

A Position on Security Assurance for Compute-in-the-network Cyber-Physical Systems

Saad Mubeen*, Mohammad Ashjaei*, Per Strandberg[†], Björn Leander[‡]

*Mälardalen University, Västerås, Sweden

[†]Westermo Network Technologies AB, Västerås, Sweden

[‡]ABB Process Automation, Process Control Platform, Västerås, Sweden

*{saad.mubeen, mohammad.ashjaei}@mdu.se; [†]per.strandberg@westermo.com; [‡]bjorn.leander@se.abb.com

Abstract—Industrial cyber-physical systems have traditionally separated application execution from communication infrastructure. Controllers, servers, and edge devices execute software, while network switches operated as passive communication devices that simply forwarded frames or packets based on protocol logic. Recent advances in compute-capable network devices challenge this separation. Applications and services that were previously executed on dedicated compute nodes can now partly or fully execute within the network infrastructure. Examples include protocol translation, security functions, and backup control functionality. This architectural change can reduce hardware costs, improve responsiveness, and simplify deployment. At the same time, it alters basic assumptions used during cybersecurity analysis. Network devices become execution platforms, introduce new attack paths, and create additional trust relationships. Existing security engineering practices do not adequately capture these changes. This paper argues that compute-in-the-network systems require new approaches to security assurance.

I. INTRODUCTION

Communication networks in industrial cyber-physical systems (CPS) are evolving beyond their traditional role as data transport mechanisms. Modern network devices (switches) used in these systems increasingly combine communication and computation capabilities, allowing software services and application functionality that traditionally execute on dedicated nodes to be deployed directly within the network infrastructure [1], [2]. As a result, network devices are evolving from passive data forwarding elements into active components of the system's computing platform.

Several practical use cases motivate this trend. Functions that traditionally required dedicated hardware can instead execute within network devices using, for example, containerization technology [3], [4]. Examples of such functions include firewall functionality, protocol translators, security services, and backup implementations of controller logic. This can reduce hardware costs while bringing computation closer to data sources. For time-sensitive applications, shorter communication paths may improve responsiveness [5].

While these advantages are attractive, they also change fundamental architectural assumptions. The traditional distinction between computational nodes and communication channels becomes less clear. Network devices are no longer passive infrastructure. They become active participants in the execution of CPS applications [6].

The discussion around compute-in-the-network has so far focused largely on performance, deployment flexibility, and cost reduction. We argue that a more fundamental issue remains insufficiently explored. Current security assurance methods for industrial CPS applications implicitly assume that network infrastructure is not part of the application execution platform. Compute-in-the-network invalidates this assumption by turning network devices into hosts for software and system services. As a result, established security analyses, trust models, and assurance arguments may no longer provide an adequate basis for assessing system security. **This paper takes the position that a new approach to security assurance is needed for industrial CPS that integrate computation directly within the communication infrastructure.**

II. WHY EXISTING ASSUMPTIONS NO LONGER HOLD

Security assessments for industrial CPS often treat network switches as passive infrastructure that implements data forwarding and switching functions. Security mechanisms are typically applied to endpoints and communication channels, while application logic executes on dedicated computational platforms. Compute-in-the-network architectures challenge these assumptions. Network devices may host software containers, execute critical services, process application-level information, and directly compute control decisions. As a result, the network infrastructure becomes part of the application execution platform rather than merely supporting it.

Compromising a network device can therefore have consequences beyond communication disruption. An attacker who gains control of a compute-capable switch may influence both data transport and software execution. Vulnerabilities in hosted services may create opportunities for privilege escalation, while configuration mistakes may expose sensitive functionality to unauthorized users.

Trust assumptions also become more complicated. In many existing architectures, switches are treated as communication channels without direct access to application semantics. In contrast, compute-capable switches may inspect, process, generate, or modify information. They may also host services with their own software dependencies, credentials, and access privileges. Consequently, security analyses that model network devices as passive elements may overlook important attack paths, trust relationships, and potential failure modes.

III. EMERGING SECURITY CHALLENGES

We identify several challenges that should receive greater attention from both researchers and practitioners.

1) **Protecting the Network Platform:** The network fabric itself becomes a target. If software running on a network device is compromised, an attacker may gain influence over communication services and hosted applications simultaneously. Protecting the integrity of the network platform therefore becomes a primary concern rather than a secondary consideration.

2) **Isolation Between Services:** Multiple services may execute on the same network device. A weakness in one service should not affect other hosted functions or core networking capabilities. Strong isolation mechanisms are needed to contain faults and prevent privilege escalation.

3) **Configuration Governance:** The management of compute-capable network devices introduces new operational challenges. Different users may require different levels of access for network management, application deployment, maintenance, and security oversight. Authentication, authorization, and accountability become central requirements. The situation becomes even more complicated when such infrastructure must **coexist with legacy industrial systems** that were not designed for fine-grained access control.

4) **Trust, Certificates, and Lifecycle Management:** Long-lived systems require mechanisms for credential renewal, certificate rollover, software updates, and trust management. Existing approaches may not be sufficient when network devices act as execution platforms and participate directly in security-critical operations. Questions arise regarding which certificates should be trusted, how trust relationships evolve over time, and how updates can be performed without affecting system availability.

5) **Secure Loading and Platform Integrity:** Hosted software must be protected against unauthorized modification. Secure boot mechanisms, trusted onboarding procedures, and remote attestation may become essential building blocks for establishing confidence in deployed functionality.

6) **Impact on Security Analysis:** Existing security assessments frequently treat switches as black-channel components, meaning that their internal operation is assumed to be outside the scope of application-level analysis and they are expected only to forward information between endpoints. This abstraction becomes difficult to justify when network devices can access and process application-level information. A comparison between traditional architectures and compute-in-the-network architectures reveals that threat models, trust boundaries, and attack paths change significantly. Security analyses must therefore explicitly account for software executing inside network infrastructure.

7) **Regulatory Compliance:** Emerging cybersecurity regulations place strong emphasis on security throughout the operational lifecycle of products and systems. Compute-in-the-network architectures raise new questions concerning responsibility allocation, maintenance obligations, evidence generation, and compliance assessment. As network devices become execution platforms, it becomes necessary to understand how security responsibilities are shared among infrastructure providers, system integrators, and system operators.

IV. CALL FOR NEW SECURITY ASSURANCE APPROACHES

Compute-in-the-network should not be viewed merely as a deployment optimization. It represents a fundamental architectural change with direct implications for security assurance. We argue that future work should focus on:

- architectural models that explicitly represent computation hosted within network infrastructure;
- security analysis methods that capture new trust relationships and attack paths;
- deployment techniques that provide verifiable isolation and software integrity;
- lifecycle management approaches for credentials, updates, and monitoring; and
- assurance frameworks that support cybersecurity compliance in systems that combine communication and computation within the same infrastructure.

Without these developments, the operational benefits of compute-in-the-network may be accompanied by risks that are difficult to identify, evaluate, and manage.

V. CONCLUSION

Compute-capable network infrastructure is beginning to reshape the architecture of industrial CPS. Services that previously required dedicated computing nodes can now execute directly within network devices, providing greater flexibility, reduced hardware requirements, and improved responsiveness. This shift challenges long-standing assumptions concerning trust, isolation, configuration management, and security analysis. The network is no longer merely a communication medium. It becomes part of the computing platform itself. Our position is that current security assurance practices are not yet prepared for this transition. The community has an opportunity to develop the concepts, models, and assurance techniques needed to support secure and dependable compute-in-the-network CPS before large-scale adoption becomes common.

Acknowledgments: This work is supported by the Swedish Governmental Agency for Innovation Systems (VINNOVA) via the FLEXATION project (Grant No. 2024-01731), and by the Swedish Knowledge Foundation (KKS) via the SEINE project. The authors thank the industrial partners: Arcticus, ABB, and Westermo.

REFERENCES

- [1] Vector, “Microsar switch – scalable software for switch-ecus,” *Smart Solutions for Smart Switches*, 2025, <https://www.vector.com/se/en/products/products-a-z/embedded-software/microsar-switch/>.
- [2] Westermo, “Application hosting,” *WeOS 5.26.1*, 2025, <https://docs.westermo.com/weos/weos-5/General/app-hosting.html>.
- [3] T. Goldschmidt, S. Hauck-Stattelmann, S. Malakuti, and S. Grüner, “Container-based architecture for flexible industrial control applications,” *Journal of Systems Architecture*, vol. 84, pp. 28–36, 2018.
- [4] M. Sollfrank, F. Loch, S. Denteneer, and B. Vogel-Heuser, “Evaluating docker for lightweight virtualization of distributed and time-sensitive applications in industrial automation,” *IEEE Transactions on Industrial Informatics*, vol. 17, no. 5, pp. 3566–3576, 2021.
- [5] M. Ashjaei, L. Lo Bello, M. Daneshmand, G. Patti, S. Saponara, S. Mubeen, “Time-Sensitive Networking in automotive embedded systems: State of the art and research opportunities,” *Journal of Systems Architecture*, vol. 117, pp. 102–137, 2021.
- [6] S. Mubeen, M. Ashjaei, and J. Lundbäck, “Architecting automotive software for the emerging compute-in-the-network paradigm,” in *48th IEEE/ACM International Conference on Software Engineering*, 2026.