UTEE: A Secure, Efficient, and Portable Distributed Bigdata Computing System on Heterogeneous Trusted Execution Devices

Heming Cui
Computer Science/Engineering

May 8, 2022

Summary of the Impact

- Impact case: a preliminary prototype of UTEE, Dr. Cui's secure big-data system, has commercially become **the core system** of Huawei Cloud's TICS (Trusted & Intelligent Cloud Services)
  - UTEE (interchangeable, in ppt) == TICS
  - Huawei claims that UTEE is **world-leading** (i.e., leading the current global industries' standards and levels in the big-data security area), which matches a **RAE 4-star**
  - UTEE won **ITF** grant (PI, **HK $3.3M**)
  - Dr. Cui won **nine grants (PI)**, totally HK **$16.5M**, in his 7-year HKU career
  - Filed 13 patents in past 3 years
  - Publish 14 international best-tiered (**CCF A**) journals and conferences, all papers are good candidates of **RAE 4-stars**

www.huaweicloud.com/product/tics.html

Impacts Achieved (Best Collaborating Scientist Award 2021 from Huawei)

- "**The Best Collaborating Scientist Award Medal 2021**" of the Huawei Theory Research Department, in Oct 2021
  - Huawei Theory Research Department (or the "华为中央研究院理论研究部" in Chinese) is Huawei's strategic research department on conducting information theory, optimization, and system research with worldwide competitive universities and their professors
  - This award selects this award for the most-productive professor among this Huawei research department's global leading university collaborators around the world, including Europe, Canada, and Asia

Underpinning Research (4 Objectives)

- [Objective 1]: Create a new secure and high-level (easy-to-use) programming language for enabling TEE features on CPU and protecting big-data analytics.
- [Objective 2]: Develop a new secure programming language for enabling TEE features on commodity GPUs.
- [Objective 3]: Invent new security protocols and runtime systems for protecting big-data/AI inference tasks (consisting of secure heterogeneous computations, i.e., computations on both CPU and GPU).
- [Objective 4]: Invent the first security protocols and integrated runtime systems for supporting all three popular types of big-data/AI tasks in industrial practice, including training/inference/analytics

Underpinning Research (Role)

I (Dr. Heming Cui) am the PI of this project

Underpinning Research (Partners)

In this big-data and AI era, people usually conduct computations on sensitive user data using intelligent software (e.g., Spark, Flink, and PyTorch), where the computation tasks (typically, analytics/training/inference) of each software are spitted and ran on a massive number of "heterogeneous" hardware (i.e., CPUs and GPUs) in public clouds (e.g., **Amazon Cloud, Huawei Cloud, and Hospitals that conduct big-data/AI computations**). To protect the confidentiality (privacy) and integrity of these tasks and the user data against malicious/flawed software and administrators with high privileges (e.g., OS kernels with flaws) in public clouds, in recent years, Trusted Execution Environments (TEE) are introduced in high-end processors (e.g., Intel SGX and Arm TrustZone).

However, despite much effort and systems in worldwide academia and industries (including Microsoft's VC3, European Union's SGX-Spark, UC Berkeley's Opaque, Google's Confidential Computing, and UNC's Civet), TEE still faces **three significant open research problems**: (1) it is still a low-level processor feature and extremely hard to use by regular software developers; and (2) such TEE features are merely mature in CPU, but not in commercial GPU; (3) existing research and industry competitors suffer from slow performance and cannot support typical big-data datasets. Our proposal aims to tackle these three problems via developing the system of **UTEE (Ubiquitous Trusted Execution Environments)**.

Underpinning Research (Innovations)

- Collaborated with Huawei to submit 13 Patents to WIPO or CNIPA, from 2019~2022
  - First-author of all the submitted patents are Dr. Cui's RPG students
- Published 14 international best-tiered papers (China Computer Federation, or CCF, A class journals or conferences)
  - First-author of most papers are Dr. Cui's PhD students
  - CCF A papers are all good candidates for RAE 4-stars
- Dr Cui got nine competitive PI Grants (HK 16.5M)
  - E.g., an ITF ITSP Platform grant (HK 3.3M) in May 2022

13 Patents submitted to CNIPA/WIPO

Note: all these patents' first authors are Dr. Heming Cui's RPG students

- **CRONUS: Fault-isolated, Secure and High-performance Heterogeneous Computing for Trusted Execution Environments, CNIPA app ID to be allocated**
- Fold3D: High-performance 3D Parallel DNN Training via Parallelizing Computation and Communication Tasks, CNIPA app ID to be allocated
- NASPipe: High Performance and Reproducible Pipeline Parallel Supernet Training via Causal Synchronous Parallelism, CN 202210138879.0
- **SOTER: Guarding Black-box Inference for General Neural Networks at the Edge, CNIPA app ID to be allocated**
- Themis: Automatic and Efficient Deep Learning System Testing with Strong Fault Detection Capability, CN 202111372034.X
- vPipe: A high-performance DNN training system with efficient and scalable pipelined parallelism on GPUs, 92000692CN01
- Dast: A system in achieving low tail-latency and high scalability for serializable transactions in edge computing, CN 2021101523346.3
- **BIDL: A High-throughput, Low-latency Permissioned Blockchain Framework for Datacenter Networks, CN 202111080651.2**
- **Upa: An Automated, Accurate and Efficient Differentially Private Big-data Mining System, CN 202010506698.X**
- **Daenet: A decentralized, secure and reliable network communication system via SGX, CN 202110048599.6**
- **Uranus: An Efficient, Secure Big-data Processing and Programming System based on Trusted Execution Environment, CN 202010366539.4**
- **Eges: An Efficient, DoS Resistant Consensus Protocol for Permissioned Blockchains, CN 202010247629.1**
- Plover: A Distributed Fault-tolerant Storage System via Virtualized State Machine Replication, 85714660PCT01

Trusted Execution Environment (TEE) Become Available in ARM/AMD/Intel CPUs, But Extremely Hard to Use

- Security problems on public clouds (AWS/Huawei clouds) are severe
  - Cloud management software (e.g., Virtual Machine, VM) and Operating Systems (OS) have high privileges on tenants' normal data server threads
  - VM and OS are developed by 1K+ diverse organizations worldwide
    - Software defects: Minnesota Linux Hypocrite Commits incidents in Apr 2021
  - Also, some admins of the clouds may be careless or malicious
- TEE's hardware-level strong isolation prevents privileged admin/software peeking data being queried by SQL/big-data engines
  - A thread running outside does complex assembly to transit into a TEE, decrypts sensitive data, processes it, encrypts result, and returns to outside
- TEE's hardware nature makes it ***extremely hard*** to use
  - TEE products in Arm/AMD/Intel supports software are in C/assembly
  - But cloud software is often written in Java (e.g., big-data engine, Spark)
  - Approaches: (1) **shielding** an entire server in a TEE (including server software defects and occupying too much TEE hardware resources), or (2) **rewriting** Java into C/assembly (error-prone) and run in TEE

A computer running OS, VM, and data servers on clouds

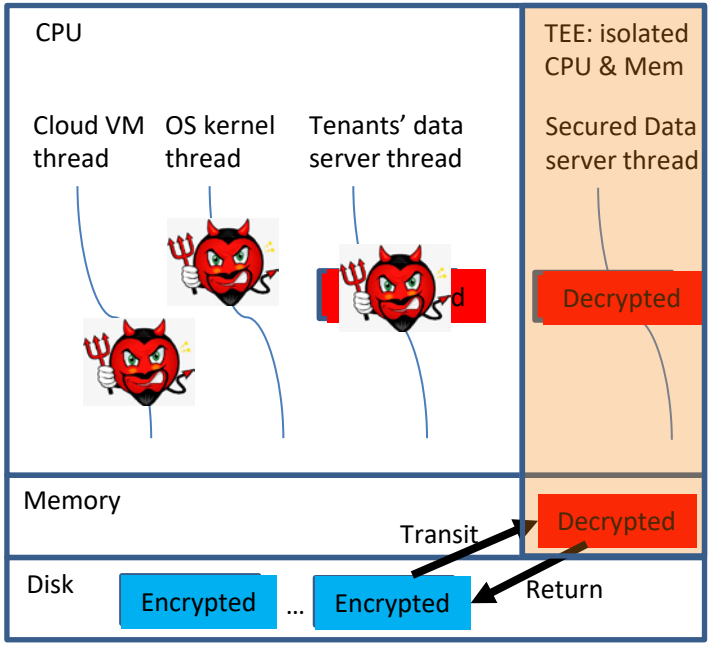

Ubiquitous **TEE**: A Synergistic Eco-system Between TEE & Java (UTEE's own components are in green)

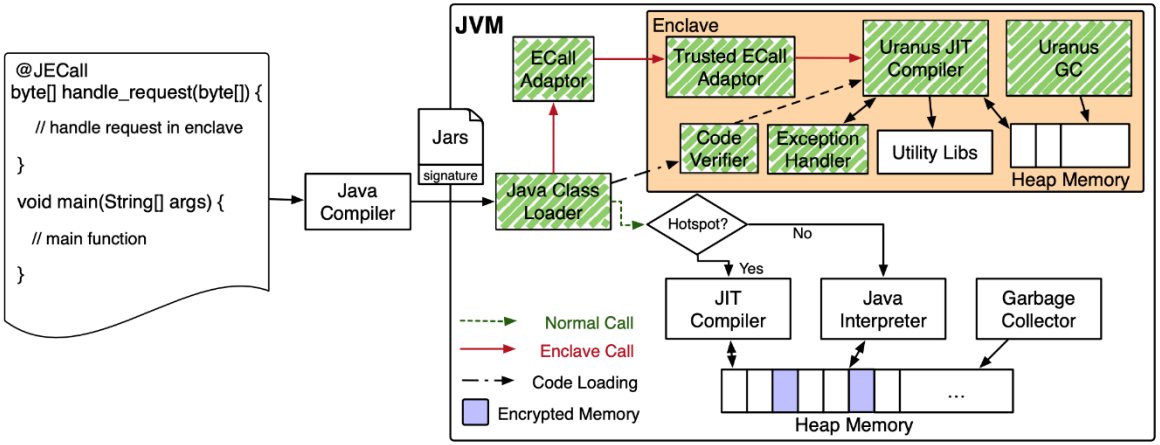

- For cloud software's JVM-compatible high-level languages (e.g., Java/Python/Scala)
- JECall: annotate queries which processing sensitive data, and UTEE automatically splits the queries and places them to run inside a TEE
  - Secure: **minimize the risk** (Lines of Code) of including software defects in TEE
  - Portable: Can **reuse the mature queries** (algorithms) developed by SQL/big-data experts
  - Easy-to-use: Java type safety, UTEE invents **new safe isolation techniques**: "queries ← → outside"
  - Fast: UTEE invents **new big-data query aware memory management**, first among competitors

Underpinning Research (Findings)

- **Easy to use**: UTEE is the world's first software system that allows big-data/AI experts to focuses on developing only the logic of their big-data/AI algorithms/functions using high-level programming languages (e.g., Java/Python)
- **Automatic security enforcement**: UTEE automatically deploys these algorithms/functions within high privileged but low-level CPU/GPU hardware and protect them from various attackers, including various admin software in clouds and malicious/careless cloud admin people
- **General**: UTEE can support various big-data/AI computing platforms (e.g., Spark/PyTorch/TVM) and various high privileged CPU/GPU hardware (e.g., Intel SGX and ARM TrustZone)
- **High-performance**: UTEE schedules to run only the big-data/AI algorithms/functions within high privileged CPU/GPU hardware, but not the entire big-data/AI platforms, greatly saving hardware resources and improving big-data/AI computing performance
- **Big-data/AI friendly**: UTEE develops a series of big-data/AI computing resource management protocols to greatly reduce/reuse the computing hardware usage, making **UTEE the world's first software system supporting typical large-scale big-data/AI computing workloads in the same security area**; worldwide competitors support only small-scale big-data/AI computing workloads

Engagement (Collaboration Workflow Between HKU and Huawei)

- Step 1: HKU and Huawei reserchers jointly identify research problems
- Step 2: The HKU team (Dr. Cui's PhD students) propose ideas, solutions, and new software architecture designs
- Step 3: The HKU team jointly applies for patents with Huawei researchers
- Step 4: The HKU team implements the solutions and designs into research prototype (software systems), and write research papers
- Step 5: The HKU team and Huawei researchers jointly submit the papers to top venues
- Step 6: The HKU team and Huawei researchers discuss with the Huawei software production teams' engineers (different from the Huawei researchers) for suggestions and commercialization opportunities of the research prototype, including jointly identifying potential VIP customers (e.g., governments, banks, and hospitals) on Huawei Clouds
- Step 7: The HKU team refines/enriches the research prototype developed by HKU and deliver the refined/enriched protype to Huawei engineers
- Step 8: The refined/enriched protype go throughs testing and is released as commercial software (e.g., https://www.huaweicloud.com/product/tics.html)
- Step 9: if the research papers are published, Dr. Cui's papers include open-source software links (https://github.com/hku-systems, managed by Dr. Cui) of the research prototypes (sometimes the versions are older than the ones delivered to Huawei), so that world-wide researchers and industries can reproduce the papers' results and extend the prototype for their own purposes

Engagement Partners

- Huawei Researchers
  - E.g., the Huawei Hong Kong Research Institute in the Shatin Science Park, and the Huawei Theory Research Department worldwide
- Researchers in the **HKU Medical School**
  - E.g., Dr. Teng Zhang https://www.aimed.hku.hk/teng-zhang
  - Dr. Cui's team have discussed several rounds with Dr. Zhang's team on applying UTEE to protect the big-data/AI computing on sensitive patents' data within both the Queens Mary Hospital and in HKU Med computing clusters
  - The HKU CS (Dr. Cui) and HKU Med collaborations aim to short for top medical journal papers (e.g., The Lancet)

Impacts Achieved (Products/Awards)

- New software release (UTEE → Huawei TICS): Huawei Cloud is **world's top-5 cloud vendor**
  - Huawei claims that UTEE is **world-leading** (www.huaweicloud.com/product/tics.html)
  - Huawei's official acknowledgement letter confirms that UTEE is usable is **already usable by Huawei Cloud's 3 million customers (including enterprises and individuals) in 170 countries**
- "**The Best Collaborating Scientist Award Medal 2021**" of the Huawei Theory Research Department
- **ITF ITSP Platform grant** award (**HK $3.3M**, Huawei donated 10%), PI (Dr. Cui), grant title "**UTEE: A Secure, Efficient, and Portable Distributed Bigdata Computing System on Heterogeneous Trusted Execution Devices**", in May 2022
- Other relevant grant 1, PI, "ParaNAS: High-performance, Scalable, Reliable and High-precision Multi-GPU Pipeline Parallel DNN Training Systems", **HK $2.3 million, Huawei Theory Lab Flagship**, 2021 - 2023
- Other relevant grant 2, PI, "A Blockchain-powered, Trustworthy Internet Layer (System) and its Decentralized and Efficient Applications", **HK $2.2 million, Huawei Innovation Research Program (HIRP) Flagship**, 2018 - 2020. Finished, the deliverables received an outstanding (highest) score from Huawei
- Other relevant grant 3, PI, "New Systems and Algorithms for Preserving Big-data Privacy in Clouds", **HK $490K, Hong Kong RGC GRF** (Ref: HKU 17202318), 2019 - 2022
- Other relevant grant 4, PI, "Achieving Strong Fault-tolerance for General Storage Applications via Fast, RDMA-powered PAXOS", **HK $544K, Huawei Innovation Research Program (HIRP) Open**, 2017 - 2018. Finished, the deliverables received an outstanding (highest) score from Huawei
- Dr. Cui's [Kakute ACSAC 2017] paper about **big-data security** won the **best paper award** in the ACM ACSAC 2017 conference (an international top security conference held by the USA Applied Computer Security Association). https://www.acsac.org/archive/

Impacts Achieved (TICS on Media)

- **HDC.Cloud 2021**: “Huawei Releases Six Groundbreaking Products to Supercharge the Cloud and Intelligent Transformation of Business” (https://www.huawei.com/en/news/2021/4/hdc-cloud-2021-six-products-cloud-intelligent-transformation)
- **Global Times**: “Huawei launches six new products to bolster growth in cloud sector” (https://www.globaltimes.cn/page/202104/1222038.shtml)
- 中国日报: “华为云可信智能计算服务TICS，安全释放数据价值” (https://caijing.chinadaily.com.cn/a/202104/29/WS608a7f1ea3101e7ce974cf35.html)
- 中国证券报官方网站: “第三届上海金融科技国际论坛：强化数据安全 释放融合价值” (https://www.cs.com.cn/xwzx/hg/202112/t20211205_6225416.html)
- “12月18日，2020数据资产大会在北京召开。会上，中国信息通信研究院为通过基于可信执行环境的数据计算平台的产品颁发证书。华为云可信智能计算服务 TICS基于鲲鹏TrustZone机密计算，结合硬件TEE和软件SMPC算法，实现了软硬结合的计算加速，同时支持跨信任域的联邦SQL分析和联邦学习能力，以九大测试项全部通过，树立隐私计算产品的新标杆。” (https://icode.best/i/87602637140406)

Three Impactful Application Scenarios of the Proposed UTEE Project

- Fore more details, please look into the Application Scenarios ("应用场景") in TICS
  - https://www.huaweicloud.com/product/tics.html
- All the three application scenarios of UTEE are designed by Dr. Cui, adopted by Huawei, and included in the TICS product page above
  - See the next three pages for the three scenarios

UTEE App 1/3: An Enterprise Stores Its Users' Data on A Government's Cloud

- **The "云上贵州" Mode in China and Europe**
  - https://technode.com/2018/01/10/apple-icloud-guizhou/
- Application service providers (e.g., Apple) run big-data/AI application within a government cloud's UTEE instances
- Users submit application requests (e.g., fetch or submit data) executed in UTEE
- Each government supervises the cloud and enforces data access policies (e.g., Euro Union's GDPR Law)
- UTEE can prevent "2018 iCloud data leakage caused by insiders of the Guizhou Cloud"
  - www.anzhixun.com/news/201810/15094812.html

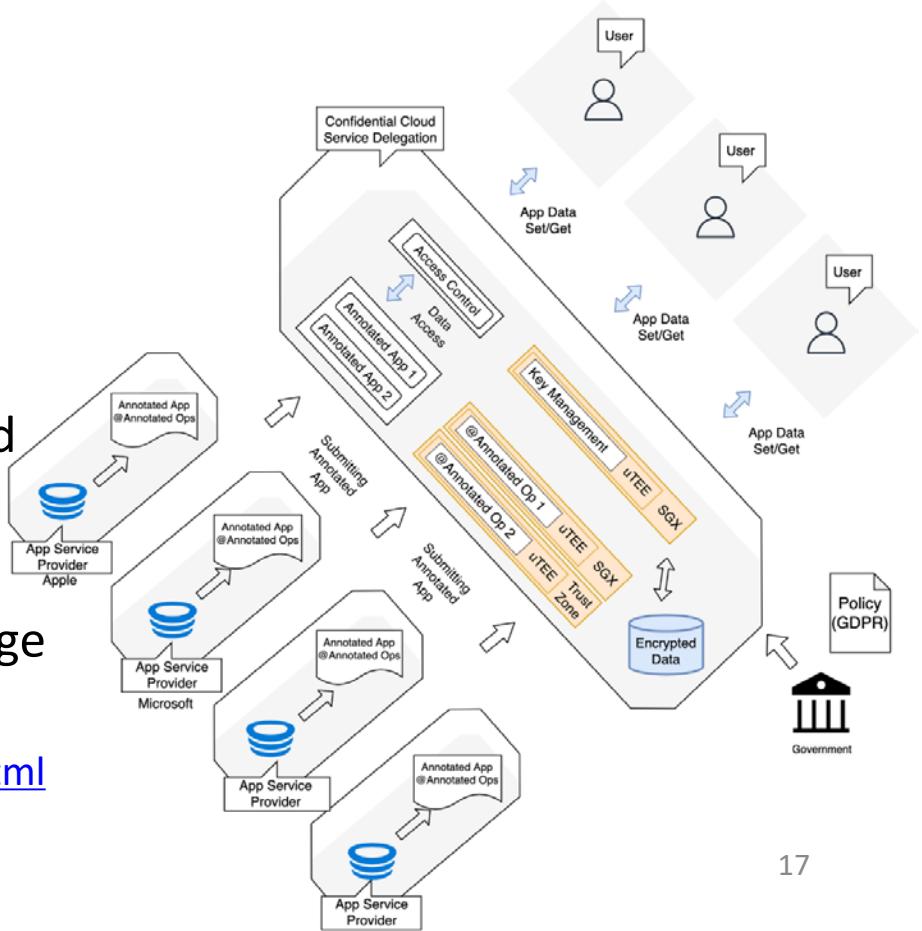

UTEE App 2/3: Enterprises and Governments Share Data with Others

- **The multi-party secure computing mode**
- Data owners (e.g., governments and big enterprises) provide sensitive data (e.g., citizen IDs, social security information and credit card data)
- Data users (e.g., small enterprises and researchers) submit data queries to query the aggregated data stored in the public cloud with security, enforced by UTEE

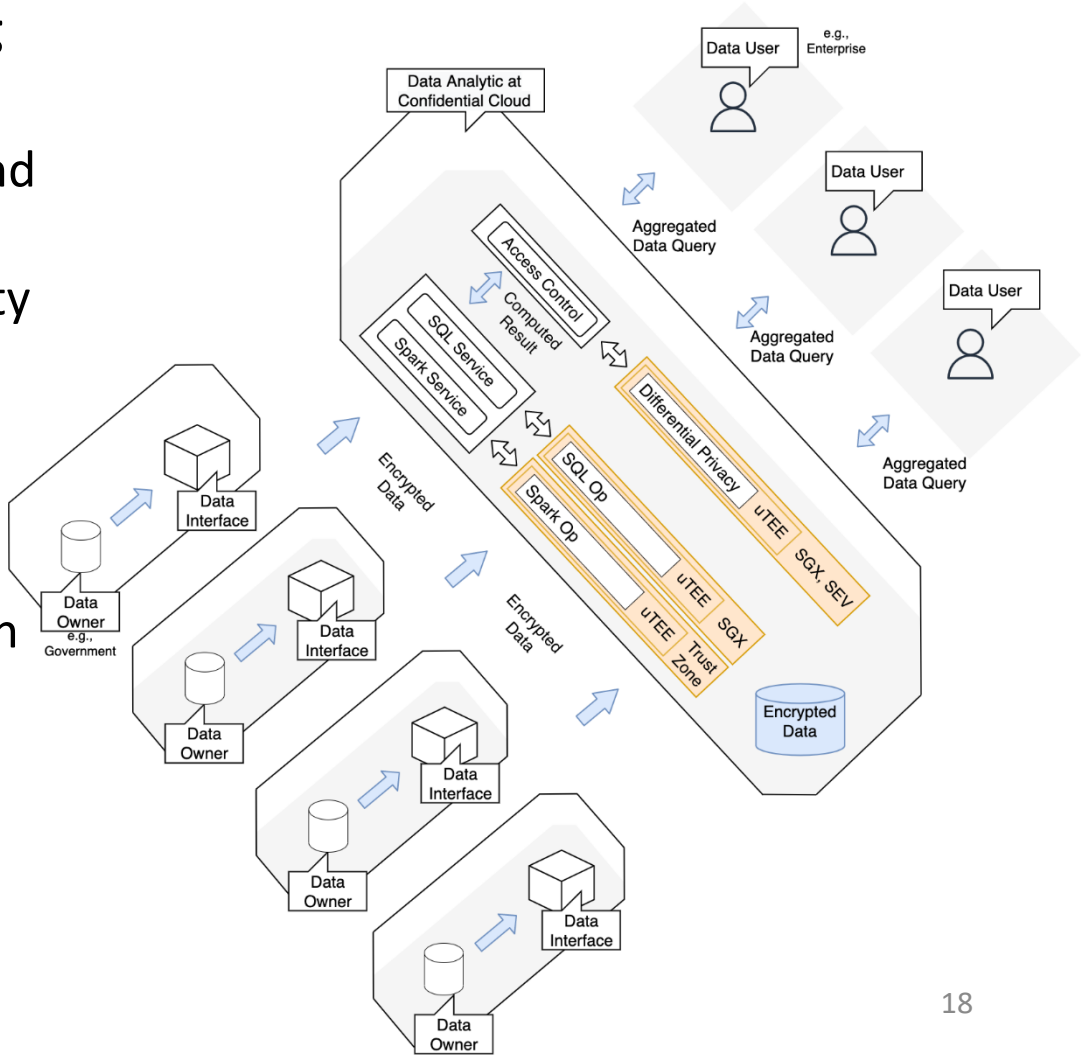

UTEE App 3/3: Users Securely Share Their Own Data on Edge and Cloud

- **The secure mobile computing mode**
- App service providers (e.g., Tik-Tok) provide an application with confidential and offline operators
- Users runs interactive applications (e.g., games or covid contact tracking) at edge devices and submits app requests to clouds
- All personal information are computed within UTEE securely in both clouds and edge devices