The zero-tax data center: a use case through quantum resilient communications

Daniel Lawo[1,*], Michał Podleś[1], Raphael Frantz[1], Abraham Cano Aguilera[1], Dismothenis Iliadis-Apostolidis[2], Jeronimo Sanchez[2], Sokol Kosta[2], Idelfonso Tafur Monroy[1], José Luis Imaña[3], J.J. Vegas Olmos[4]

[1]*Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands*

[2]*Department of Electronic Systems, Aalborg University, Copenhagen, Denmark*

[3]*Department of Computer Architecture and Automation, Complutense University of Madrid, Spain*

[4]*Software Architecture, NVIDIA Corporation, Ofer Industrial Park Yokneam, Israel*

**e-mail: d.c.lawo@tue.nl*

ABSTRACT

Data centers, cloud and edge computing platforms are absorbing heavy AI-based workloads. This raises two concerns: how to reduce the networking overhead by host processors in order to focus only on AI-processing, and how to provide quantum resilient level protection for any connection running such workloads. In this paper, we will show case the concept of zero-tax data center, which is achieved through offloading of networking functions from the host to the network. In particular, the implementation of postquantum cryptography algorithms providing quantum resilient links and how the processing its complex processing is moved from host to network through the utilization of data-processing units.

Keywords: data center, networking, accelerated computing, high performance computing, postquantum cryptography.

1. INTRODUCTION

Artificial Intelligence has shaken the Accelerated Computing ecosystem by stimulating the converge of high-performance computing (HPC) and cloud-computing. This is adding pressure to data centers, regardless of their size and role (edge computing or cloud computing): how do we make sure the processing capabilities are utilized to its maximum? In other words, how do we ensure that processing power is not idle or wasted on overhead, but rather, fully utilize to support the users and applications operating on top? The zero-tax data center emerges from this question and serves as a driving force to conduct research that combines low-level hardware design and firmware to all the way up the software stack. In a networking context, the concept of zero-tax data center might relate to the idea of creating a network architecture that maximizes efficiency, reliability, and security while minimizing costs. This could involve several strategies, related to for example: network virtualization - leveraging technologies such as virtual LANs (VLANs) or software-defined networking (SDN) [1] allows data center operators to create flexible and scalable network architectures. This can enable efficient resource utilization and rapid provisioning of network services; traffic optimization - Zero-tax data centers may implement traffic engineering techniques to optimize the flow of data across their networks. This could involve load balancing, route optimization, and prioritization of traffic to ensure that critical applications receive adequate bandwidth and latency requirements are met; security - implementing robust security measures is essential to protect data and ensure compliance with regulations. This includes firewalls, intrusion detection/prevention systems (IDS/IPS), encryption, and access controls. Network segmentation and micro segmentation can also enhance security by isolating sensitive workloads and limiting the blast radius of potential breaches.

Zero-tax data center, in the context of this paper, relates to the offloading of networking operations from the host to the network (i.e. the network interface card as building block handling networking). This allows the host to focus on its core functions and pay no networking tax (aka. zero-tax data center).

This paper will describe some on-going efforts in the context of the Doctoral Network QUARC, in which the zero-tax data center concept is explored from the perspective of quantum-resilience.

2. QUANTUM RESILIENCE

A commercial quantum computer is expected to be available within the upcoming years while prototype systems or digital and quantum annealers that operate on similar terms are already available. That poses a severe thread to our nowadays used communication systems based on classic cryptographic infrastructure and methodologies.

Asymmetric cryptography, such as for example Rivest–Shamir–Adleman (RSA), is said to be broken by a quantum processor [2], whereas symmetric cryptography, e.g. Advanced Encryption Standard (AES), is expected to be quantum safe as opposed to asymmetric cryptography as long as its key size is doubled [3]. In line with this challenges, large organizations have recently recommended to increase the RSA encryption from 2000bit to

3000bit, whereas computer and mobile phone providers have recently started introducing post-quantum cryptography (PQC) tools to their communication tool kits. Overall, introducing quantum resilience features to all communication networks is an evolutional step that will come over time.

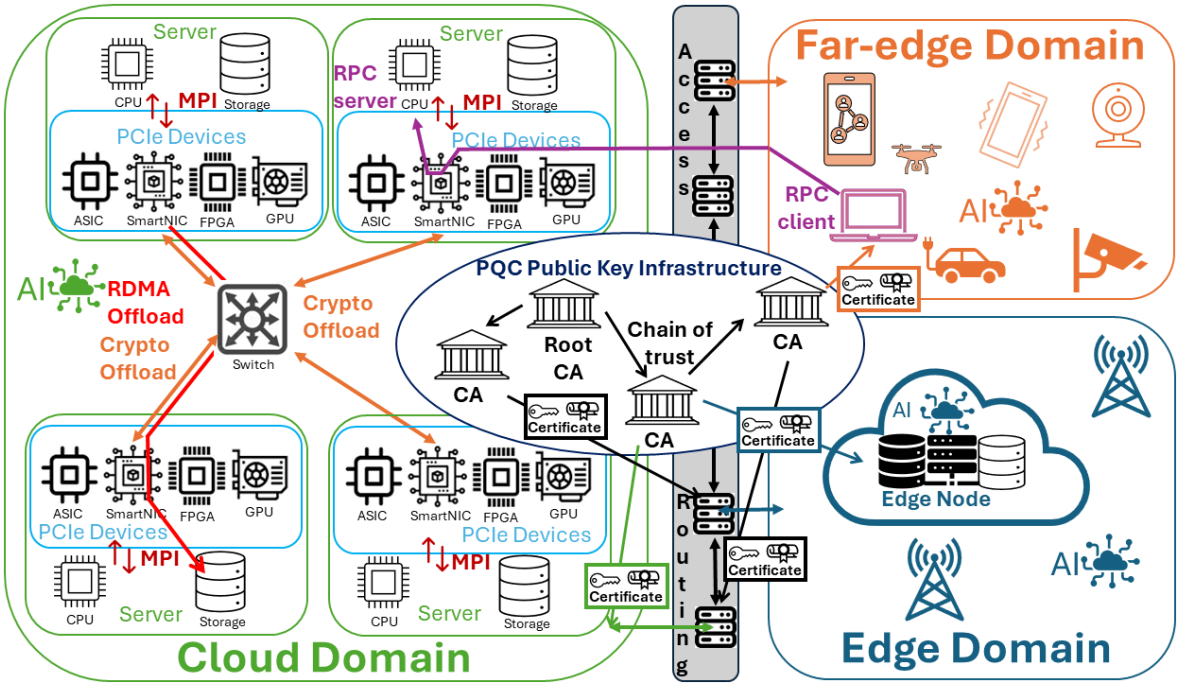

Figure 1. General architecture of network: The Far-edge domain communicates with the Edge domain and the cloud domain using pqc-secured chain of trust that is established by pqc-certificates issued by certificate authorities (CA). In the cloud key features are offloaded to peripheral component interconnect express (PCIe) devices. Those key features include remote direct memory access (RDMA), remote procedure calls (RPC) issued by clients outside of the cloud domain, and message passing through interface (MPI).

2.1 Offloading PQC premises to the network

Post-Quantum Cryptography (PQC) addresses two essential components of secure communication between parties: digital identification and key encapsulation/signature verification. The third component, data exchange, relies on traditional encryption methods such as AES-128 or AES-256. Historically, PQC research has focused on showcasing algorithms or implementations on conventional host devices like CPUs, FPGAs, GPUs, or ASICs. However, as PQC integrates into the networking stack, its ideal execution platforms are NICs, smartNICs, or DPUs. Only recently, with the inclusion of onboard CPU and GPU processors, have NICs possessed sufficient processing power to effectively run PQC algorithms. Figure 1 represents our vision of the evolution of future networks towards offloading of specific tasks to the network.

2.2 Offloading PQC premises to the network

PQC algorithms comprise various components, including the number theoretic transform (NTT), inverse NTT, matrix expansion, and key pairing. Within QUARC, there's ongoing exploration into determining which functions should be offloaded and whether complete offloading consistently optimizes resource utilization, latency, and energy consumption. Deciding which network functions should be offloaded to dedicated hardware accelerators involves a comprehensive evaluation process. Firstly, assess the performance requirements of network functions, considering factors such as throughput, latency, and computational intensity. Functions demanding high performance are strong candidates for offloading. Next, analyze the resource utilization of network functions, identifying those consuming significant CPU, memory, or power resources. Such functions may benefit from offloading to specialized hardware. Evaluate the capabilities of available accelerators, including FPGAs, GPUs, and ASICs, and match them with the functionality and performance demands of network functions. Conduct a cost-benefit analysis, weighing the expenses associated with hardware accelerators against the performance gains achieved through offloading. Additionally, consider factors like flexibility, scalability, security implications, and ease of software integration when making decisions. By systematically considering these factors, network

administrators can make informed decisions to optimize performance, resource utilization, and cost-effectiveness through the strategic offloading of network functions to dedicated hardware accelerators.

2.3 The CPU vs GPU dilemma

Offloading PQC computations to either a CPU or a GPU involves harnessing distinct hardware architectures to execute cryptographic algorithms resistant to quantum attacks. CPUs, as general-purpose processors, excel in managing diverse tasks with optimized sequential processing, making them suitable for cryptographic operations necessitating intricate decision-making and branching. Conversely, GPUs specialize in parallel processing, offering significant advantages in handling PQC algorithms amenable to efficient parallelization. Their multitude of cores facilitates simultaneous execution of numerous simple tasks, potentially resulting in speed enhancements. CPUs feature a memory hierarchy that facilitates quicker access to frequently-used data through registers and cache memory. Factors such as power consumption and the specific characteristics of cryptographic algorithms inform the decision between CPU and GPU offloading for PQC, underscoring the importance of aligning hardware capabilities with the computational demands of cryptographic operations.

2.4 Algorithm election

The NIST PQC standardization process continues, with evaluation ongoing for multiple cryptographic algorithms. Among the finalists are lattice-based, hash-based, code-based, multivariate polynomial-based, and isogeny-based cryptographic schemes. Each finalist relies on a unique mathematical foundation and set of security assumptions [4]. Lattice-based schemes exploit the complexity of lattice problems, while hash-based schemes depend on the collision resistance of hash functions. Code-based schemes utilize error-correcting codes, multivariate polynomial-based schemes entail solving systems of multivariate polynomials, and isogeny-based schemes are rooted in the challenge of computing isogenies between elliptic curves. Practically, each PQC algorithm requires varying processing cycles for its sub-functions and yields keys of different sizes. This presents two challenges addressed in QUARC: determining the processing requirements for each algorithm, including whether they can be executed in a DPU, and assessing whether algorithms are suitable for different segments of the network based on their processing demands and key sizes. Initial experiments at line-rate demonstrate effective execution of PQC algorithms on DPUs, yet some may be better suited for intensive intra-rack communications, while others may be more appropriate for deep-edge or Internet-of-Things (IoT) ecosystems.

2.5 Open Programmability

The preceding sections have strongly suggested the pivotal role of DPUs in revolutionizing the implementation of PQC within expansive communication systems. DPUs, capable of serving as Data Plane Control Devices, offer computational advantages for PQC while simultaneously processing packets containing requisite data for PQC algorithms [5]. By positioning computations close to the network level, transmission latency for cryptographic parameters and data can be minimized, facilitated by their integration on the same board. DPUs are evolving to be reconfigurable both at the software and hardware levels, enhancing the flexibility and adaptability of packet processing. This evolution contributes to improved network performance in terms of bandwidth and latency. On the software front, P4 is emerging as an industry standard for packet processing, enabling programmers to define Match+Action Tables (MATs) and other necessary components for packet processing without altering the underlying hardware [6]. Runtime adaptations to traffic and packet requirements offer potential performance enhancements through runtime optimizations. Meanwhile, on the hardware side, DPUs leverage ASICs for packet processing tasks, utilizing the Disaggregated Reconfigurable Match+Action Table (dRMT) architecture to boost inter-packet concurrency compared to its predecessor, the RMT architecture. Additionally, dRMT facilitates better resource utilization through its multiprocessor design and shared memory clusters. However, packet processing performance on DPUs and SmartNICs remains non-deterministic, lacking performance guarantees. They operate in a Run-to-Completion manner and may encounter bottlenecks based on different programs and inputs, necessitating optimizations during runtime such as caching and reordering. This dynamic capability enables efficient handling and routing of mixed PQC and classical PKC traffic in Data Centers, leveraging the reconfigurability and flexibility provided by DPUs' hardware and software architectures.

2.6 Network Acceleration Solutions and Deserialization Challenges

In the preceding section, we established the inherently network-centric nature of PQC. To facilitate discussion, a generic network can be conveniently divided into far-edge, edge, and cloud segments to delineate their primary characteristics. Regarding PQC, it is pertinent to note that in scenarios with numerous connections, the overhead introduced by PQC can be substantial. For instance, a far-edge device initiating a connection to an edge computing platform might need to execute digital authentication algorithms once and a key-exchange mechanism periodically based on data transferred or elapsed time. However, within intra-data centers or high-performance computing clusters, the proliferation of I/O operations magnifies the overhead introduced by PQC processes. Consequently, there is a pressing need to expedite not only the PQC generation itself, as discussed previously, but also the general

mechanisms governing network control and management, such as the message parsing interface (MPI) or deserialization offloading.

As previously mentioned, integrating PQC into HPC systems could potentially compromise their performance. One approach to mitigate this issue is MPI acceleration: utilizing dedicated network hardware to offload specific MPI communication tasks from CPUs can minimize or even eliminate the performance penalty imposed by PQC algorithms on communication. This is particularly significant in distributed PQC implementations, where such hardware facilitates the creation of innovative RDMA operations that concurrently execute computations while retrieving data—an area of active exploration for QUARC. Additionally, this dedicated hardware offers an opportunity to reassess conventional MPI challenges, such as enhancing MPI tag matching performance through hardware offloading, which is another research focus within QUARC.

PQC involves intensive I/O processes, especially when exchanging keys across extensive networking fabrics. Consequently, deserialization offloading emerges as a crucial step in reducing latency. QUARC will investigate methods to offload the deserialization of keys and other PQC-related operations using dedicated hardware (DPUs). Leveraging RDMA hardware, we can implement rapid remote procedure call (RPC) protocols or reimagine existing protocols atop RDMA. The inherent difficulty of parallelizing deserialization can be mitigated by parallelizing the deserialization of multiple requests, utilizing nodes equipped with suitable hardware such as GPUs or DPUs. One of the primary targets for this form of acceleration is the key management system.

3. CONCLUSIONS

The convergence of high-performance computing and cloud computing driven by Artificial Intelligence has intensified pressure on data centers, irrespective of their size or role. This has prompted the exploration of the concept of a "zero-tax" data center, driving research efforts across hardware, firmware, and software stacks to maximize processing capabilities efficiently. In a networking context, strategies such as network virtualization, traffic optimization, and enhanced security measures are essential components of the zero-tax data center paradigm. Offloading networking operations from hosts to the network interface card (NIC) represents a pivotal aspect of this concept, allowing hosts to focus on core functions without incurring networking overhead. Meanwhile, the impending advent of commercial quantum computing poses significant challenges to conventional cryptographic infrastructure, necessitating the integration of quantum-resilient solutions. Post-Quantum Cryptography (PQC) addresses these challenges, with ongoing efforts focusing on algorithm evaluation, offloading PQC operations to dedicated hardware accelerators, and optimizing network acceleration solutions while addressing deserialization challenges. Through strategic integration of PQC and network acceleration techniques, data centers can navigate the complexities of emerging technologies while ensuring efficient and secure communication infrastructures.

ACKNOWLEDGEMENTS

This work was partly funded by EC-funded QUARC (101073355) and CLEVER (101097560) projects.

REFERENCES

[1] Jalal Bhayo et al. Towards a machine learning-based framework for ddos attack detection in software-defined iot sd-iot) networks. Engineering Applications of Artificial Intelligence, 123:106432, 2023.

[2] Moolchand Sharma et al. Leveraging the power of quantum computing for breaking rsa encryption. Cyber-Physical Systems, 7(2):73–92, 2021

[3] Xavier Bonnetain et al. Quantum Security Analysis of AES. IACR Transactions on Symmetric Cryptology, 2019(2):55–93, June 2019

[4] RGorjan Alagic et al. Status report on the third round of the nist post-quantum cryptography standardization process, 2022-07-05 04:07:00 202

[5] A. Cano Aguilera (2023). First end-to-end PQC protected DPU-to-DPU communications. Electronics Letters, 59(17), Article e12901. https://doi.org/10.1049/ell2.12901

[6] H. Harkous et al., "Performance-Aware Orchestration of P4-Based Heterogeneous Cloud Environments," in IEEE Transactions on Network and Service Management, vol. 20, no. 4, pp. 4765-4778, Dec. 2023, doi: 10.1109/TNSM.2023.3267983