

GRIFIN: cognitive and programmable security for resilient next-generation networks

Gregory Blanc¹, Thomas Silverston², and Sébastien Tixeul³

¹SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, gregory.blanc@telecom-sudparis.eu

²LORIA, CNRS, Inria, Université de Lorraine, thomas.silverston@loria.fr

³LIP6, CNRS, Sorbonne Université, sebastien.tixeul@lip6.fr

GRIFIN (ANR-20-CE39-0011) is a research collaborative project (PRC) funded by the French National Research Agency (ANR). It started in April 2021 and ends in September 2025. It is led by **Télécom SudParis**, and involves **Université de Lorraine** and **Sorbonne Université**. External advisory is provided by **Montimage**, a French SME, and British research centers at **University College London** and **Middlesex University London**. A prototype of a resilience framework to protect IoT networks (TRL < 4) will be developed.

I. CONTEXT AND PROBLEM STATEMENT

IoT devices' security and dependability is paramount to the safety of a number of future usages such as e-Health or Industry 4.0. Yet, IoT devices are suffering from a number of limitations that make them hard to upgrade with security functions, let alone secure them by design. We focus on damages caused by compromised devices at the network layer.

In GRIFIN, we consider a resilience framework superimposed on the IoT infrastructure, integrating machine learning (ML) and software-defined networking (SDN) technologies. First, the introduction of ML approaches to classify threats or capture intrinsic patterns of legitimate behaviour could be applied to IoT with some expected success. However, applying ML to anomaly detection incurs a number of challenges by itself, which are amplified in IoT networks.

We propose to centralize the detection at the perception level and at the network level. The Internet gateway allows the capture of IP traffic from a set of devices. *How can the gateway modulate its monitoring to balance detection and scalability? Is there a way to alleviate the load of this task?* The resulting intrusion detection system may perform well, depending on the dataset. But a more thorough evaluation is expected, in particular when dealing with deep-learning (DL) methods, which are often considered black boxes. *How can we reliably assess (ML-based) intrusion detection systems?*

Once detected, threats need to be dealt with in order to ensure the resilience of the IoT network. SDN enables our gateways to be reconfigured to enforce countermeasures. However, alerts raised by anomaly detection do not inform on the type of threat and hence on the appropriate measure. Moreover, countermeasures may also have adverse effects on the network. *How can we reason on the detected anomalies, the*

network state and the countermeasures' impact to propose the best reaction? How can we characterize these impacts?

Finally, the decoupling of the control and data planes in SDN involves an additional, overlooked, challenge: decisions taken at the control plane need to be conveyed to the data plane. *How can we ensure that what is being enforced at the data plane actually fulfills what has been decided at the control plane? How can we reduce the semantic gap between the control and the data planes?*

II. RESEARCH LOCKS

The resilience framework proposed in GRIFIN attempts to tackle the scientific and technological locks described above as follows: 1) leveraging DL methods relies on massive IoT datasets that are often hard to come by; 2) anomaly detection is only efficient if errors can be investigated, or if at least decisions leading to alerts can be interpreted, so that the human in the loop can make an informed decision. This is even more critical when the human is absent from the loop; 3) Intrusion response requires extensive domain knowledge to be effective and not have adverse impacts on critical environments, which cannot be achieved by current situation awareness; 4) SDN applications abstract away the low-level network operations and though authentication measures exist (but are never enforced), there is no formally proven way to ascertain that the execution of these applications are translated into equivalent data plane configurations.

III. RESEARCH AVENUES AND APPROACHES

A first activity focus on leveraging the IoT infrastructure to provide an edge-based yet distributed intrusion detection. Using a deep-learning approach, we will break away from legacy supervised multi-label classification based on expert features to propose a robust and transferable hybrid deep anomaly detector. A generic data-driven evaluation method will provide a more sound assessment of our detector. A second activity will aggregate alerts from distributed edge detectors and select appropriate countermeasures through a reinforcement learning based approach to characterize their impacts and a multi-objective optimization to factor in non-security requirements. A third activity leverages the increased data plane programmability to provide an automated yet verifiable countermeasure deployment.