D2.3

Final Report on Requirements and Trustworthiness

Funded by
the European Union

© COREnext 2023-2025

Revision v1.0

Work package	WP2
Tasks	T2.1-2.3
Dissemination level	PU – Public
Deliverable type	R – Document, report (excluding periodic and final reports)
Due date	31-12-2024
Submission date	18-12-2024
Deliverable lead	WINGS
Version	v1.0
Authors	Andreas Georgakopoulos (WINGS), Mohand Achouche (IIIV), Michael Roitzsch (BI), Nils Asmussen (BI), Markus Ulbricht (IHP), Arantxa Echarte (AUS), Gian Michele Dell'Aera (TIM), Julien Lallet (NNF), Renaud Santoro (NNF), Mamoun Guenach (IMEC), Björn Debaillie (IMEC), Enrico Guarino (TIM) Patrick Pype (NXP), Viktor Razilov (TUD)
Contributors	Work package partners (see below)
Reviewers	Mamoun Guenach (IMEC), Mohand Achouche (IIIV)

Abstract

This report provides the final outcome of WP2 and uses inputs from all tasks. This document highlights additional updates compared to previous WP2 deliverables (D2.1 and D2.2) to requirements, further analyses trustworthiness requirements for both digital and analog components, and the underlying lab validations to validate the architecture's robustness against emerging threats.

Keywords

Requirements; Architecture; Threats; Trustworthy Digital Components; Trustworthy Analog Components; Lab validations.

Document Revision History

Version	Date	Description of change	Contributor(s)
v0.1	07-10-2024	ToC	WINGS
V0.2	15-11-2024	First edited version ready	WINGS and partners
V0.3	25-11-2024	Second edited version for review	WINGS and partners
V0.4	10-12-2024	Updated version following review by IMEC and IIIV	WINGS and partners
v1.0	17-12-2024	Final version	WINGS and partners

Contributing Partners

Abbreviation	Company name
BI	BARKHAUSEN INSTITUT
AUS	AUSTRALO
CHAL	CHALMERS TEKNISKA HOGSKOLA
CEA	COMMISSARIAT AL ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES
EAB	ERICSSON
ETZ	ETH ZURICH
CYB	CYBERUS TECHNOLOGY
EUR	EURECOM
IFAG	INFINEON TECHNOLOGIES AG
IMEC	INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM
NXP	NXP SEMICONDUCTORS
RAD	RADIALL
SEQ	SEQUANS
TUD	TECHNISCHE UNIVERSITAET DRESDEN
TIM	TELECOM ITALIA
WINGS	WINGS ICT SOLUTIONS
IMS	INSTITUT POLYTECHNIQUE DE BORDEAUX
ETHZ	EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH
IHP	IHP MICROELECTRONICS
NOK	NOKIA NETWORKS GERMANY
NNF	NOKIA NETWORKS FRANCE
IIIV	NNF/IIIV LABS
IFAT	INFINEON TECHNOLOGIES
KAL	KALRAY

Disclaimer

The information, documentation and figures available in this deliverable are provided by the COREnext project's consortium under EC grant agreement **101092598** and do not necessarily reflect the views of the European Commission. The European Commission is not liable for any use that may be made of the information contained herein.

Copyright Notice

©COREnext 2023-2025

Executive Summary

The COREnext project aims to establish a cutting-edge 6G hardware platform that prioritizes trustworthiness and energy efficiency. By addressing critical 6G requirements, including ultra-low latency, high-capacity connectivity, and secure execution environments, COREnext lays the foundation for end-to-end scalable, virtualized, and cloud-native solutions. **This document highlights additional updates from previous deliverables D2.1 and D2.2 to requirements, further analyses trustworthiness requirements for both digital and analog components, and the underlying lab validations to validate the architecture's robustness against emerging threats.**

Use cases and architectures were analyzed to refine the technical requirements for three primary scenarios: Extended Reality (XR), Automotive Infrastructure, and Smart Cities. The project emphasizes privacy, reliability, and energy efficiency across the aforementioned use cases. For instance, XR requires ultra-low latency for seamless experiences, while smart cities depend on energy-efficient sensors to sustain long-term data collection. Trustworthiness across distributed systems in e.g. radio access as well digital computing infrastructure, including secure communication and execution, remains central to the project's goals.

Digital components include innovations like the RISC-V many-core accelerator, which enhances signal processing efficiency, and advanced MAC scheduling techniques for radio resource management using AI-driven as well as well-defined innovative solutions. These components address performance bottlenecks in latency and throughput, enabling robust 6G base station and core operations. Additionally, virtualization frameworks ensure secure and efficient resource sharing, supporting trust and scalability in distributed environments.

On the analog front, the project advances radio link authentication and ultra-high-speed data interconnect systems for wider range of spectrum. These additional analog trustworthy components further enhance the end-to-end communication reliability and reduce vulnerabilities across the network. Notable contributions include RF fingerprinting for enhanced security and sub-THz data links, paving the way for high-speed connectivity.

Lab validations verify the feasibility of theoretical advancements by simulating real-world scenarios. The M^3 platform, for example, demonstrates its ability to handle hardware vulnerabilities securely, while maintaining operational efficiency. Similarly, sub-THz plastic waveguides validate innovative interconnects, achieving high data rates and energy efficiency. These validations enhance COREnext's capacity to meet its performance and trustworthiness benchmarks.

In conclusion, COREnext's progress toward building a trustworthy 6G platform underscores its commitment to addressing critical challenges in connectivity, security, and sustainability. The integration of cutting-edge digital and analog components, coupled with rigorous validation, ensures that the project contributes significantly to Europe's leadership in 6G development.

Table of Contents

List of Figures

List of Tables

Acronyms and Definitions

AI	Artificial Intelligence
CPU	Central Processing Unit
DoS	Denial of Service
DSP	Digital Signal Processor
FPGA	Field Programmable Gate Arrays
IoT	Internet Of Things
ITS	Intelligent Transport System
HW	Hardware
MMU	Memory Management Units
RCE	Remote Code Execution
RF	Radio Frequency
SoC	System On Chip
TA	Trusted Authority
TCB	Trusted Computing Base
TCU	Trusted Communication Unit
TEE	Trusted Execution Environments
WP	Work Package
XR	Extended Reality

1 Introduction

COREnext with its strong and ambitious goal of building a disruptive 6G hardware platform in Europe aims to provide an efficient and trustworthy disaggregated compute architecture as to be able to unleash the full potential of virtualized and cloud native solutions with a focus on end-to-end operating model. COREnext has set Trustworthiness as a foundational critical Key Value Indicator and is deploying through its working program extensive efforts to advance this main 6G research challenge. B5G and 6G enabled new use cases and applications and will rely on further digital transformation that requires connectivity with high level of security and low latency for a variety of devices such as sensors, robots, cameras, tablets, head-mounted displays, etc.

This deliverable aims to revisit the initial requirements (from D2.1) and trustworthiness aspects (from D2.2) for the digital, analog and lab validator solutions of COREnext. Specifically, it proceeds to the reviewing of the requirements of the key digital and analogue components (accelerators, etc.) designed to reach high isolation and to orchestrate for trustworthiness. This will prevent external threats or attacks that could occur through the various end-end-end building blocks of a 6G system (base stations, terminals, etc.). Finally, there is a review of the requirements for lab validations to demonstrate their efficiency against the project's initial external attack scenarios.

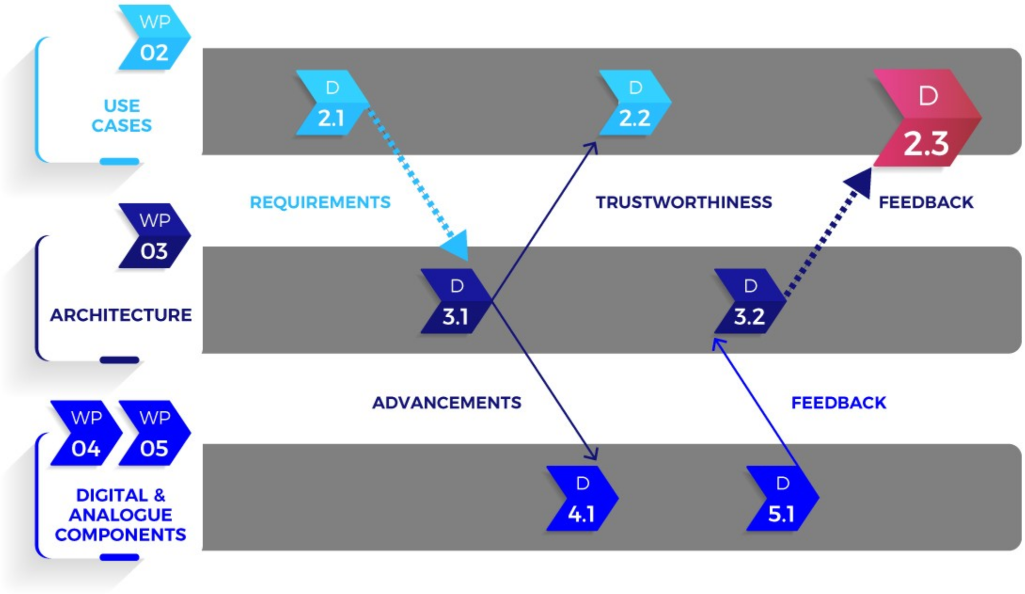

Figure 1: Information flow between deliverables

2 Use Cases and Architecture

In this section we revisit the use cases and architectures and demonstrate how the proposed architecture and the individual technical contributions linked to the use cases considered in the COREnext project.

2.1 Use Cases and Technical Properties

We start with the individual technical contributions of COREnext within the proposed architecture and discuss their relevance for the selected use cases, which is summarized in Table 1.

	XR	Automotive Infrastructure	Smart City
Energy-efficient signal processing	High	High	High
Heterogeneous compute platform with TEEs	Medium	High	Medium
Power-efficient high-throughput interconnect	Low	High	Low
Radio link authentication and infrastructure attestation	Medium	High	High

Table 1: Use Cases and technical properties overview

2.1.1 Extended Reality

The Extended Reality (XR) use case has a high demand for low-latency and high-throughput communication as significant amounts of data need to be transferred over the network. For that reason, signal processing needs to be accelerated to fulfill these requirements. Depending on the concrete use case, a trustworthy compute platform is also important. For example, when remotely collaborating on confidential data protecting this data is critical, which can be addressed by the compute platform proposed by COREnext. In contrast, the high-throughput interconnects are more suitable for base stations or edge clouds and thus less critical in typical XR scenarios. Finally, we deem the radio link authentication as potentially helpful for the XR use case as it provides additional confidence when authenticating participants in security-critical scenarios. This for instance will allow low-latency security defense mechanisms without the need to rely on higher layer protection strategies incurring higher latency.

2.1.2 Automotive Infrastructure

The automotive infrastructure involves several different devices and thus has a high demand for all contributions within COREnext. At first, highly efficient base stations will be required to support the potentially high demand of traffic in dense areas, which in turn asks for energy-efficient and accelerated signal processing. In contrast to the XR use case, high-throughput interconnects can be employed in disaggregated base stations to combine a high resource utilization with efficient communication. As one key aspect of the automotive infrastructure is the ability to offload computations from cars to the infrastructure, base stations require a trustworthy compute platform with trusted execution environments (TEEs) to perform these securely. Finally, with several different devices such as cars, mobile devices of passengers, and sensors in the infrastructure, the radio link authentication provides an important additional factor in the authentication of devices.

2.1.3 Smart City

The smart-city use case also involves several different kind of devices and in particular low-energy devices, but is less demanding regarding the base station in terms of throughput. Therefore, it shares the high demand for a radio link authentication with the automotive-infrastructure use case, but benefits less from high-throughput interconnects. As the devices are typically very resource and energy constraint, power-efficient signal processing is also critical. However, as these devices often only run a single software component, isolation of components with the trustworthy hardware platform is less important.

2.2 Threat Analysis of the Architecture

After revisiting the link between the individual parts of the architecture that COREnext contributes to, we now want to focus on the M³-based Digital Compute Architecture as one of COREnext's critical contributions. We study the effectiveness of the M³ architecture by systematically walking through all building components and their associated attack vectors and analyzing whether the architecture can either prevent or mitigate these attacks. The attacks including real-world examples and their coverage by architecture are shown in Table 2.

No.	Attack	Covered
1.	Attacks on physical properties	-
2.	Bugs in trusted computing base (M³ kernel, TCUs, etc.)	-
3.	DoS attacks over NoC/DRAM	✓
4.	Side-channel attacks via NoC/DRAM	-
5.	Transient execution vuln. in user core	✓
6.	Arbitrary bug in user core	✓
7.	Bugs in OS	✓
8.	Bugs in user software	✓
9.	Software timing bugs	-

Table 2: Potential attacks and their coverage in proposed architecture

We additionally show the locations of the individual attacks in Figure 2 on potential locations of attacks (green components are part of the TCB, red ones are not).

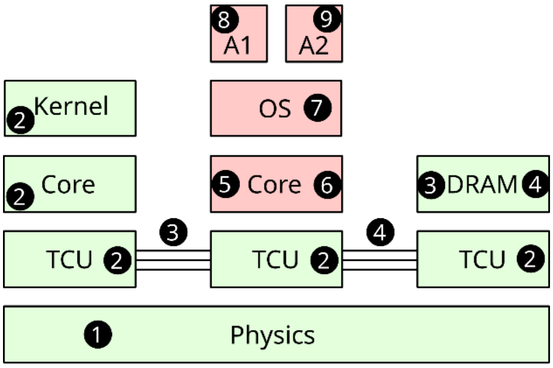

Figure 2: Potential locations of attacks

The first class of attacks leverages physical properties of the hardware labeled (1) in the figure such as voltages or undesirable side effects between DRAM cells These attacks cannot be prevented by the M³ architecture and thus need additional countermeasures. The architecture can also not prevent attacks that exploit vulnerabilities of TCB (2) as it assumes that these components work correctly. Due to the simplicity of the TCU and the M³ kernel, formal verification could be used to increase the trust in these components. Denial-of-service (DoS) attacks over a shared NoC or DRAM (3) can be prevented by a credit system as demonstrated by M³. However, timing attacks based on these shared resources (4) are still left for future work.

In contrast to existing point mitigations such as KPTI [38], the M³ architecture defends against arbitrary bugs in user cores (5) and (6) if the victim runs on a different tile than the attacker. In other words, although the bugs itself might still exist in user cores, the architecture prevents that they can be used to attack other tiles. Similarly, the OS that runs on a user tile has no permissions beyond its own tile, limiting the impact of bugs (7) to the same tile. Bugs in user software (8) are also contained within the same tile and traditional means such as address spaces and different CPU modes can be used to shield them from other software on the same tile (as shown by M³'s tile-local multiplexer).

Software-only timing issues (9) like operations requiring varying periods of time cannot be prevented by the M³ architecture though and thus need additional protection.

3 Updated Requirements and Trustworthiness

3.1 Initial Requirements from D2.1

The project's deliverable D2.1, elaborated on the initial requirements by focusing on three main use cases namely: XR, Automotive Infrastructure and Smart cities. These are used as reference use cases in the COREnext project. This section provides a brief recap of the different requirements needed for the use cases described above, considering the level of performance required and how relevant is the contribution of COREnext in the use-case for each specific area.

Requirements	XR	Automotive infrastructure	Smart city
R1: Privacy, reliability, integrity	High	High	High
R2: Trustworthy analogue access	High	High	High
R3: Trustworthy distributed code execution	Medium	High	Medium
R4: Energy-efficient connectivity	High	High	High
R5: Energy-constrained devices	Low	Low	High
R6: Ultra-low latency	High	Medium	Low
R7: High-Capacity Connectivity	High	Low	Low

Table 3: Use case requirements and relevance

The requirement R1 of **reliability, integrity and privacy** are relevant for all the reference use cases because the data needs to be always correctly transmitted without error. The information must be protected and not accessible. COREnext wants to guarantee the level of security introducing, for all the use cases, a relevant level of trustworthiness – even at the analogue access.

R2 and R3 on Trustworthiness is crucial in analogue access and distributed code execution architecture. This includes ensuring that code is executed securely, reliably, and with integrity across a distributed system. For example controlling access to distributed components, controlling the integrity of code executed across, monitoring and auditing distributed code execution, and keeping the distributed system up to date with security patches, bug fixes, and software updates.

In the context of technology and networks, including 6G networks, R4 and R5 on **energy efficiency** plays a crucial role, it is relevant in all the reference use cases to maintain a sustainable system: reducing ICT carbon footprint, saving costs for system owners and promoting a positive public image.

Performance requirements for data transmission can be summarized by a **latency** requirement (R6) and a **capacity** requirement. A delay requirement refers to the maximum tolerable latency in the reception of transmitted data that can be acceptable for the user experience. The requirement of **High Connection capacity** refers to the maximum number of simultaneous connections that a network or system can support with a high amount of transmitted data . It represents the ability of the network to accommodate a high volume of data from various devices or users both in radio access and core network. Overall, the COREnext project will enable novel **trustworthy-by-design** platforms and computing architectures, capable of efficiently and securely integration of third-party accelerators, supporting even the most demanding 5G/6G processes in cloud servers, base stations, and client-side devices.

3.2 Trustworthy Digital Components

As presented in the introduction, the COREnext architecture is composed out of four component clusters:

- power-efficient signal processing,
- power-efficient high-throughput interconnects,
- radio link authentication and infrastructure attestation, and
- a heterogeneous computing platform with trusted execution environments (TEEs).

Of these four component clusters, the first two relate to efficiency, and the last two to trustworthiness. The first and the last components require innovation in the digital domain, and the middle two in the analogue one. Since WP4 is addressing the digital domain, this section concentrates on how we fulfill the requirements for power-efficient signal processing with our components for the heterogeneous acceleration of 6G processing capabilities and how we aim to implement a heterogeneous computing platform with trusted execution environments with components for trustworthy computation and orchestration.

3.2.1 Components for Heterogeneous Acceleration of 6G Processing Capabilities

3.2.1.1 Many-core RISC-V Accelerator for Low-PHY Processing

The growth in wireless spectrum and size of Massive MIMO in B5G telecommunications brings parallelization opportunities in the base station digital signal processing. At the same time, the rapid evolution of 5G standards and technologies demands flexible programmable processing engines. The TeraPool Many-Core cluster developed by ETH-Zurich as a contribution to COREnext combines the two requirements, offering 1024 individually programmable lightweight cores and 4MiB of memory for large-scale parallelization of the lower-PHY workloads.

As part of the project, we placed&routed a TeraPool design instance in an advanced technology node (GF12 LPPLUS FinFet), running at > 875MHz in the typical corner. A fully software-defined uplink application (the data-processing operations of the PUSCH in a high-load use-case with 64 antennas, 274 resource-blocks and 4 transmitting-UEs MIMO) runs on the cluster in 1.7ms. TeraPool consumes on average 5.4W on this task. Considering the base station average power consumption ~2kW and 90% of the power dissipated in power amplifiers and air-conditioning, this result well matches the upper bound of <100W for digital signal processing.

Funded by the European Union

Requirements	Many-core accelerator
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	High
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 4: Contribution of the Many-core accelerator to fulfilling the requirements

3.2.1.2 Programmable Vector Processing Accelerator

Flexible and scalable solutions will be needed for B5G communications processing systems. RISC-V processors enhanced with vector processing capabilities as specified by the RISC-V vector extension pose an interesting base for such systems because they provide an efficient means of exploiting data-level parallelism, which is heavily present in communications kernels.

As a contribution to COREnext, TUD studies the programming of vector processors for communications signal processing and how to improve their architecture for the application. The technology-independent goal is to maximize the utilization of the vector processor's functional units. For a given workload, this is equivalent to reducing the cycle count. At first, we introduced dual load to speed up binary vector operations and demonstrated a cycle count improvement of up to 21 %. Then, we introduced dynamic bank layout in the vector register files to reduce detrimental access conflicts. It either allows to decrease the cycle count by 3 % or to save significant area for operand queues or banks.

Requirements	Programmable Vector Processing Accelerator
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	High
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 5: Contribution of the vector processor accelerator to fulfilling the requirements

3.2.1.3 FEC Accelerator

Forward Error Correction (FEC) improves data transmission reliability by adding redundancy to data. Parity-Check (LDPC) codes, offers strong error-correction with efficient performance, widely used in high-throughput systems like wireless communication (e.g., IEEE 802.11n). Due to increasing demand for throughput, data rates for FEC increase as well, and the barrier of 100 Gb/s wireless communication has already been exceeded several times[1]. We expect data rates of 1 Tb/s or higher in the next few years

As a contribution to COREnext, IHP developed a high-performance LDPC decoder for the IEEE 802.11n WLAN standard, achieving up to 1 Tb/s (for the unencoded data). Verified in 28 nm CMOS, the design operates at a high clock frequency of up to 2 GHz, even in worst-case conditions (slowest process, 0.9V, 125 C), due to a fully unrolled, pipelined architecture with seven pipeline stages per iteration. It supports 648-bit codewords with 4-bit quantized Log-Likelihood Ratios (LLRs) and processes 2592-bit words per cycle using a min-sum algorithm with a (648,540) parity matrix. With efficient minimum search through binary tree comparators, this design meets high-throughput requirements for modern wireless systems.

Requirements	FEC accelerator
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	Low
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 6: Contribution of the FEC accelerator to fulfilling the requirements

3.2.1.4 MAC Scheduling Accelerator

In 5G/6G networks, MAC scheduling plays a key role in efficiently managing downlink and uplink data streams, balancing ultra-high throughput and strict Quality of Service (QoS) requirements. This involves optimizing various parameters like latency, throughput, and channel conditions, making scheduling a complex but critical function.

IHP is developing an advanced MAC architecture that addresses these demands using deep reinforcement learning (RL) for intelligent, low-latency radio resource management. The architecture includes separate control and data processing domains: the control plane on CPUs for resource allocation decisions, and a data plane with a specialized accelerator for high-performance data handling. The control plane leverages the NVIDIA Deep Learning Accelerator (NVDLA) within IHP's RISC-V-based CRISPY platform. Designed for a single-cell multiuser MIMO system, it supports resource block group (RBG) allocation to the user equipment (UE) with focus on TDMA and OFDMA multiple access methods. Initial evaluation and RL training of this system will be conducted in NS-3, which is integrated with OpenAI Gym.

[1] A. Karakuzulu, W. A. Ahmad, D. Kissinger and A. Malignaggi, "A Four-Channel Bidirectional D-Band Phased-Array Transceiver for 200 Gb/s 6G Wireless Communications in a 130-nm BiCMOS Technology," IEEE Journal of Solid-State Circuits, vol. 58, no. 5, pp. 1310-1322, 2023.

Requirements	MAC Scheduling Accelerator
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	Low
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 7: Contribution of the MAC Scheduling Accelerator to fulfilling the requirements

3.2.2 Components for Trustworthy Computation and Orchestration

3.2.2.1 FPGA Multi-tenancy

Deployment of functions in the Cloud will be the rule in 6G networks. To ensure real-time, low latency and energy efficient processing, FPGA will have a key role to play. FPGA are already widely used in Cloud infrastructures (Microsoft, Amazon etc.) but are still underexploited due to the lack of trusted and secure space sharing solutions. To even more increase the power efficiency achieved by the use of FPGAs, those platforms need to be available for several users at the same time aka multi-tenancy.

On the one hand, multi-tenancy brings more flexibility and more power efficiency, on the other hand it brings new security challenges. Indeed, users having access to shared resources inside the FPGA open possibilities to potential attacks or data leaks, in the same way when multiple users are communicating with a base station.

The main interest of the contribution is to prevent data leaks between offload functions inside FPGA logic. To reach this goal, we propose an encryption framework which should ensure data integrity between software and hardware offload function target. Each offload process should be able to use a personal AES key which is shared to the offload function logic by an asymmetric elliptic encryption key generated and provided by the FPGA.

Requirements	FPGA multi-tenancy
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	Low
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 8: Contribution of the FPGA multi-tenancy to fulfilling the requirements

3.2.2.2 Digital Signal processor Virtualization

Digital Signal Processors (DSPs) are specialized processors for implementing efficiently some demanding signal processing functions like for example Forward Error Correction (FEC) encoding and decoding. Such components may be used on edge nodes of mobile network computing infrastructures that can be shared among multiple base station instances. DSPs may therefore have to be shared by multiple processes to be used efficiently. This contribution aims at studying the mechanisms that enable sharing DSPs between multiple processes.

The main interest of this contribution is to enhance the efficiency of computing in term of computing resource and energy usage while being able to guarantee some quality of service to the users. Through its attention toward efficiency, this contribution addresses the requirements on energy-efficient connectivity while the attention given to the quality of service enables to provide ultra-low latency.

Requirements	Digital Signal Processor Virtualization
R1: Privacy, reliability, integrity	Low
R2: Trustworthy analogue access	Low
R3: Trustworthy distributed code execution	Low
R4: Energy-efficient connectivity	High
R5: Energy-constrained devices	Medium
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	Low

Table 9: Contribution of DSP virtualization component to fulfilling the requirements

3.2.2.3 M³ – Microkernel-Based System for Heterogeneous Many-Cores

M³ is a compute platform developed with a hardware/software co-design approach, which provides strong isolation between the compute workloads it is hosting. Trustworthiness is realized by the system building on a very small trusted computing base (TCB). The TCB is comprised of those hardware and software components, which must operate correctly for the system to securely provide basic compute functionality. Because M³ is based on a microkernel approach, this TCB does not include device drivers, network stacks, or file system implementations, removing a large attack surface from the TCB.

In terms of compute hardware, M³ builds upon a tiled architecture, with a Trusted Communication Unit (TCU) connecting the tiles via a Network-on-Chip (NoC). Only this TCU, the NoC, and the compute tile executing the microkernel are part of the TCB. These components are of low complexity and therefore amenable to being formally verified. Taken together, these mechanisms allow M³ to offer general-purpose compute functionality and secure integration of accelerators, while leveraging its isolation features to tolerate vulnerabilities in software or hardware.

Within the COREnext project, M³ is developed further by adding functionality for Trusted Execution Environments (TEEs). Such TEEs are execution containers with strong isolation — which M³ already provides – and which can be attested to outside parties. Remote attestation is a process, where a platform like M³ uses a hardware-embedded root-of-trust to prove its authenticity and integrity to an external entity. Attestation therefore allows outside parties to establish a verified trust relationship to an M³ system and thus allows the construction of secure and verified distributed client-to-edge-to-cloud compute environment.

Requirements	M³ Platform
R1: Privacy, reliability, integrity	High
R2: Trustworthy analogue access	Medium
R3: Trustworthy distributed code execution	High
R4: Energy-efficient connectivity	Medium
R5: Energy-constrained devices	Medium
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 10: Contribution of M³ to fulfilling the requirements

3.2.2.4 IoT Management

As mentioned also in D4.1, the IoT has revolutionized the way we interact with technology, transforming everyday objects into smart devices interconnected through the internet. IoT management and devices are now pervasive in our homes, workplaces, and public spaces, providing us with convenience and efficiency. As the IoT landscape continues to expand, so does the importance of ensuring the trustworthiness of these devices (e.g., are they reliable; secure; ensure privacy of IoT devices and the data they collect and transmit). With these devices becoming increasingly integrated into critical systems and handling sensitive data, their trustworthiness has a direct impact on user safety, data security, and overall system integrity.

As a result, a Trust Manager is proposed and implemented. Edge operational devices, such as drones and collaborative robots (cobots), communicate with the Trust Manager Orchestrator to ensure efficient and secure operations. These devices generate a plethora of data related to network Key Performance Indicators (KPIs) and device-specific metrics. This data is essential for understanding the operational status and capabilities of each device within the network. Machine Learning techniques, particularly K-means clustering, are employed to analyze this data. By doing so, devices with similar characteristics are grouped into clusters, facilitating more streamlined and effective management. Finally, the Trust Manager uses the trust index and the specific prerequisites of each task to determine the optimal input node for the task at hand. This decision-making process is crucial for ensuring that tasks are assigned to the most trustworthy and capable devices, thereby enhancing the overall security and efficiency of the network. The table below summarizes the main requirements with respect to this component.

Requirements	IoT Management
R1: Privacy, reliability, integrity	High
R2: Trustworthy analogue access	High
R3: Trustworthy distributed code execution	Medium
R4: Energy-efficient connectivity	High

Funded by the European Union

R5: Energy-constrained devices	Low
R6: Ultra-low latency	High
R7: High-Capacity Connectivity	High

Table 11: Contribution of IoT Management component to fulfilling the requirements

3.3 Trustworthy Analog Components

This section provides an overview of the efforts and innovations in developing trustworthy analog components, a critical element in enhancing the security and efficiency of next-generation communication systems. These components, designed and implemented through a collaborative effort led by WP5 and supported by its partners, address the challenges of physical layer security and ultra-high-speed data interconnects.

Subsection 3.3.1 explores components for trustworthy radio links, focusing on RF fingerprinting techniques for physical layer security. These findings contribute to the validation activities detailed in subsection 3.4.1. Subsection 3.3.2 addresses components for ultra-high-speed data interconnect systems, encompassing a range of innovations from material development to full-scale prototyping. The integrated approach includes components such as Polymer Microwave Fibre (PMF), sub-THz transceivers, and advanced packaging solutions. These efforts aim to maximize data rates and energy efficiency while ensuring seamless integration and system linearity, particularly at higher frequencies. The application of these components is further elaborated in subsection 3.4.4, focusing on high data-rate interconnects leveraging sub-THz plastic waveguides.

This comprehensive approach underscores the commitment of WP5 and its partners to delivering innovative and reliable analog components that align with the overarching objectives of the project.

3.3.1 Components for Trustworthy Radio Links

The RF fingerprinting work is one of the components in the context of physical layer security. The work involves analyzing acquisition techniques and identifying suitable RF non-idealities (imperfections) that can be used for fingerprinting. Both SW- and HW-based (sub-6GHz, sub-10GHz) platforms jointly designed in WP5 and WP6 have been developed to generate RF fingerprint training data for testing and validating the concepts.

The proposed sub-THz multiuser MIMO link level simulator will be used to analyze beamforming-based security defence mechanisms in the sub-THz frequency band, considering different architectures and RF impairments in this challenging frequency band.

All findings feed into the validation activities described in subsection 3.4.1 where they are extensively discussed and related to the project requirements.

3.3.2 Components for Ultra-high-speed Data Interconnect

For ultra-high-speed data interconnect, work has been carried out at all needed system components, going from device characterization and measurements to full bottom-up design and prototyping. Those components are:

- the Polymer Microwave Fibre (PMF),

Funded by
the European Union

- waveguide-to-PMF transitions,
- Sub-THz SiGe BiCMOS transceivers,
- Substrate Integrated Waveguides (AFSIW),
- Embedded Wafer Level Ball Grid Array (eWLB) packages,
- PMF holder concepts and prototypes.

The close cooperation between the project partners aimed at achieving the highest data rate at the lowest energy consumption for the complete data interconnect system. In RF designs, especially in the high frequency ranges targeted, the matching between the components is crucial and dominates system linearity, achievable bandwidth and power efficiency.

Relating the individual components to the project requirements was seen as of little value. Therefore, these aspects are covered in the subsection 3.4.4 which follows and is related to high data-rate fly-over interconnects using Sub-THz-over-plastic waveguides showcase.

3.4 Lab Validations

This section provides details on the lab validators which are developed by partners. Specifically, in the project we have 4 validators which are:

- Trustworthy radio link validation
- M3 platform lab showcase
- Accelerated signal processing capabilities based on a RISC-V platform
- High data rate fly-over interconnects using sub-THz over plastic waveguides showcase

3.4.1 Trustworthy Radio Link Validation

Validator 1 provides realistic assessment of the theoretical concepts and components developed in T5.1 and T4.3 with "hardware-in-the-loop" demonstrators. In other words, we aim through the validation setups in this task to (i) assess the high-throughput sub-THz link robustness to RF non-idealities after compensation in the sub-THz PoC, (ii) generate experimental data to validate the HW fingerprint for a wider range of RF components and spectrum. We pursue a mixed approach in which some SW components of the PHY models in T5.1 and T4.3 are replaced by HW RF components. In the nutshell we have three experimental setups as sketched in **Figure 3**, namely:

- Sub-THz PoC where the proposed end-to-end link simulator from T5.1 with most relevant sub-THz hardware imperfections is further extended with 140GHz RF front-end that is non-ideal and reconfigurable. Note that the Hardware development activity of the D-band chip is outside of our project scope.
- Sub-6/10GHz end-to-end link-level demonstrators are developed between T5.1, T4.3 and T6.1 with concepts, signal processing and algorithms (AI-based) investigated in T5.1 / T4.3, while the two demonstrators for the validation in T6.1 generates validation data for AI neural networks training and inference. There two setups as described next:
  - The Sub-6GHz PoC with SDR and different external RF HW components is flexible and allows to test different RF components, while
  - The Sub-10GHz 'BS TX PA -in-the-loop' is dedicated to capture manufacturing variations of the Doherty PAs widely used in the base stations.

With the different experimental setups, different KPIs are targeted that can be summarized as

- In the sub-THz PoC, the aim to control RF non-idealities after compensation and ensure robust sub-THz link with a target bit error ratio of 10^{-6} and target throughput in the range of 10-20Gbps,
- Validation of HW RF fingerprints over a wider spectrum range involves identifying RF devices with an accuracy that depends on three factors: (i) the device type (e.g., PAs), (ii) the number of devices, and (iii) the device manufacturer, since devices from the same manufacturer or model may have more subtle differences in their RF fingerprints. The minimum target acceptable accuracy for distinguishing between two devices of the same manufacturer is set at 80%, which decreases as the number of devices increases.
    - For the sub-6GHz PoC the target to identify 2 devices from the same manufacturer
    - For the sub-10GHz the target is to be able to identify the base station Doherty PA with the manufacturing variations (HW spread).

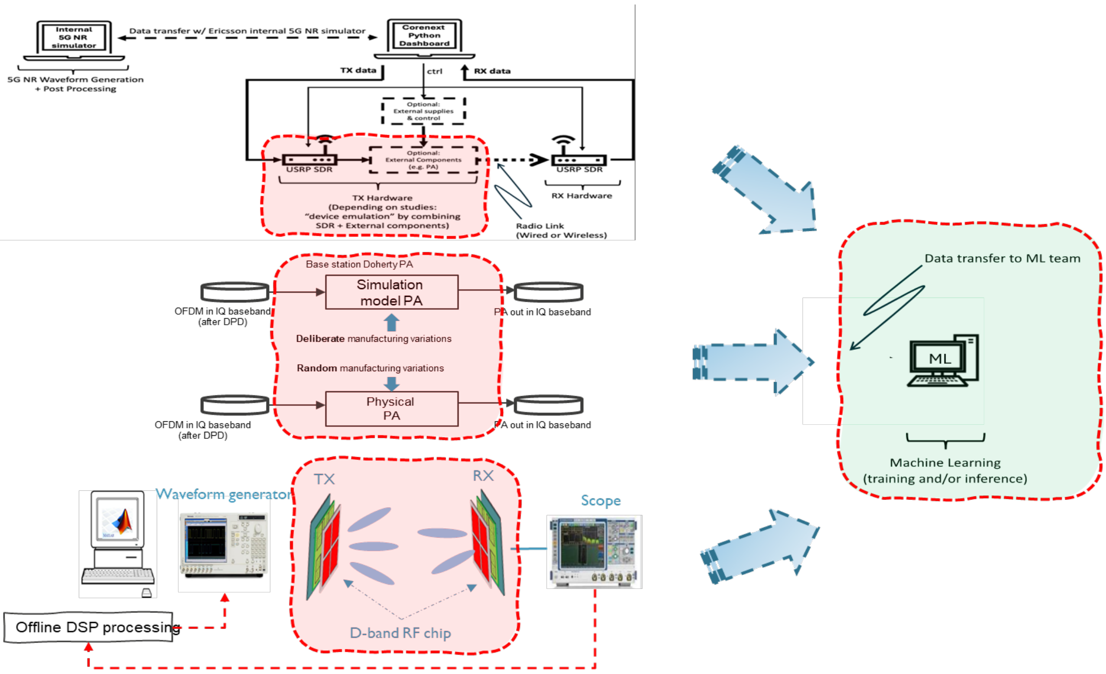
Internal 5G NR simulator
Data transfer w/ Ericsson internal 5G NR simulator
Corenext Python Dashboard
5G NR Waveform Generation + Post Processing
TX data
ctrl
RX data
Optional: External supplies & control
USRP SDR
Optional: External Components (e.g. PA)
USRP SDR
TX Hardware (Depending on studies: "device emulation" by combining SDR + External components)
RX Hardware
Radio Link (Wired or Wireless)
Base station Doherty PA
Simulation model PA
OFDM in IQ baseband (after DPD)
PA out in IQ baseband
Deliberate manufacturing variations
Random manufacturing variations
Physical PA
OFDM in IQ baseband (after DPD)
PA out in IQ baseband
Waveform generator
TX
RX
Scope
D-band RF chip
Offline DSP processing
Data transfer to ML team
ML
Machine Learning (training and/or inference)

Figure 3: Block diagram of the sub-6GHz, sub-10GHz and sub-THz PoCs

3.4.2 M3 Platform Lab Showcase

This validation effort demonstrates the COREnext improvements to digital trustworthiness. A cornerstone of these efforts is the M^3 platform, which is developed as one component within WP4. Within COREnext, M^3 is augmented with hardware and operating system functionality to provide Trusted Execution Environments (TEEs). The validation of these efforts follows two paths:

1. We will demonstrate the unique selling point of M^3 itself, which consists of providing isolation against vulnerabilities in software and in hardware. Our demonstrator shows how M^3 handles a RISC-V CPU with a deliberately placed hardware-level vulnerability, which a malicious software application could exploit. However, this protection comes at the cost of additional isolation hardware, which also adds latency to compute workloads. Our goal is to

evaluate the platform quantitatively by running simulated workloads from mobile communication signal processing as it would run in a 6G base station. Running such workloads on M^3 will allows us to gauge, how much the additional latency is impacting these workloads and whether the security benefits justify the additional latency costs.

2. We will illustrate conceptually, how other components developed within WP4 can be integrated into the M^3 platform. Some components primarily address energy efficiency and can be integrated as individually isolated accelerators into components into an M^3 chip. Other components contribute to architectural trustworthiness such as AI radio fingerprinting or FPGA multi-tenancy. These components are co-designs of hardware and software building blocks. With selected components, we can show how M^3 can securely execute the software part within its novel Trusted Execution Environment facilities.

Altogether, we will qualitatively evaluate M^3's contributions to the COREnext goal of offering a trustworthy digital compute platform. At the same time, we quantitatively evaluate that M^3 does not have a prohibitive negative impact on the overall system's energy efficiency.

3.4.3 Accelerated Signal Processing Capabilities Based on a RISC-V Platform

Validator 3 aims at delivering a porting of the mobile network software OpenAirInterface (OAI) on a RISC-V based System-on-Chip (SoC) accelerated with Kalray's VLIW core (KVX). Indeed, OAI is a piece of software with heavy computing power requirements that is usually achieved by relying on powerful multi-core AMD64 processors with 512 bits SIMD instruction set extensions. Porting OAI to a RISC-V based SoC is an opportunity to replace powerful general-purpose processors with well-chosen efficient accelerators for demanding functions. On a SoC, OAI could rely on LDPC coding core and on a KVX tile for channel estimation, channel equalization and other demanding tasks.

The lab validator aims to validate the ability to run the physical layer of OAI by achieving some defined quality of service and while showing an improvement in efficiency in resource and energy usage. A reference for Quality of Service could be a state of the art high-capacity OpenAirInterface O-DU – Distributed Unit in the O-RAN standard –. This is the lower layers of a 100MHz 4x4 and 4-layers TDD base station where physical procedures of one slot are done in less than 0.5 millisecond. The available instances of OpenAirInterface allow us to assess different components of the demonstrator with regards to the reference so that the demonstrator could be evaluated whatever its maturity by the end of the project. It addresses therefore the requirements regarding quality of service which are the **ultra-low latency** and **high-capacity connectivity** requirements as well as **energy-efficient connectivity** and **energy-constrained devices** requirements.

3.4.4 High Data-rate Fly-over Interconnects Using Sub-THz-over-plastic Waveguides Showcase

Lab validator 4 will show the capability of interconnection between RF front-end and antenna and the propagation of the sub-THz signal over plastic fiber. RX and TX modulator are built in Infineon SiGe technology leveraging on the high frequency and high-power capability of the technology. The two separated dies will be integrated in eWLB package in which a Vivaldi antenna is integrated to allow the conversion of the electrical signal in an electromagnetic wave that can be propagated through the plastic waveguide (PMF). On the other end of the fiber the signal will be demodulated in base band frequency and distributed.

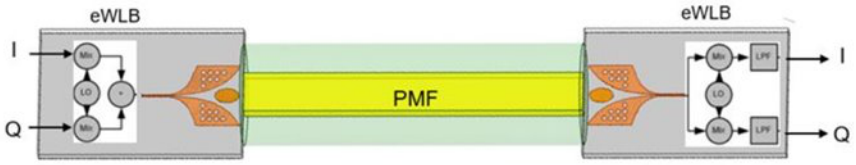

Figure 4: Concept system

A dedicated evaluation board (EVB) will be designed and fabricated to allow the demonstration in lab. The board will be 6 layers stacked printed circuit suitable for 30GHz signal propagation. Input signal will be applied via a K-connector and baseband channels will be applied via SMA connectors.

KPIs are high –data rate (high-capacity connection); data transmission will be over 200GHz frequency. High energy-efficiency per bit rate.

The following table summarizes the contribution of all 4 validators to fulfilling the Requirements.

Requirements	Validator 1	Validator 2	Validator 3	Validator 4
R1: Privacy, reliability, integrity	High	High	Low	Medium
R2: Trustworthy analogue access	High	Medium	Low	Low
R3: Trustworthy distributed code execution	Medium	High	Low	Low
R4: Energy-efficient connectivity	High	Medium	Medium	High
R5: Energy-constrained devices	Low	Medium	High	High
R6: Ultra-low latency	High	High	High	High
R7: High-Capacity Connectivity	High	High	High	High

Table 12: Contribution of Validators to fulfilling the requirements

4 Conclusions

The COREnext project demonstrates significant advancements in developing a robust, trustworthy, and energy-efficient 6G hardware platform. By addressing critical use cases–Extended Reality, Automotive Infrastructure, and Smart Cities–the project refines requirements that balance performance, reliability, and sustainability. Innovations in digital and analog components, such as heterogeneous signal processing, trustworthy execution environments, high-speed interconnects, and the HW fingerprint ensure that COREnext is well-equipped to tackle challenges posed by next-generation applications and infrastructure needs.

A key achievement of COREnext is its comprehensive approach to trustworthiness and validation, combining theoretical innovations with many practical demonstrations. The M³ platform exemplifies how secure computing can integrate hardware and software solutions to enhance both performance and security. Lab validations, such as sub-THz interconnect testing and AI-driven RF fingerprinting, verify the project's solutions against real-world conditions, ensuring scalability and resilience across diverse use cases. These efforts not only meet high standards of data privacy and reliability but also provide low-latency and energy-efficient solutions critical for future networks.

As the project progresses, COREnext is positioned to set a new benchmark for 6G development by combining cutting-edge technology with a clear focus on sustainability and security. The integration of validated components underscores the project's commitment to creating a transformative, trustworthy platform that will shape the future of communication and connectivity in Europe and beyond.