Case Study

Deep Learning
Confidential Computing

Deep Learning Combined with Confidential Computing to Offer Enhanced Security and Full Functionality

This article explores how Intel® Software Guard Extensions (Intel® SGX) plays an important role in building the Baidu MesaTEE* confidential computing platform, which is now linked with the Baidu PaddlePaddle open-source deep learning platform to achieve confidential computing capabilities – specifically customized for deep learning. It also offers a glimpse into the platform's potential applications in various fields.

Introduction

We live in a new era driven by big data and AI. Throughout the world, industries are realizing the significance of data and are constantly driven to explore its enormous potential. However, in recent years, the circulation and aggregation of data have also become topics of concern. As data continues to expand and diversify, companies are eager to aggregate data from multiple sources, with greater dimensions and of higher quality, then leverage the data using advanced technologies such as deep learning. With sensitive data being transmitted, shared and processed, there has been growing concern about data security. From a regulatory perspective, both worldwide and in China, more and more stringent data security regulations are being introduced, making compliance an ever-challenging business issue.

Confidential computing has grown in prominence, in response to the demand to share data while safeguarding data at the same time. Intel has launched the Intel® Software Guard Extensions (Intel® SGX) technology – ever since the 6th Generation Intel® Core™ processors – as a confidential computing solution. Intel SGX is a hardware-based security solution that bypasses a system's operating system (OS) and virtual machine (VM) software layers, partitions sensitive application scripts and data into hardened enclaves, giving them more protection from disclosure or modification.

This article explores how Intel SGX plays an important role in building the Baidu MesaTEE* confidential computing platform, which is now linked with the Baidu PaddlePaddle open-source deep learning platform to achieve confidential computing capabilities – specifically customized for deep learning. It also offers a glimpse into the platform's potential applications in various fields.

Building the Baidu PaddlePaddle Confidential Deep Learning Platform

About PaddlePaddle

Deep learning is the acclaimed "operating system" of the AI era. As a renowned technology giant in China, Baidu has been committed to deep learning technology research for many years, with PaddlePaddle as its vital offering. Designed as a cloud-based distributed open-source deep learning platform, PaddlePaddle's

mission is to make the innovation and application of deep learning technology easier and more accessible to users. Today, PaddlePaddle has a large and active developer community in China. It provides a rich collection of models as well as a full range of high-performance deployment and integration solutions for developers. It has been widely used in manufacturing, agriculture, and service industries.

Ever since opening its source in 2016, PaddlePaddle is continuously being optimized and upgraded. In the past two years, in order to help companies and developers apply deep learning models to explore the value of data integration in a more secure manner, PaddlePaddle has built a software- and hardware-based solution for confidential computing within a contained ecosystem.

The software-based solution is achieved through MPC (Secure Multi-Party Computation), with Paddle FL (an open-source federated learning framework) as its foundation. In PaddleFL, several federated learning strategies are provided with applications in computer vision, natural language processing, algorithm recommendations and more. Applications of traditional machine learning training strategies are also provided, such as transfer learning in federated learning settings and multi-task learning. Based on PaddlePaddle's large-scale distributed training, and the flexible scheduling of training jobs on Kubernetes, PaddleFL can be easily deployed on full-stack open-sourced software.

The hardware-based solution is achieved by creating a Trusted Execution Environment (TEE) with Baidu's confidential computing framework, MesaTEE, as its foundation. MesaTEE can provide TEEs enabled by Intel SGX. By avoiding software-level encryption, this solution allows PaddlePaddle to run in TEEs in a fast and secure way, capable of fighting against server-side malicious attacks.

MesaTEE and Intel SGX

In September 2018, Baidu Security and Intel jointly released MesaTEE, a memory safe computing framework for confidential computing. MesaTEE was developed through the combination of Rust and Intel SGX, and has two versions, a commercial version (MesaTEE) and community version (Teaclave). Teaclave opened its source on GitHub in mid-2019 and was donated to Apache to become an Apache Incubator project at the end of that same year. The commercial version MesaTEE builds on Apache Teaclave to offer customized commercial solutions to corporate users. While supporting TEEs enabled by Intel SGX, it has also created a first-of-its-kind collaborative confidential computing engine. These two computing engines make privacy-protected big data analysis and machine learning possible. With MesaTEE offering enhanced privacy protection, organizations across the industry chain can break down data barriers, solve legal compliance and data security issues, and achieve data sharing and collaboration, which unleashes the value of untapped data. While the value and insights behind big data can be discovered and shared, sensitive data itself remains truly "available but not visible".

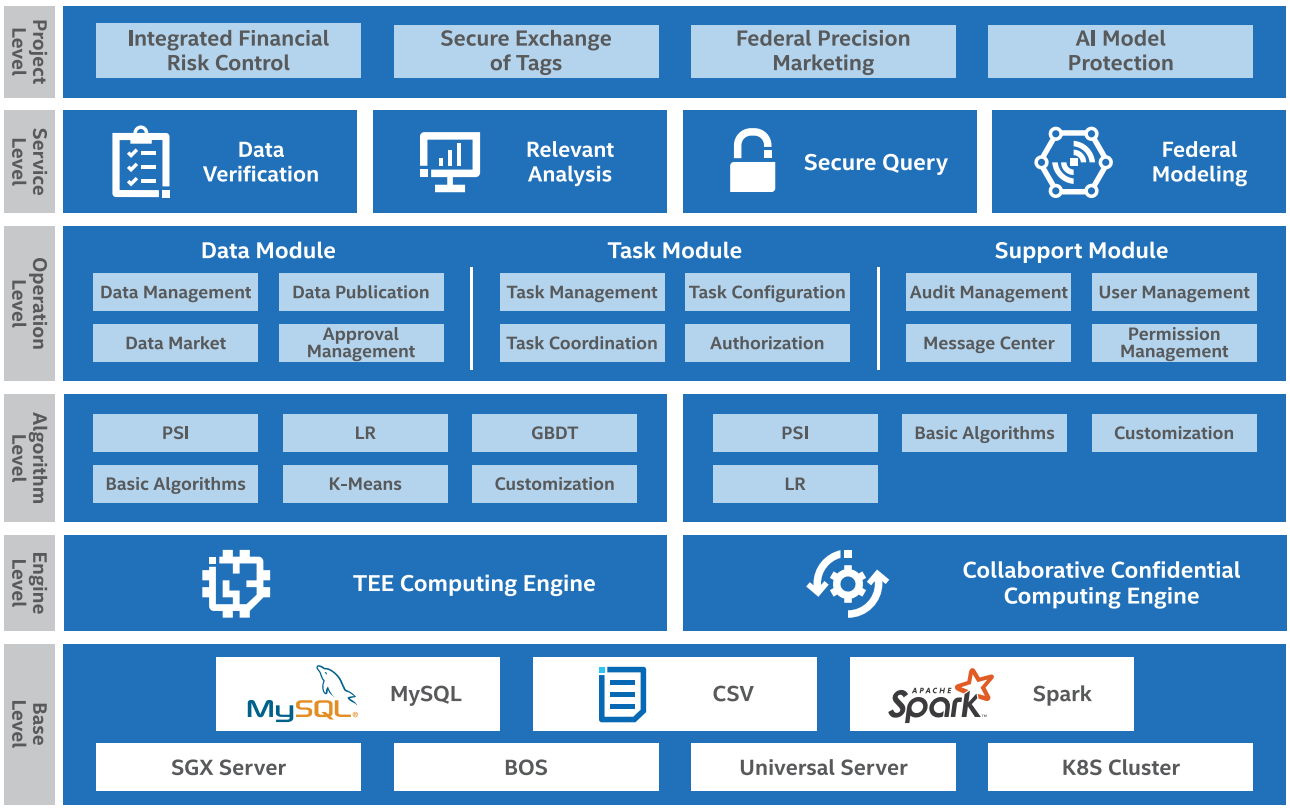

Figure 1. Overall Structure of Baidu MesaTEE

MesaTEE and PaddlePaddle

As a universal confidential computing platform, MesaTEE runs entirely in a TEE and is developed in the Rust language. It helps protect the complete process - from data upload, data processing to result acquisition. In addition, MesaTEE is a collaboration platform that allows multiple participants to collaborate on tasks at the same time. It manages the collaboration process through tasks, similar to discrete contracts, and eliminates interference from external environments. This helps to ensure that all participants can cooperate without having to reveal confidential information, and still equally obtain the necessary results they need from the model.

Real-world user application scenarios are extremely diverse, and everything from security, function, performance, and even user habits, need to be considered. To suit various needs, MesaTEE provides five security levels. While L1 offers the most basic security, and is therefore the lowest security level, it also offers the most comprehensive functions and the fastest speed. L5 represents the highest security level, providing the increased security protection, but with some restrictions on functions and performance. By default, MesaTEE is set at L5, the highest data security level, which is paired with an optimized computing speed. And in order to adapt to a wide range of application scenarios and security requirements, MesaTEE can also provide different types of execution environments through Executor plug-ins, including TEEs such as SCONE, Occlum, and Graphene.

In order to run PaddlePaddle in TEEs that are similar to LibOS, its underlying runtime library needs to be modified. The enclaves of Intel SGX are resource-constrained environments that strictly restrict system calls, data entry/exit, network access, and memory. PaddlePaddle was originally designed to run in glibc without the ability to run in a resource-constrained environment. Migrating glibc to TEEs proved challenging, often requiring cropping or patching. What's more, very few vendors support glibc, making it difficult to maintain. In contrast, most vendors choose to use musl (a runtime library that is newly written for restricted environments and implements POSIX standards) as the calling interface for TEEs and upper-level programs. In order to run in existing TEEs, PaddlePaddle, with collaboration with Baidu Security, modified its underlying call interface library, and provide WITH_MUSL compilation options as well as fully automated compilation scripts. PaddlePaddle will continuously audit the security of its underlying runtime library and reduce potential risks. The modified PaddlePaddle is fully available to the community and constantly updated.

As a collaboration platform and task scheduler, MesaTEE uses PaddlePaddle as a task execution environment through Executor plug-ins and delivers specific deep learning tasks to PaddlePaddle in TEEs. In addition, MesaTEE works with different TEE vendors to adapt the validity of its Remote Attestation process and to integrate it into the communication protocol to ensure that the security of its operating environment is measurable. In terms of file access, remote file storage protocols such as S3 are supported, and all file contents are encrypted. There is also a special file system for PaddlePaddle to access the required cyphertext in TEEs during the training process without affecting the usage mode. The capability is designed to resist malicious attacks by preventing access to any computing content from outside the TEE.

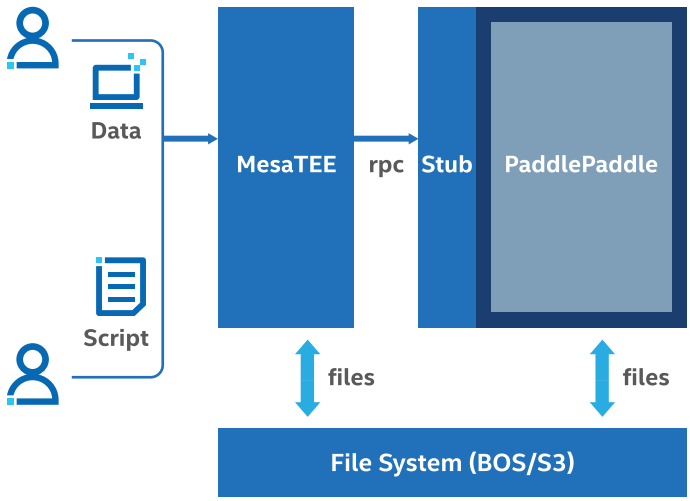

Figure 2. The interaction between PaddlePaddle and MesaTEE

Benefits of the Solution

There are several benefits of this hardware-based solution to enable PaddlePaddle confidential deep learning capabilities:

A. **Efficient:** Empowered by the hardware-enhanced Intel SGX, PaddlePaddle can operate in an embedded TEE at a higher speed. When running on the new generation Intel® Xeon® Scalable processors, it can operate in a high speed. In addition, the hardware-based solution is free of network restrictions, eliminating the risk of large-scale training tasks being interrupted by network failures.

B. **Low cost:** The cost advantage is self-evident. Traditionally, companies invest heavily in solutions such as specific encryption devices to enhance data security. The combination of PaddlePaddle and TEE is like an extensive encryption device. Users who run the applications on compatible Intel processors do not need to pay additional costs.

C. **Easy to use:** Installation time is as short as three minutes with only three simple steps.[1]

D. **Compatible with multiple containers:** It can support most SGX TEE containers, including SCONE, Occlum and Graphene.

Potential Applications of the PaddlePaddle Confidential Deep Learning Platform

Across industries, traditional machine learning models can no longer keep up with user needs, especially with the exponential growth of data and the complexity of data formats across text, images, audio, multimedia, and so on. In contrast, deep learning possesses many advantages. At the same time, in order to break the barriers between data providers, model developers, and data users, confidential computing has become a necessity. By combining deep learning with confidential computing, PaddlePaddle is flexibly responding to industry trends and will continue to play a critical role in many industries.

Medical Image Screening

Medical image processing concerns images, rather than feature vectors in the traditional sense. To recognize images, a combination of image processing techniques and machine learning technologies are used under the traditional machine learning approach. This approach is inefficient and unable to cope with huge amounts of graphic data. The advancement of deep learning technology pushes medical image screening capabilities to new heights. It leverages computer vision to offer an integrated and comprehensive solution, from feature extraction to overall image recognition. With deep learning performing image classification, recognition, positioning, detection, and segmentation at fast speeds and enhanced accuracy, medical image recognition is becoming more efficient than ever.

However, while deep learning is the way forward, the technology alone cannot solve the problems that occur in the real world. Successful model training feeds on a large amount of patient image data, which inevitably involves a large volume of confidential information, therefore requiring extremely high standards of data security. Furthermore, when multiple medical institutions access the same model, they are obligated to strictly protect their patients' information from being used elsewhere. In addition, model developers, whether third-party developers or the medical institutions themselves, need to protect the model so they can maximize their commercial benefits. Naturally, confidential computing is the missing link to ensuring a complete medical image processing solution that offers all necessary privacy safeguards.

The PaddlePaddle platform running in TEEs offers a comprehensive solution that addresses these concerns. A case in point is an ocular fundus blood vessel image segmentation system. Typically, ocular fundus images are noisy, and the retinal blood vessel structure is complex. The system developed on PaddlePaddle can overcome these challenges to extract vascular features in ocular fundus images in real time, playing a significant role in the diagnosis of ophthalmology, cardiovascular and cerebrovascular diseases.

Another successful use case is a CT image analysis model for pneumonia diagnosis. Developed on PaddlePaddle, this system efficiently detects, identifies and outlines lesions, then runs a series of post-processing codes to analyze and output a full set of quantitative indicators such as the number, size and proportion of lung lesions. Its deep learning model has been trained on both high-resolution and low-resolution CT images, so it can recognize CT images of varying qualities. It has the potential to be an effective diagnostic tool for primary hospitals with limited medical resources and capabilities. More importantly, at such a critical time in the COVID-19 pandemic, the system can effectively help relieve the burden on medical staff, speed up patient diagnosis and treatment, alleviate the strain on limited medical resources and contribute to the successful fight against the pandemic.

At the same time, paired with the confidential computing capabilities of PaddlePaddle, these systems can also fully protect patient privacy and eliminate any barriers between data providers, model developers and data users. In the medical industry which struggles with issues such as frequent data leaks and low data usage efficiency, these solutions forge a new path for the future of smart medicine.

Financial Risk Control

Risk control is the foundation of the financial industry. Take credit risk assessments for example. The main purpose of such assessments is to conduct a detailed analysis, evaluation and prediction of the borrower's credit rating and repayment ability. For a long time, financial institutions relied on machine learning technology to model and analyze simple, two-dimensional numeric data to establish borrower profiles. The rise of deep learning technology has opened up new possibilities. Deep learning models break through the limitations of traditional machine learning models by using voice, image as well as other methods to incorporate

multimedia data and streaming data. Such data can include anything from payment data (from third-party payment service providers), purchase history (from retailers and e-commerce platforms), credit data (from third-party credit institutions), and even social contacts (from social media platforms), to employment history (from recruitment platforms) and travel information (from travel platforms).

Traditionally, financial institutions use a distributed system to establish a credit risk assessment model and consolidate the data in a single place. However, from the perspective of data providers, user privacy data is highly sensitive and may not be shared by law. As a result, in real-world situations, data silos are commonplace in the financial industry. Data providers, model developers, model evaluators and data users are all separated from each other. In order to achieve true data aggregation, confidential computing is the missing piece.

Through the PaddlePaddle deep learning platform running in TEEs, financial institutions will be able to work with data providers in a more secure manner. They can obtain data that would otherwise be unavailable to create a more robust credit risk assessment model. The PaddlePaddle confidential deep learning platform will greatly boost the confidence of all parties, reduce the risk of data leakage, divide data ownership clearly and ensure that the transmission of data complies with regulations. It is a one-stop solution that addresses the various pain points of having multiple users share and use sensitive data and satisfies the growing demand from the financial industry to tap into the value of data sharing and aggregation.

Looking Forward

In December 2020, PaddlePaddle 2.0 will be officially released. A number of important upgrades and features will be introduced, including the realization of all existing PaddlePaddle functions in TEEs, and pre-compiled musl dependences, which users can directly install in TEEs on their own.

On top of the above mentioned MesaTEE and PaddlePaddle, Baidu has also released the Dianshi big data open platform on a higher level. Dianshi integrates Baidu's full-stack technology capabilities in big data, as well as complex and diverse big data tools to form a one-stop development service for data processing and application. For data providers, system integrators, independent software developers, consulting companies, model developers and other different types of partners, Dianshi provides unique solutions to accurately meet their specific needs.

At Intel, the upcoming new generation Intel® Xeon® Scalable processors will feature the next release of Intel SGX with an Enclave Page Cache (EPC) that jumps exponentially from 128MB (single-channel) to now include up to 512GB per socket, signifying a giant leap in confidential computing performance using the Intel SGX TEE.

In terms of collaboration, Intel and Baidu will continue their long-term vision to develop innovations. By applying Intel technologies to empower Baidu platforms, services and products with higher performance and more outstanding features, which ultimately benefits the entire industry ecosystem.

intel®

[1] The results were provided by Baidu. For more complete information about these test results, please contact Baidu. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.