Solution Brief

5G Communications Service Providers
Confidential Computing

Confidential Computing for 5G Networks

The solution ecosystem for 5G network security includes the Fortanix confidential computing environment, which enables network operators to run code in Intel® SGX enclaves, deployed in containers using Red Hat® OpenShift®, without tailored development.

As communication service providers (CoSPs) refine their strategic courses for the 5G era, migrating to cloud-native architecture offers significant advantages. An architecture based on cloud-native network functions (CNFs)—developed as microservices and deployed in containers—dramatically increases the environment's flexibility, agility, and scalability. The ability to reconfigure the network on demand means that new services can be brought to market rapidly, and because changes are enacted in granular, modular microservices, discrete changes can be rolled back with minimal disruption to increase resilience and reduce business risk.

The shift to cloud-native comes at a time when CoSPs must manage multiple simultaneous transitions. In place of the single-vendor proprietary interfaces that predominated in previous technology generations, a service-based architecture is the emerging standard, combining network functions from multiple vendors. Each network function exposes a service-based interface that enables components to be interconnected using Web APIs in a software-defined infrastructure that can be recomposed at will. Rather than being centralized, as in previous network generations, 5G network functions are highly distributed, from the network core to the edge, across multiple clouds.

The Rising Importance of Confidential Computing in a 5G World

These transitions dramatically change the security landscape for 5G, as there are more vulnerable areas. A few examples of the enlarged attack surface and novel challenges for protecting data in use are shown in Figure 1. The traditional network perimeter has vanished, and data must be protected while in use at the level of the individual network function. Confidential computing protects data that would otherwise be exposed in RAM while a workload is operating, complementing protections for data at rest and in transit, which are provided by data and transport-channel encryption, respectively.

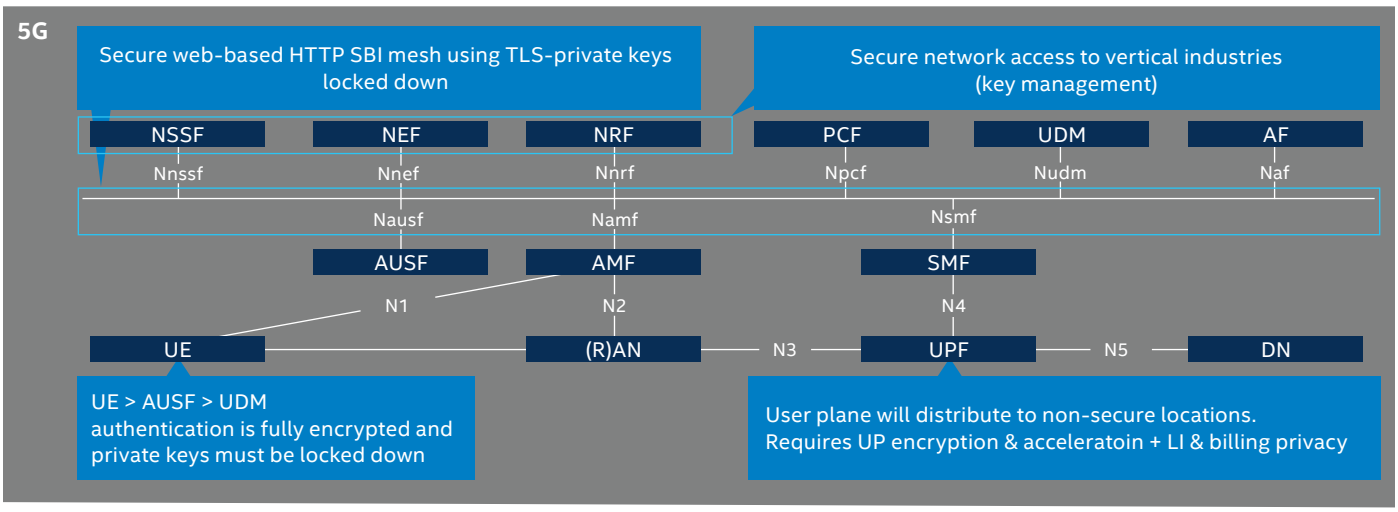

Figure 1. Increased 5G attack surface and challenges for protecting data in use.

Intel® Platform Support for 5G Use Cases

Together with other hardware security measures such as crypto acceleration, Intel® SGX enables a range of 5G use cases, including the following:

- Confidential communication for service-based architecture
- Secure authentication and secure converged edge
- Secure federated multitenant content delivery
- Secure communication edge-to-core and between networks
- Assured integrity of billing information with an audit trail

Confidential computing is chiefly concerned with isolating data from other workloads operating in their shared multi-tenant, multicloud infrastructure, preventing both inadvertent and malicious leaks. A prevalent approach is to provide an isolated trusted execution environment (TEE), where code can be run beyond the reach of outside software. CoSPs use TEEs to help meet challenges associated with protecting data in use by highly distributed 5G services in functional areas such as the following:

- **Service-based architecture** web-based integration of network functions, protected using transport layer security (TLS)
- **Key management for secure network access** extended to industry partners and customers
- **Authentication** with the Authentication Server Function (AUSF) and unified data management (UDM)
- **Distributed user plane**, including to non-secure physical locations

In addition to isolation, TEEs provide assurance to outside entities that the code running inside them has not been corrupted or tampered with through the process of attestation. By attesting to the authenticity of its workload, a TEE can therefore give the CoSP's remote partners and customers confidence that the destination they are connecting to is legitimate and trustworthy, as well as being protected from interception by third parties.

Protected Enclaves with Intel® Software Guard Extensions (Intel® SGX)

Intel SGX is a set of processor instructions supported by 3rd Gen Intel® Xeon® Scalable processors that is used for partitioning a region of memory as a secure enclave that isolates and protects specific data and code, as shown in Figure 2. The contents of the enclave are encrypted using a cryptographic key generated within the CPU, and likewise, the data and code are only decrypted inside the CPU. This arrangement protects data and code from compromise within RAM or in transit between the CPU and RAM.

The enclave prevents data from being leaked to other software running on the same platform, including other workloads as well as processes running at high privilege levels, such as the OS, hypervisor, BIOS, or firmware. The use of enclaves therefore dramatically reduces the attack surface, as shown in Figure 3.

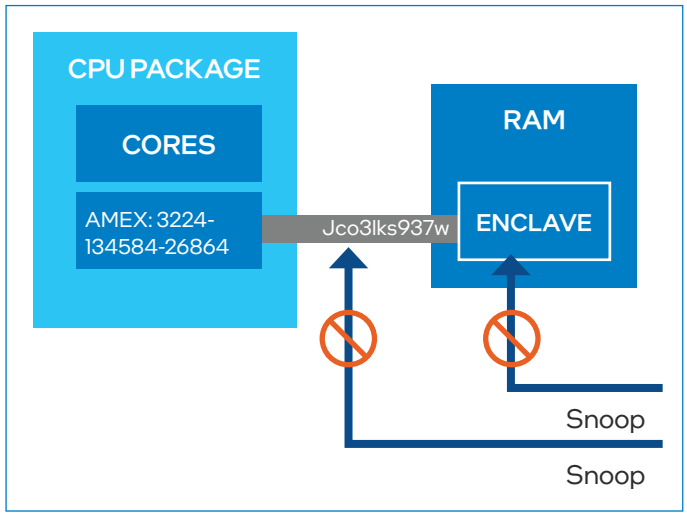

Figure 2. Data and code protection in Intel® SGX enclaves.

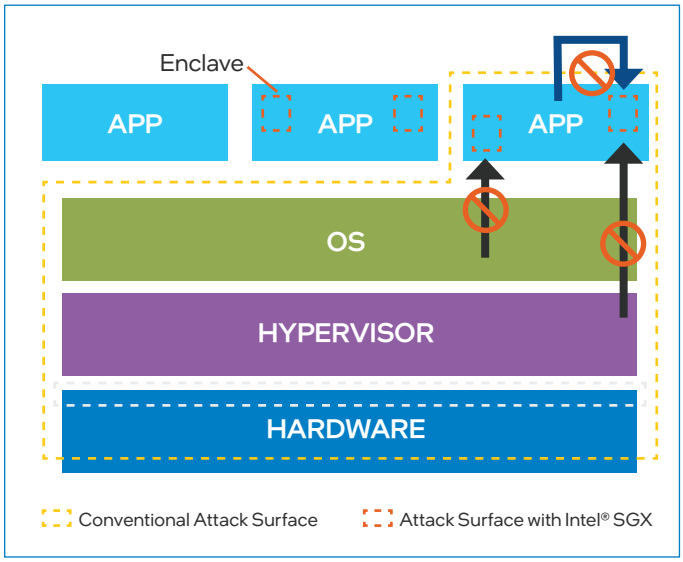

Figure 3. Reduced attack surface.

The native method of adapting software for use with Intel SGX is for the developer to identify the sensitive portion of the application—designated the trusted portion—and define an enclave for it to operate within by using the Intel SGX SDK. The Intel SGX enclave is deployed as a shared library. The remaining, untrusted portion of the application operates outside the enclave, and Intel SGX treats it as compromised by default. The trusted and untrusted portions of the application communicate with each other using specific Intel SGX instructions. The Intel® Attestation Service plays a vital role in establishing and maintaining the trusted status of code by enabling applications to verify the following factors cryptographically:

- **The code is running as-built** in a genuine enclave
- **The hardware is a secure Intel SGX-capable platform** with all needed microcode updates applied
- **All necessary Intel SGX hardware and software configurations** are made correctly

Fortanix solutions help streamline the implementation of confidential computing based on Intel SGX for 5G applications, including efficient attestation for 5G virtualized network functions (VNFs) at scale.

Platform Capabilities that Complement Intel® SGX

3rd Gen Intel® Xeon® Scalable processors incorporate multiple hardware-resident security features that work in conjunction with Intel® SGX. The following features are of particular interest to CoSPs as they deploy 5G network functions:

- **Built-in crypto acceleration.** To help CoSPs handle the performance impact of pervasive encryption in 5G, platform results include up to 4.2x higher TLS encrypted connections per second.[1]
- **Intel® Platform Resilience.** To protect fundamental platform firmware components, this Intel® FPGA-based solution establishes a chain of trust and verifies firmware images before execution.

Confidential Computing with Fortanix Solutions

Fortanix technologies implement Intel SGX across a range of confidential computing functions, as illustrated in Figure 4. The Fortanix Data Security Manager implements Intel SGX in its key-management service (KMS), which provides secure generation, storage, and use of cryptographic keys, certificates, and secrets. Runtime Encryption Technology provides a comprehensive environment for developing, operating, and maintaining Intel SGX enclaves.

Enclave OS is the runtime environment for code to run inside enclaves, based on simply repackaging existing images without requiring any changes to application binaries. The ability to run existing software without modification dramatically reduces the time, cost, and complexity associated with deploying it in a confidential computing environment. Enclave OS operates a root of trust established in the CPU to create a region of memory that is inaccessible to any process outside the application itself, regardless of privilege level. The CPU autonomously generates the key that it uses to encrypt this memory on the fly, using a secret provisioned at the time of manufacture, so the key is never exposed outside the CPU itself.

The Confidential Computing Manager is a cloud-native SaaS environment that provides a single pane of glass for managing secure enclaves and confidential computing nodes, on-prem or in any cloud or hosted environment (on nodes that support Intel SGX). It controls the enclave lifecycle, including enablement for policy enforcement measures on running applications. For example, whitelisting allows applications to run only on a predefined list of hosts, providing added assurances of integrity and protection. Geofencing restricts the geographic area within which an application can run in a multicloud environment, for regulatory and audit purposes. The Confidential Computing Manager also performs highly efficient attestation services, with minimal burden placed on developers and network operators.

The Enclave Development Platform (EDP) is an open source environment for writing Intel SGX enclaves from scratch, using the Rust programming language. The design of the EDP is optimized by years of in-house use by Fortanix to develop various products, making it exceptionally efficient and developer-friendly. Rust combines high computational performance with built-in code safety measures, especially for safe concurrency and memory safety. The Rust compiler

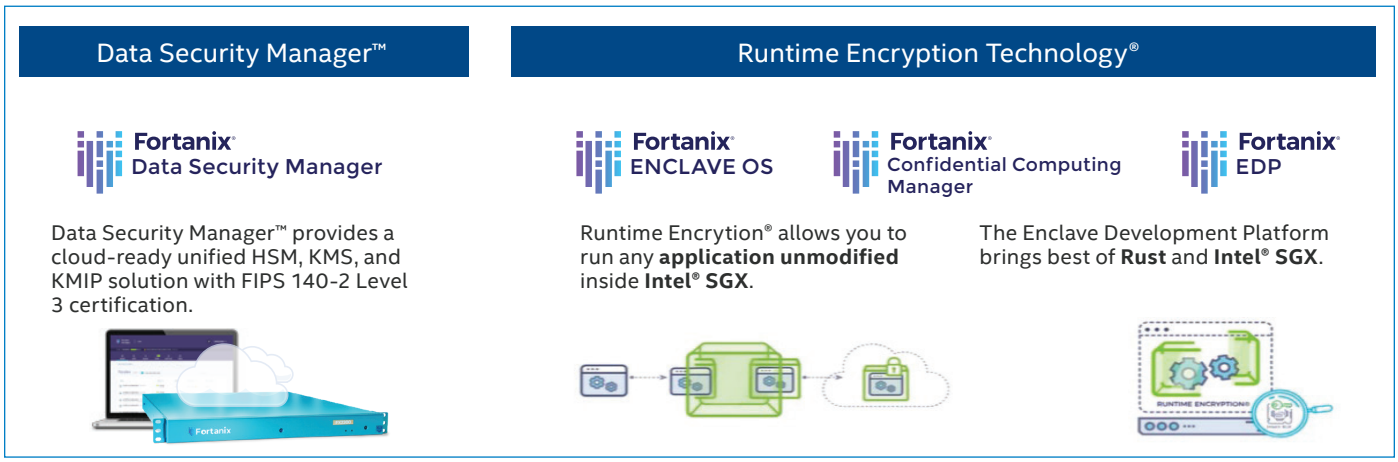

Figure 4. Fortanix technology portfolio for confidential computing.

allows for code to be compiled specifically for Intel SGX and incorporates static code analysis to provide another level of verification for code safety.

The Fortanix Node Agent software runs on physical or virtual compute nodes, as an intermediary between unmodified applications and Fortanix services. It enables compute nodes to register with the Confidential Computing Manager and manages the workloads operating in secure enclaves. The Node Agent is instrumental in the establishment of trusted compute pools, validating compute-node hardware and platform software. It also assists with application attestation and visibility for the Confidential Computing Manager. Fortanix is working with the cloud ecosystem to enable deployment of the Node Agent through a variety of platforms, including Microsoft Azure and Red Hat® OpenShift® Container Platform.

Automated Deployment on Red Hat OpenShift

Red Hat OpenShift is an enterprise Kubernetes platform, shown in Figure 5, that automates management functions for hybrid cloud, multicloud, and edge deployments, helping CoSPs increase operational efficiency. It incorporates a Linux OS and container runtime, hardened by Red Hat security engineers, as well as networking, monitoring, registry, and authentication and authorization components.

OpenShift encapsulates discrete capabilities as Operators, which are software entities that automate and accelerate management tasks. Functionally, OpenShift Operators are custom, application-specific Kubernetes controllers running on the Kubernetes master nodes that commonly fill roles such as dynamic configuration and tuning of compute and network parameters.

Red Hat sustains the open source Operator Framework project to provide tooling and support to the Operator ecosystem. This includes the Operator SDK, which abstracts away Kubernetes API complexities from developers, as well as the Operator Lifecycle Manager, which oversees the lifecycles of all Operators on a Kubernetes cluster, including installation, configuration, and updates. Red Hat also validates the functionality and soundness of Operators on OpenShift, providing developers and cluster administrators with a library of workloads as a service in the form of Red Hat OpenShift Certified Operators.

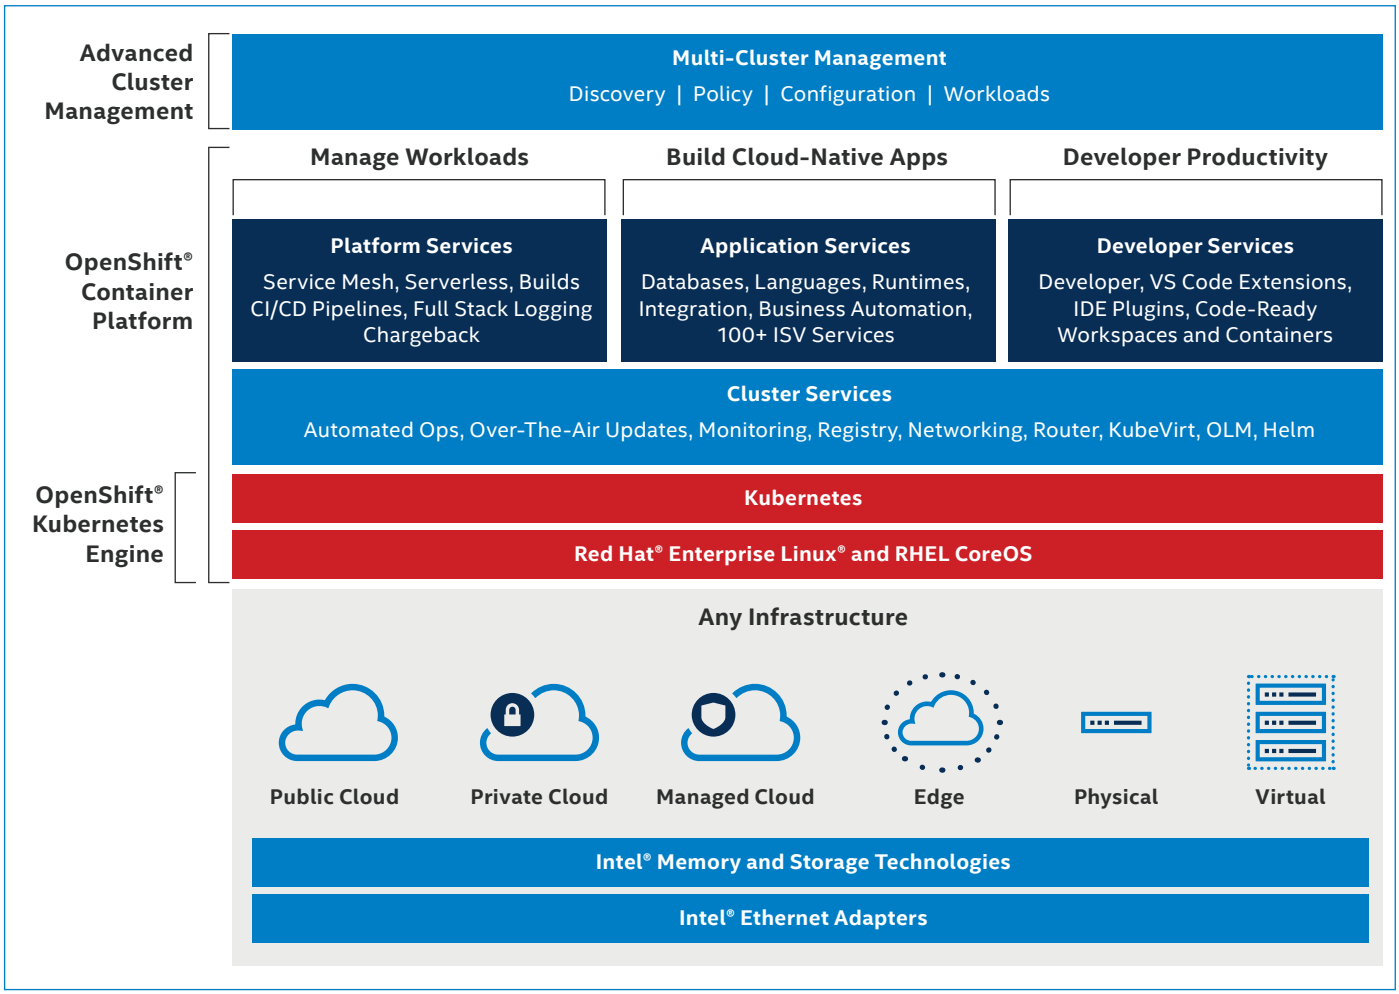

Figure 5. CoSP implementation of Red Hat® OpenShift®.

Casa Systems: Deploying 5G Services with Fortanix on OpenShift®

Intel, Fortanix, and Red Hat are working with Casa Systems to enable 5G network services—including the 5G SA Core—to be deployed on OpenShift® using the Fortanix Node Agent Operator. This proof of concept solution provides a blueprint for the industry to operate CNFs on OpenShift® that benefit from protection by Intel SGX secure enclaves.

The Fortanix Confidential Computing Manager Node Agent is now offered as a Red Hat OpenShift Certified Operator through the Red Hat Embedded Operator Hub, which is included in Red Hat OpenShift. This simple deployment path allows for quick installation and helps streamline maintenance. The Operator is a key contributor to the ability to easily provide confidential computing for 5G workloads.

Conclusion

In addition to the need to protect data at rest and in flight, the 5G Standalone (SA) core creates new imperatives for the protection of data while in use, from the edge to the network core. Confidential computing rises to this challenge, with Intel SGX providing secure enclaves that define isolated memory space for data and executing code. Fortanix enables existing applications to operate in secure enclaves without modification, making the adoption and utilization of Intel SGX significantly easier.

To help streamline the process of packaging, deploying, and managing applications that make use of Intel SGX, Fortanix offers the Confidential Computing Manager as a Red Hat OpenShift Operator. This combination of technologies makes it easy for CoSPs to deploy CNFs that execute using Intel SGX secure enclaves, without added development burden.

More Information

Intel® SGX: **intel.com/content/www/us/en/architecture-and-technology/software-guard-extensions.html**

Fortanix Confidential Computing: **fortanix.com/solutions/use-case/confidential-computing/**

Red Hat® OpenShift®: **redhat.com/en/technologies/cloud-computing/openshift**

Fortanix Red Hat® Operator: **catalog.redhat.com/software/operators/detail/60d6590772a23a5230f9d1a2**

Casa Systems: **builders.intel.com/docs/networkbuilders/casa-systems-framework-provides-private-5g-network-functionality.pdf**

Solution provided by:

intel. Fortanix® Red Hat

[1] See [70], [90], [71], and [69] at 3rd Generation Intel® Xeon® Scalable Processors - 1 - ID:615781 |Performance Index. Testing by Intel as of August 4, 2020. Performance comparisons relative to 2nd Gen Intel® Xeon® Scalable processors using a single buffer algorithm versus multi-buffer algorithms for 3rd Gen Intel Xeon Scalable processors. Results have been estimated based on pre-production tests at iso core count and frequency as of August 2020. Performance gains are shown for individual cryptographic algorithms.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for configuration details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a nonexclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

1021/RKM/MESH/346428-001US