Whitepaper / Technical Report

Enemy in the clouds: protecting your cloud assets from powerful adversaries

Confidential Cloud-Native Computing in Large Kubernetes Clusters

T · · Systems ·

intel

SCONTAIN

Imprint
Whitepaper-Publication August 2020
T-Systems Multimedia Solutions GmbH
Riesaer Strasse 5, 01129 Dresden

Authors
Dominik Nägele, Dr. Ivan Gudymenko,
Dr. Christof Fetzer, Heqing Zhu, Kapil Sood,
James Greene

Organisation
Project management: Julia Kunert
Layout: Peter Brücker

Content

Management Summary

A demand for secure and trustworthy processing of data is rising in the market. The data should not only be processed in adherence with the General Data Protection Regulation (GDPR), but also remain under the full control of the data owner (data sovereignty). Customer data sovereignty should be protected at all times – even during and after successful cyberattacks on the underlying cloud infrastructure. Increasingly, companies should design security with a “zero trust” principle – thereby assuming threats to the data exist everywhere and that no entity should intrinsically be assumed trustworthy. Therefore, this principle should be valid both for external attacks and for attacks by insiders (such as administrators) with full software and hardware access. Such a guarantee can practically only be upheld by enabling the uninterrupted encryption and integrity protection of data, programs and all keys (Confidential Computing). This addresses two of the three classic IT security protection goals: Confidentiality (hence „confidential computing“) and integrity. High availability in this context could for example be accomplished by leveraging native Kubernetes capabilities.

The broad use and acceptance of Confidential Computing requires simplicity and low overhead. It should not impair the usability nor the development friendliness of the IT solution. In general, despite the proven benefits of enhanced security, many organizations only permit a slight increase in the cost for development and operation. Modern cloud-native applications use Containers, Kubernetes, Microservices, Continuous Integration/Continuous Deployment, DevOps and a combination of open-source and closed-source software developed in a variety of programming languages. Confidential Cloud-Native Computing uses the same methods, and additionally provides uninterrupted encryption without the need to change the source code of the application. The workloads running in the cloud can thus be secured transparently with minimal effort for developers and service providers.

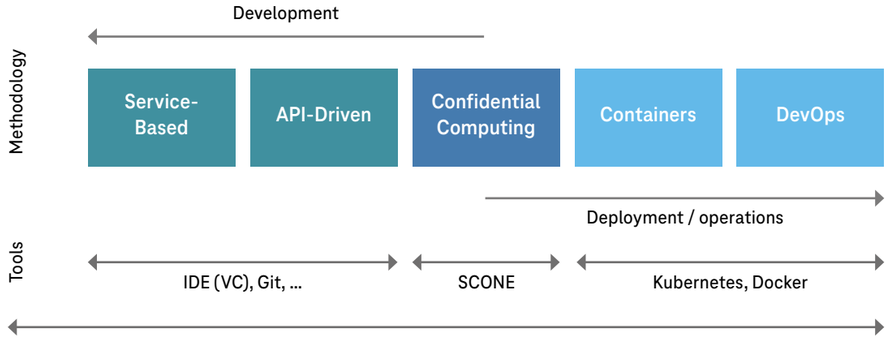

Figure 1: Development of Confidential Cloud-Native Applications

Over the last few months, T-Systems, Intel and Scontain have carried out a proof of concept (PoC) for Confidential Cloud Native Computing. The project investigated how the optimal scaling of a container-based microservice architecture can be realized while simultaneously encrypting data and programs during execution using Intel® Software Guard Extensions (Intel® SGX) which has the unique ability to isolate designated application code and data in-memory. Intel® SGX does this with hardware-based memory encryption, allowing application developers to partition their applications into CPU-hardened ‚enclaves‘ or encrypted areas of execution in memory. The SCONE platform developed by Scontain enhances the native capabilities of Docker and Kubernetes to provide the software platform. This enables the applications to be run in the cloud in a performant and secure way without modifications to their source code.

The aim of the PoC was to show how the operation of an Intel® SGX-based container platform can be efficiently scaled, while taking into account multiple different load profiles. For this purpose, load & performance tests with different usage profiles were used. This process helped to identify configurations on the most performant and resource-optimized scaling of a future operating platform.

The findings enable efficient and optimal use of the hardware with different load profiles. One of the key objectives was to maximize the number of parallel microservices per physical node while maintaining predefined performance characteristics.

For this purpose, a secure document management system was implemented on the T-Systems container platform, including system monitoring tools and scripts. Altogether, this provides reliable information about the stable and scalable behavior of the container platform (including document handling) for about 6000 parallel micro-services distributed over 75 physical nodes.

Background and problem definition

Our problem was to investigate the practicability of a realistic Confidential Cloud-Native application.

Confidential Cloud-Native Application Development	
Security	data, code, and keys are always encrypted - at rest, in transit, in memory -
Focus	Speed to market
Development Methodology	Agile developement, DevOps
Teams	Collaborative DevOps team
Delivery Cycle	Short and continuous
Application Architecture	Loosely coupled, service-based, API-based communication
Infrastructure	Container-centric, portable, scales horizontally, on-demand capacity
Enterprise-grade SLAs	ensuring end-to-end performance, stability and availability

Figure 2: Properties of Confidential Cloud-Native Applications

The application should scale horizontally with load requirements while supporting all aspects of cloud-native applications:

- Container-based deployment,
- continuous integration and continuous deployment,
- Micro-Service based,
- free choice of programming language for each Micro-Service,
- support of DevOps teams, and
- orchestration by means of Kubernetes.

Confidential Cloud-Native application also requires that all critical data, program code and keys remain encrypted without interruption:

- **at rest** – persistently stored on discs,
- **in transit** – during data transmission over the network, and
- **at runtime** – during data processing in main memory.

These first two points are quite well understood and widely implemented. The third point is FAR more challenging—especially with the growing use of cloud, virtualized or otherwise shared computing environments. The critical data, the code and the keys must therefore never be in plain text, or otherwise accessible by anyone outside of the data or application owner.

In today's world, personal information is increasingly classified as data requiring special and continuous protection. Against this background, the management of documents including their metadata was examined in this PoC. Possible applications of this PoC include the financial, insurance and health sectors.

We have therefore chosen a secure document management system as our example application. Clients can store their documents via an encrypted connection to the cloud (e.g. TLS). The data is stored encrypted in a database, and the metadata is also encrypted continuously

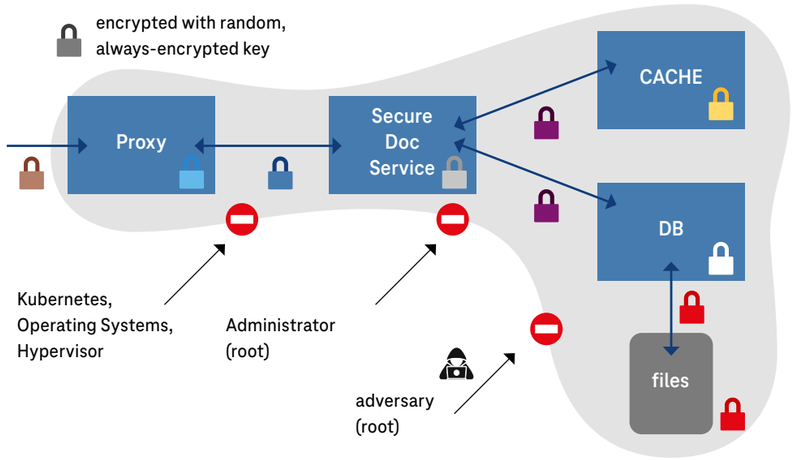

Figure 3: SCONE ensures that data, code, and keys are continuously encrypted

The PoC is based on the following technologies:

- **Intel® SGX:** An architecture extension residing in the hardware designed to increase the security of select application code and data, protecting it from disclosure or modification.
- **SCONE:** A platform that allows for transparent execution (i.e., without modification of the program code) in enclaves. Data is encrypted transparently for storage (at rest) and communication (in transit).
- **Docker:** Micro services are carried out in docker containers. The associated container images are decrypted transparently by SCONE during runtime.
- **Kubernetes:** The execution of the docker containers is orchestrated by Kubernetes. The PoC ran on a Kubernetes cluster with 90 nodes – of which 15 were used to generate the load, spinning up the 6000 microservices running on the other 75 nodes.
- **Ceph:** All data is stored in a distributed and fault-tolerant storage system.
- **Helm:** The simplified installation of the application is accomplished via Helm.
- **Grafana:** The resource utilization in the cluster is visualized in real time by Grafana.
- **Prometheus:** All metrics are stored in the Prometheus database.

Intel® SGX: technological overview

Intel® Software Guard Extensions (Intel® SGX) offers hardware-based memory encryption that isolates specific application code and data in memory. Intel® SGX allows user-level code to allocate private regions of memory, called enclaves, which are designed to be protected from processes running at higher privilege levels. Only Intel® SGX offers such a granular level of control and protection.

Intel® SGX can help to prevent attacks such as memory bus snooping, memory tampering and cold boot attacks against memory in RAM. It helps protect against software-based attacks even if the operating system, drivers, BIOS, Virtual Machine Manager are compromised.

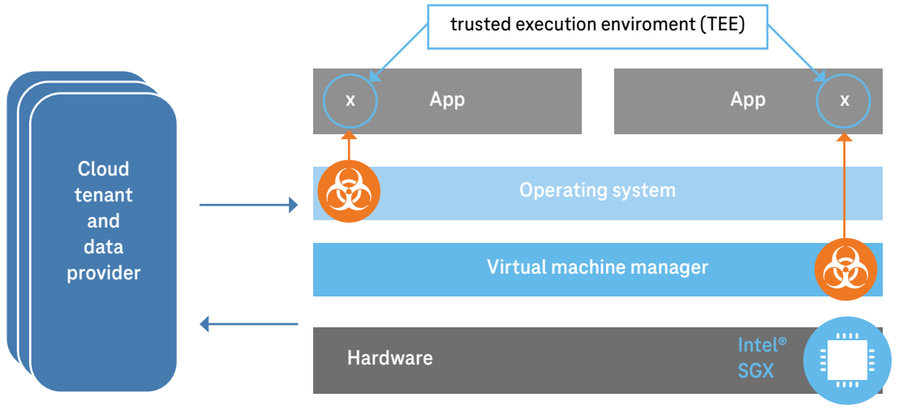

Figure 4: Intel® SGX: Trusted Execution Environment

Intel® SGX also provides an option for a hardware-based attestation capability to measure and verify valid code and data signatures. Remote attestation allows a remote provider (also known as a relying party) to have increased confidence that the software is running:

- Inside an enclave
- On a fully updated Intel® Software Guard Extension (Intel® SGX) system at the latest security level (also referred to as the trusted computing base [TCB] version)

Attestation is necessary in order to make remote access secure, since very often the enclave's contents may have to be accessed remotely, not from the same platform.

Security vulnerabilities can occur in today's extremely complex modern computer systems. These vulnerabilities are often mitigated through updated software or firmware from the hardware provider. Intel® SGX architecture is designed to recover from security vulnerabilities through a Trusted Computing Base (TCB) recovery and to re-establish trust in the recovered platform. It does this in two ways. First is the update of the mutable components of the TCB, completed by issuing a CPU firmware update and/or a new version of the attestation software that forms part of the Intel® SGX platform software (PSW) component. The second is producing a new TCB key set that identifies the TCB for a specific platform, which is then used in re-provisioning the platform's Intel® SGX attestation key.

As with any security architecture new attacks on both software and hardware and especially on Intel® SGX itself are possible. The advantage of the SCONE platform is that all vulnerabilities and mitigation mechanisms can be monitored. The transparent attestation of software and hardware helps to ensure that only up-to-date software is executed on up-to-date hardware. In case of vulnerabilities that cannot yet be sufficiently mitigated by hardware or software updates, additional measures in the SCONE platform are taken to mitigate these vulnerabilities. Only if automatic migitation is not possible, will DevOps need to activate system-level mitigation mechanisms.

While there are several methods for enabling VM-level encryption, they are not sufficient for our use cases for two reasons. First, we deploy the application using containers (and not VMs) and need to isolate the application from the administrators who have access to the VMs and the host. Secondly, we also need to isolate the individual containers – both from containers of other applications and clients and from containers of the same application. This helps to ensure that an attacker who breaks into one container does not have automatic access to the data in the other containers. A VM-based encryption model is useful for some scenarios but requires trust in the hypervisor (including its host operating system). If this is compromised, the entire system is at risk. This enhanced security-related aspect has made the use of Intel® SGX for cloud-native confidential computing particularly interesting.

Implementation and architecture

For the scaling tests, a client was implemented to generate the load on the container platform. This called various REST and SOAP commands on the Intel® SGX secured document management system. Execution was completely automated in order to more easily examine and visualize different application scenarios. In addition to the performance optimization for scaling tests, this enabled the precise prediction of resource requirements as the number of clients increased. Furthermore, the ideal configuration parameters were identified to minimize the required computer resources and response times and maximize throughput.

A transparent online monitoring of cluster resources has been developed to identify and solve performance problems and resource bottlenecks as quickly as possible. Several optimizations have been implemented that enable efficient use of Intel® SGX enclaves in the cluster.

The program code of the application services can be executed in enclaves without modification to source code. SCONE is integrated into the CI/CD (Continuous Integration/Continuous Delivery) process for this purpose. The SCONE cross-compiler compiles the applications so that they are executed in enclaves. This approach enables support for all popular programming languages. Key management is carried out transparently.

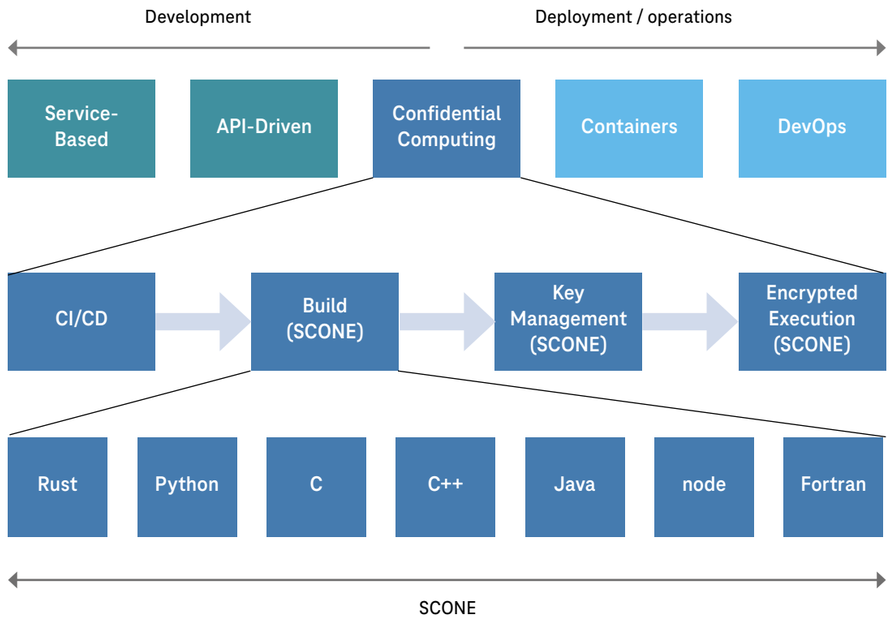

Figure 5: SCONE: Transformation of the development process for services in SGX enclaves

Results

A horizontal scaling of the secure document management system for up to 75 computer nodes was successfully demonstrated. 15 additional computers of the Kubernetes cluster were used for the execution of 6000 parallel clients (load generation). For the secure isolation of documents, a separate document server was started for each client. This means that up to 6000 document services were executed in parallel. Service response times were below 100ms on average, and the startup times of the document services remained well below 1 second.

Even under extreme load, response times remained stable and far below the required maximum values. Furthermore, the system can be scaled elastically to dynamically adapt to changing load requirements. It adapts not only to the external load changes, but also to the changes on the individual computer nodes to guarantee the specified SLAs (Service Level Agreements).

Figure 6: Secure document management dashboard

Conclusions

This PoC implementation has shown that practical use of Confidential Cloud-Native Applications is already possible today, even for applications with high requirements in terms of throughput and response times.

- **Throughput:** number of requests that can be processed per second, and
- **response times:** the time needed to answer a request even during periods of high throughput/load.

Applications executed in Intel® SGX-enabled containers do not have to be rewritten, and all modern tools used for the development and operation of cloud-native applications can be still used.

Further application possibilities: Data protection in the Blockchain environment

Though not part of the previously described test scenario, blockchain-based systems can also benefit from Confidential Computing technology. The implementation of data protection requirements in the blockchain environment remains a challenge. This is one of the major hurdles that such systems will need to be able to overcome on the way to production and operational readiness, as well as to their acceptance in society. Confidential Cloud-Native Computing allows end-to-end (E2E) encryption of data and workloads in the blockchain system, helping prevent any unauthorized data access.

As part of the scientific cooperation, T-Systems and Scontain have already developed and practically evaluated another PoC addressing this issue. It attempts to implement the secure and trustworthy execution of blockchain applications (smart contracts in the form of fabric chaincode) in a consortium blockchain network based on Hyperledger fabric technology. Within the PoC, the use case addressing the producer-consumer model in a smart energy e-car environment was developed. The users of such system can thus rent out free capacities at their charging stations at home, and receive tokens in return. The latter can be further used at foreign charging stations in the ecosystem to enable the charging process. The customer data such as the wallet credit, as well as the corresponding chaincode (smart contract), which reflects the business logic, are protected by the Intel® SGX enclaves in terms of confidentiality and integrity.

Summary

Confidential cloud-native computing is one crucial way to bring new trust models to cloud solutions and providers to help meet growing security requirements. The ability to encrypt data end to end throughout its entire lifecycle reduces the threshold to enable more workloads to gain the benefits of cloud deployment models.

A further use case for this concept can be found in key management, for example. Here, Intel® SGX is used to manage cryptographic keys and provide HSM-like (HSM: Hardware Security Module) functionality for the workloads. This would provide significant cost and complexity savings compared to traditional hardware-enforced key protection models. In the field of edge computing, the concept can be used to secure the computation in IoT edge devices and their communication with the cloud.

Also, in the field of Artificial Intelligence, the SCONE platform can be used by various stakeholders to secure their respective intellectual property (IP). SCONE can help protect the entire Python code, AI or Machine Learning models, training data and operational data from unauthorized access using Intel® SGX. This enables collaboration in new areas where sensitive code and training or operational data are owned by different stakeholders. Confidential Computing thus creates a basis of trust for new collaboration models in the field of artificial intelligence.

Authors

Dominik Nägele
T-Systems International GmbH

Dominik Nägele (M.Sc.) has been working in various positions in the IT sector for more than 7 years. As an Enterprise Architect in pre-sales, he is currently responsible for complex ICT projects, mostly in the automotive, insurance, banking, retail and public sectors. He supports companies in their digital transformation by leveraging Cloud Services, Big Data, DevOps and Confidential Computing.
dominik.naegele@t-systems.com

Dr. Ivan Gudymenko
T-Systems Multimedia Solutions GmbH

Having a technical background in IT security, distributed systems and telecommunications, and after finishing his PhD in technical privacy preservation in the area of NFC/RFID-based e-ticketing systems, Dr. Ivan Gudymenko joined T-Systems Multimedia Solutions to support the company as an IT security architect and technical consultant. Since then, he has been involved in various projects in IT security and blockchain-based systems. These include the innovation projects where he combined the experience gained from classic IT security field and technical project management with innovation-driven consulting to facilitate the process of raising new technologies to the level of enterprise-ready systems.
Ivan.Gudymenko@t-systems.com

Dr. Christof Fetzer
Scontain UG, scontain.com, community: sconedocs.github.io

Prof. Fetzer is the COO of Scontain and co-founder of Scontain UG, Cloud&Heat Ltd. and SIListra Systems Ltd. Since 2004, he has been teaching at the TU Dresden, Germany, after receiving his Ph.D. at the University of California, San Diego. His research and developmental work focusses on trusted execution and cloud computing.
christof.fetzer@scontain.com

Heqing Zhu
Intel Corporation

Heqing is leading the network and security platform strategy and solution definition. In his 15 years' service at Intel, he managed open source projects such as DPDK, Hyperscan and network virtualization and container data plane acceleration. He authored the book Head First, DPDK in China.
heqing.zhu@intel.com

Kapil Sood
Intel Corporation

Kapil Sood is Platform Security Architect at Intel's Data Center Group, driving security strategy, technologies and research, and setting product security direction for Intel's NFV, 5G & Networking BU. Kapil has 25+ years of technology leadership experience, spanning telecom, mobile, networking, and cloud.
Kapil.sood@intel.com

James Greene
Intel Corporation

Jim is the Market Development lead for the Communications Services Providers market segment, focused on data center modernization, cloud security, service assurance and data plane processing technologies and solutions. Jim is co-author of the book Intel® Trusted Execution Technology for Servers: A Guide to More Secure Datacenters
james.j.greene@intel.com

Paul O'Neill
Intel Corporation

Paul is Senior Director of Strategic Business Development in Intel's Confidential Computing Group. He is responsible for eco-system development and coordination of Confidential Computing solutions like Intel® SGX and strategies across market segments from Enterprise, Government, Cloud, Healthcare and IoT with strategic partners. Paul has been working at Intel since 2014 and has 20+ years' experience in technology companies from start-ups to large enterprises.
paul.oneill@intel.com

Guided. Digital.
About T-Systems Multimedia Solutions

T-Systems Multimedia Solutions facilitates the digital transformation of large corporations and medium-sized companies. With an annual 2019 turnover of € 176 million, it enables its customers to develop new digital business models for Industrial IoT, Customer Experience, New Work and Digital Reliability. Leveraging its consulting and technical expertise with some 2,100 employees at seven locations, the digital service provider also offers dynamic web and application management. As market leader it offers the first certified test laboratory for the internet and multimedia industry, delivering the highest standards of software quality, accessibility and IT security.

T-Systems Multimedia Solutions has been honored several times with the “Social Business Leader Award” backed by the Experton Group, as well as the “iF Design Award” and has been 2017 a winner of the Outstanding Security Performance Awards. Headquartered in Dresden, the company has also received the “Great Place to Work Award” several times as one of Germany’s best employers and was awarded as “Bester Berater 2018” (best consultant) by the german business journal “brand eins“.

For further information: www.t-systems-mms.com

Would you like personal contact with our experts?
Please direct your inquiry:

Julia Kunert
Marketing Manager
E-mail: Julia.Kunert@t-systems.com

www.t-systems-mms.com