of concepts from authorization, authentication, to specialised concepts of homomorphic encryption and Trusted Platform Module are available in space.

Data is ubiquitous these days across enterprises and systems. Before Confidential Computing, there existed a key gap in the data security landscape. The data in use has been under constant threat from abuse and theft making many industry sectors apprehensive of cloud strategy.

Data can exist in three different states:

a. **Data-in-transit:** In motion into, out, or within the cloud.

b. **Data-at-rest:** Stored in one of the many cloud services like File System, Containers, Database, etc.

c. **Data-in-use:** Active data undergoing analysis, change, or other manipulation.[1] This is also referred to as active data, which is stored in a non-persistent digital state typically in RAM, CPU Caches, or CPU registers.[2]

When a user makes a request to the cloud, all data in transit is recommended to be secured using HTTPS, TLS, or IPSec Virtual Private Networks. The best practices for protecting data in transit on cloud are well-defined. For instance, all traffic to Google is routed through globally distributed Google Front-end which encrypts the traffic using BoringSSL (an open-source implementation of TLS protocol).

There are several approaches available for customers for securing the data at rest. Approaches range from:

- Server-side or client-side encryption and several key management options from Cloud Service Providers. For instance, AWS provides KMS for key management for server-side encryption.
- Application or field-level encryption for their data.
- Standard libraries available from application frameworks like .net and java
- Transparent Data Encryption (TDE) functionality in Microsoft SQL or Oracle
- Through integration of third-party solutions in their applications
- File-level or Full disk encryption (FDE)

1. The Cyberwire. Definition of data in use, available at https://thecyberwire.com/glossary/data-in-use (accessed March 27, 2021).
2. Wikipedia. Data In Use, available at https://en.wikipedia.org/wiki/Data_in_use (accessed March 27, 2021).

In addition to these techniques for securing data at rest and data in motion, there are other overarching concepts which help further securing the systems like Identity and Access Management (IAM), Role Based Access Control (RBAC), Cloud Access Security Brokers (CASB), and Data Encryption.

For data in use, Conventional computing infrastructure and concepts like homomorphic encryption are limited in its ability and, at times, found wanting. Confidential computing aims to address the gap in securing data in use.

Confidential Computing and How It Helps

Confidential computing is a cloud computing technology that isolates sensitive data during processing. There are many industries like Healthcare, BFSI, Oil and Seismic where confidential computing will help expand cloud adoption and provide use cases to;

- Protect customer information (patient information in Healthcare Industry, customer information in Financial Sector)
- Secure health or financial data
- Blockchain financial operations
- Run ML/AI processes on sensitive information
- Perform algorithms on encrypted data sets from multiple sources

For example, Microsoft announced Azure as the first cloud to offer Confidential Computing back in 2017.[3] Which has since been adopted by other cloud providers with their own proprietary implementations.

The industry initiative Confidential Computing Consortium (CCC)4 started under The Linux Foundation to accelerate the adoption of Trusted Executive Environments (TEE) technologies and standards to help secure data in use. CCC defines Confidential Computing as:

"The protection of data in use by performing computations in a hardware-based Trusted Execution Environment (TEE)."[4]

TEE (also known as an Enclave) is an environment that enforces execution of authorised code only. Data in the TEE environment cannot be read or tampered with outside its boundaries. The data being processed and the techniques used to process it are accessible to authorized programming code only. This information is invisible or known to anyone including the cloud provider providing greater confidence and assurance to customers on security of their data in the cloud.

IBM Cloud defines confidential computing as "hardware based technology that allows for physical partitioning of the memory at the server level."

3. Microsoft Azure. Introducing Azure confidential computing, available at https://azure.microsoft.com/en-us/blog/introducing-azure-confidential-computing/ (accessed April 17, 2021).
4. Confidential Computing Consortium. What is the Confidential Computing Consortium? available at https://confidentialcomputing.io/ (accessed April 17, 2021).

How Confidential Computing Works

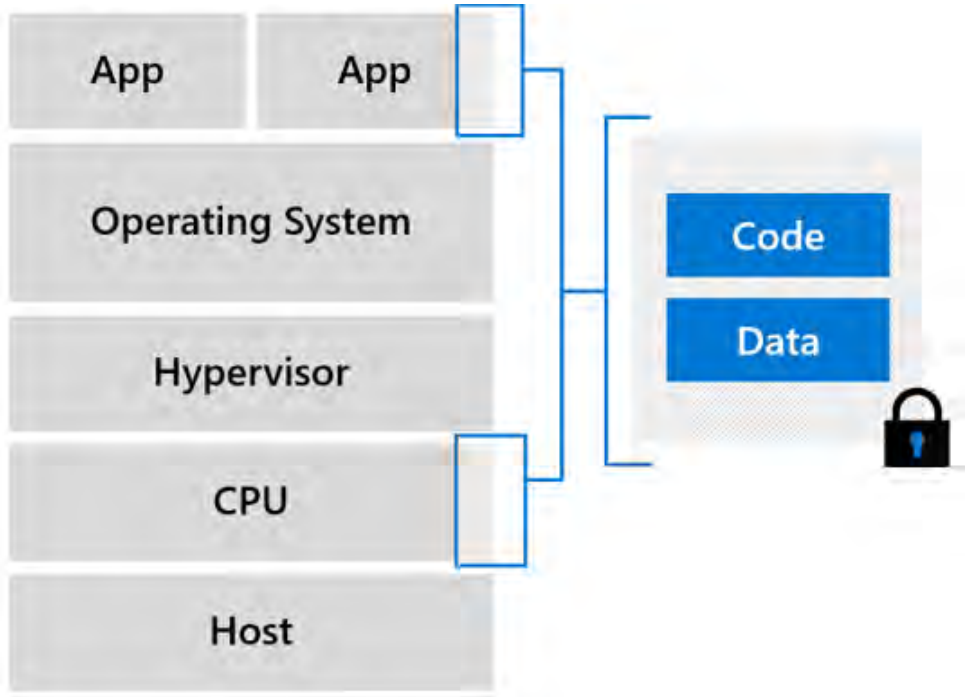

Figure 1: Confidential computing: In Azure confidential computing virtual machines, a part of the CPU's hardware is reserved for a portion of code and data in your application | Image: Microsoft Azure

The Trusted Computing base (TCB) of a computer system consists of all hardware, firmware, and/or software components critical to its security.[5] TEE aims to lower the TCB to trusted runtime binaries, libraries, and code. TEE is a memory black box which contains data and code.

The content and processing inside TEE is completely insulated from unauthorised access. TEE provides data code integrity by:

a) reducing visibility to authorised applications,
b) prevent unauthorised modifications, and
c) provide validation through attestation reports.

TEEs are enabled through hardware based isolation technologies like Intel SGX or ARM TrustZone or through software-based options like Microsoft's Virtual Secure Mode (VSM) implemented by Hyper-V. For example, Microsoft DC-series is backed by the latest generation of Intel Xeon E-2176G 3.7GHz processor with SGX Technology. AWS Nitro Enclaves are available in bare metal and virtualized options. Google Confidential VM runs host based on second generation AMD Epyc processors code-named "Rome" and provides confidential computing through AMD Secure Encrypted Virtualization (SEV).

5. Wikipedia. Trusted computing base, available at https://en.wikipedia.org/wiki/Trusted_computing_base (accessed April 17, 2021).

Attestation verify trustworthiness of TEE and integrity of the binaries running inside. Attestation provides increased confidence to the relying party that their software is (1) running in an enclave and (2) that the enclave is up to date and secure. Microsoft Azure Attestation, now generally available[6], enables the attestation process for code running in TEE. AWS Nitro enclaves use cryptographic attestation. Integrity monitoring is enabled, by default, in Google Confidential VM instances.

Most of the major technology players are contributing to confidential computing in some way. For example, Microsoft is contributing to OpenEnclave[7] aimed at creating a single unified enclave abstraction to build TEE based applications. Google Asylo is an open and flexible framework to build applications that can run in TEE. RedHat Enarx enables hardware independence for applications running in TEEs.

How Confidential Computing is Helping Address Industry Concerns

There are three classifications of concerns where confidential computing can help the most:

- **Business**, including protecting sensitive information critical to the business while leveraging Cloud computing or collaborating with partners. For instance, related to IP protection, PII, PHI, Company confidentiality, and Multi-party Computing.

- **Technology**, where concerns include the privacy, security and confidentiality of data while in use. Human error, compromised credentials, insider threats, memory dumps and privilege access at Cloud providers are all considerations.

- **Regulatory and compliance** such as HIPAA compliance in the health industry.

Federated learning (also known as privacy preserving analytics) is an application of confidential computing that allows for data aggregation and analytical insights while preserving the confidentiality of a stakeholder's customer base.

Specific business-related use cases can include Money Laundering, Credit qualifications, Market rate calculations, Credit Scores, Load Fulfillment, Know-your-customer (KYC) and Stock trading from BFSI Sector. Through confidential computing, banks and others will be able to run agreed-upon machine learning algorithms on the combined sensitive data set securely without comprising individual bank sensitive data, for example. This can help detect money laundering by a customer across multiple banks or loan fulfillment on the same property by a customer across multiple banks to facilitate fraud detection.

Beyond healthcare use cases of protecting health data, Confidential computing will also help in areas of federated medical research without sharing the proprietary data.

6. Microsoft Azure. Microsoft Azure Attestation is now generally available, available at https://azure.microsoft.com/en-us/updates/azure-attestation/ (accessed April 17, 2021).
7. Open Enclave SDK. What is Open Enclave SDK? available at https://openenclave.io/sdk/ (accessed April 17, 2021).

Conclusion

With the Confidential Computing Consortium driving this initiative, Confidential Computing in conjunction with homomorphic encryption and Trusted Platform Module is poised to fill the data security gap and enable businesses to confidently adopt the cloud.

Enterprises can leverage and exploit confidential computing to address security concerns that were a roadblock in cloud adoption for business use cases and data in the past.

Want to learn more? Let's work together to explore the ways in which confidential computing can solve challenging issues for your business.

About the Author

Vikas Mangla is a passionate program manager who always thinks of people and the customer first, and blends it with technology and processes. He has over 18 years of experience in Product Engineering, consulting and advisory roles, sales, and IT services. Experienced with cloud computing, Machine Learning, the latest agile processes, custom application development, and architecture patterns, he is also well-versed in Microsoft Office and Google productivity tools.