FONDAZIONE BRUNO KESSLER
MAGAZINE

When AI helps tell real threats from noise

April 21, 2026

ARES-AI: A new approach to network security

In the world of cybersecurity, one of the most difficult challenges is not only identifying cyberattacks but distinguishing real threats from so-called “background noise.” In many network infrastructures, security systems generate thousands of alerts, many of which turn out to be false positives—i.e., reports of suspicious activity that, after verification, prove to be completely legitimate.

The ARES-AI project focuses precisely on this problem. It was developed by the Center for Cybersecurity at Fondazione Bruno Kessler (FBK) in collaboration with the company Nubee, as part of the InnovAction program, an initiative co-funded by the Ministry of Business and Made in Italy (MIMIT). The aim of the program is to build a European Digital Innovation Hub to support the digital and green transformation of companies.

The goal is clear in principle but complex in execution: **to make an existing network security probe smarter** by integrating machine learning techniques capable of filtering alerts and supporting analysts in their daily work.

From static signatures to artificial intelligence

The starting point of the project is **ARES**, a security probe designed and developed by Manuel Roccon to monitor company network traffic and identify potentially malicious activities. The system is based on a traditional approach widely used in cybersecurity: **static signature** analysis.

In this model, traffic is compared to a set of rules that describe known attack behaviors. When a rule is matched, the system generates an alert. This approach has the advantage of being deterministic and easily interpretable, but it also has structural limitations: it can accurately identify known threats, but it struggles to detect new variants or anomalous behaviors that have not yet been cataloged.

Another issue concerns the volume of alerts generated. In many cases, analysts must manually review hundreds of reports, most of which turn out to be harmless. This process consumes time and resources, increasing the operational burden on security teams.

ARES-AI was created to address this critical issue **by combining the traditional detection engine with an intelligent component capable of analyzing traffic and improving the quality of generated alerts.**

Learning “normal” network behavior

To achieve this goal, the research team developed a machine learning model based on an architecture called an **autoencoder**, which is frequently used in anomaly detection systems.Instead of directly searching for attacks, the model learns to recognize normal network behavior. When it observes a traffic flow that deviates from this behavior, it flags it as potentially abnormal. This approach proved particularly suited to the project context. During data collection, it became clear that most observed traffic is legitimate, while real attacks are relatively rare. Under these conditions, training a supervised classifier—that is, a model that learns to distinguish between benign and malicious traffic from labeled examples—is difficult because it requires a sufficient number of correctly identified attacks. **Anomaly detection techniques make it possible to take advantage of the large volume of benign traffic available**.

In practice, the system builds a model of normal company network behavior. When it observes new communications between devices, it compares them to what it has learned. If the behavior differs significantly from the norm, the system flags it as potentially abnormal.

A flexible architecture for network security

One of the most interesting aspects of the project concerns the integration between the existing system and the new machine learning component, leading the team to propose and explore two possible approaches. The existing system is the ARES network probe, which analyzes corporate network traffic and generates alerts using a signature-based detection engine—i.e., rules that identify known attack patterns.

In the first approach, called **sequential**, the machine learning model analyzes only the potential anomalies already identified by the traditional system. In this way, it can verify whether they are real attacks or false positives, significantly reducing the number of reports that must be checked manually.

In the second approach, called **parallel**, the model directly analyzes all network traffic alongside the static signature system. This allows it not only to filter false positives but also to identify attacks that the traditional system may have missed.

To coordinate information from these components, a dedicated software module called the **Prediction Collector** was developed to aggregate decisions from the various systems and produce a unified view of observed security events.

A dataset built from real traffic

A fundamental element of the project was the creation of a **dataset representative of real traffic**, collected directly from the probe installed within Nubee’s infrastructure.

Over 500 alerts were recorded during one week of observation. Analysis showed that the vast majority of these reports were related to simple IP address reputation checks, while only a small portion corresponded to real attacks or scanning attempts.

This distribution confirms a reality well known to security operators: most alerts generated by automated systems result from the so-called **background noise of the Internet**—i.e., generic and often harmless traffic originating from the global network.

To train the model, the team used both public research datasets and actual traffic collected by the probe, paying particular attention to privacy and data protection.

Reducing alerts by up to 99%

The results obtained during the experimental phase are encouraging. In tests conducted on reference datasets, the model was able to **drastically reduce the number of false positives** compared to the traditional system. In some configurations, the reduction reached **99% of generated alerts** while still maintaining the ability to detect relevant anomalies in network traffic.On data collected in Nubee's real operating environment, the system also demonstrated significant improvements, with a reduction in false positives of up to 53%.

These results indicate that integrating traditional approaches with machine learning may represent a promising path for improving the effectiveness of network security systems.

Toward increasingly adaptive defense systems

The ARES-AI project represents a first step toward a new generation of security probes capable of **dynamically adapting to the operating context in which they are deployed**.

One key element is the introduction of a **human-in-the-loop** mechanism, which allows analysts to validate alerts and use this feedback to progressively improve the machine learning model. In this way, the system can be retrained over time and gradually adapt to the specific characteristics of its environment.

The prototype developed in the project reached an intermediate technological maturity level (TRL 4) and was validated in a realistic experimental environment. The modular architecture and the tools developed—from the machine learning model to the automated test infrastructure—provide a solid foundation for future industrial development.

A collaboration between research and industry

ARES-AI demonstrates how collaboration between academic research and industry can lead to practical solutions for addressing cybersecurity challenges.

By combining FBK's scientific expertise in the field of network traffic analysis and Machine Learning with Nubee's operational knowledge in monitoring corporate infrastructures, the project has developed a prototype capable of significantly improving the efficiency of threat detection systems.

In a landscape where cyberattacks are becoming increasingly sophisticated and dynamic, tools capable of learning from the behavior of the network and adapting over time represent a fundamental piece to build smarter and more proactive defense systems.

PERMALINK

https://magazine.fbk.eu/en/news/when-ai-helps-tell-real-threats-from-noise/

TAGS

- #alert
- #anomaly detection
- #ARES-AI
- #artificialintelligence
- #attacchi informatici
- #autoencoder
- #cybersecurity
- #digital innovation
- #digital transformation
- #digitalindustry
- #green transformation
- #human-in-the-loop
- #industry
- #innovaction
- #Machine learning
- #malevolo
- #mimit
- #Nubee
- #prediction collector

AUTHORS

- Martina Cicaloni