2023

L-ADS: Live Anomaly Detection System

Datasheet OT Security

Cybersecurity Unit
Research & Innovation

EVIDEN
an atos business

Introduction

The L-ADS is a soft real-time *anomaly-based network intrusion detection* system. It monitors the network traffic in segments of an OT infrastructure and detects anomalous behaviour in such communications. It offers a comprehensive network traffic monitoring from OSI layer 2 protocols to higher application layer OT protocols in ICS/SCADA substations such as Modbus-TCP, IEC 60870-5-104, DNP3, and IEC 61850 (GOOSE/SV).

The L-ADS features a high-level automation in both deployment and operation with minimal security expertise required. It automatically learns the normal behaviour of different ICS/SCADA components, such as PLCs, RTUs, HMI and SCADA servers, and monitors any deviation from the benign traffic model. The tool is driven by a custom bi-directional flow generator and two deep learning algorithms: one for the detection of anomalies and another one for recognition of potential attacks causing the anomalies.

Detection capabilities

The L-ADS inspects and analyses metadata of network packets, and for the most popular OT protocols, it provides a deep packet inspection to the scope of protocol headers and function codes utilisation. L-ADS derives a large number of behavioural characteristics from ICS/SCADA components' communications, which empower the detection of anomalies and attacks. The anomalies represent any deviation from legitimate traffic. L-ADS provides explainability of why the anomalies detected, with useful information on critical behavioral characteristics, their value difference and direction of intrusion.

L-ADS also adopts a new type of characteristics called conformity features that, in addition to the behavioral ones, indicate how much network traffic conforms (or does not) to the different legitimate dimensions indexed during training. These features reinforce L-ADS's visibility on intrusions on border areas (between legit and anomalous) where advanced or stealthy attacks mostly escape or hide behind. Importantly, the L-ADS conformity monitoring ranges from monitoring whitelist IP (v4/v6) and MAC addresses to finer-grained monitoring of traffic conformity between legitimate pairs of ICS/SCADA components' communications per protocol, per protocol and day of week, and even per working hours a day. All conformity info is automatically indexed and configured during training.

The L-ADS offers custom deep leaning models and extended OT protocol monitoring for the most popular OT protocols Modbus, IEC-104, and DNP3. It uses its own knowledge database of OT attacks based on several ICS/SCADA intrusion detection datasets published from renown in the field organisations. It alerts on anomalies even when new, "zero-day" attacks take place. Whether an attacker performs a network reconnaissance or a lateral movement, the L-ADS will detect and alert such anomalous traffic.

Our detector has been validated in several pilot OT environments (e.g., a substation of distribution power, and PV plant) under EU co-funded projects, but also on several intrusion detection datasets. A summary of LADS's detection capability is listed below.

Generic attacks (IT protocols)	Man in the Middle (MitM)	
	Packet Injection (forging/spoofing)	
	Unauthorised Host Access	
	Reconnaissance	Network Scanning, Port Scanning

	Denial of Service (DoS/DDoS)	TCP SYN Flood, ICMP Ping Flood, Blackholing
		Protocol specific
OT protocol specific attacks: IEC104, DNP3, Modbus	Unauthorised Activity	Unauthorised Function Code use or misuse
	Reconnaissance	Enumerate / Info attack, Function Code Scanning
	DoS	Protocol/Function Code Level DoS such as selective Blackholing, Packet drop, (low rate) excess num. of Headers/Function Code payloads.
	Replay attack	OT Protocol Packets Replay

Setup and operation

In contrast with classic rule-based intrusion detection tools, the L-ADS does not require creation of customized rules, load specific rulesets or complex configurations for getting it started. Its setup process is much simpler and automatic than most classic IDSs.

Step 1 – Configuration

Before running the L-ADS for the first time, some minimal configuration is required, like configuring the range of IP/MAC addresses to be monitored, activating an extended inspection of specific OT protocols, and choosing the desired processing scalability in terms of CPUs and RAM.

Step 2 – Training

Once configured, the L-ADS automatically learns the network behaviour through a training process. It learns the communications and traffic footprint variations of the different OT devices and systems for a representative period. When done, it generates an AI deep learning *model* that represents the legitimate network behaviour of the OT environment.

Step 3 – Monitoring

Once trained, the L-ADS automatically switches to monitoring mode. It uses the trained AI model for the anomaly detection in real time traffic inspection. Any network activity that differs from the modelled one (learnt during the training phase) generates an anomaly that triggers a security event by L-ADS.

For each detected anomaly, and in addition to the explainability, the L-ADS identifies the likelihood of an attack causing the anomaly. It uses the attack classifier with a pre-trained AI model of common attacks on OT protocol communications (refer table above). The classifier AI models are shipped with L-ADS installation or updated from our server.

Step 4 – Reporting

The L-ADS extends the detected anomalies with an explanatory information of why a flow is detected as anomaly, identifying the most relevant network behaviour features causing the anomaly, and justifying their value difference comparing the values of training and monitoring.

As a result, it generates a **security event log** of the detected anomalies (in JSON or CSV format) that can be used by other tools, like SIEM, CTI, or any risk contingency and decision-support systems.

Step 5 – Retraining

L-ADS is designed to **adapt** to changes in an OT environment by offering scheduled retraining of its AI models to incorporate new legitimate behaviour, or through human-in-the-loop cognitive feedback on anomalies that represent benign traffic.

Deployment

The L-ADS is provided as a Docker container with a Compose file to ease its deployment and configuration. The L-ADS can be scaled to more complex infrastructures with segmented networks by deploying multiple instances at critical (e.g., gateway) points such as those of field or telecontrol LANs. The following diagram shows the typical deployment view of the L-ADS in an OT environment:

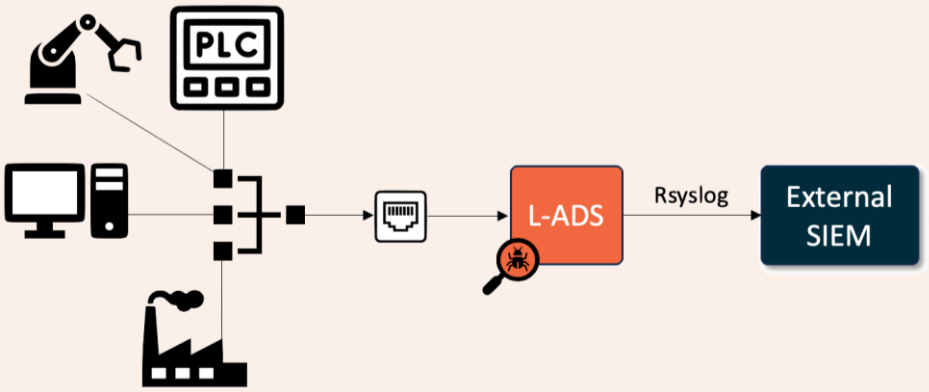

Performance & requirements

The L-ADS performance depends on the performance of the hosted server and the available resources. However, the following resources are recommended for running the L-ADS stably:

Requirements	Min	Recommended
RAM	4 GB	6 GB
vCPUs	4	6
HDD	30 GB+	60 GB+

Using the recommended resources, the L-ADS is able to process and analyse up to **5500 data flows per second** in real-time.

Additionally, the L-ADS only requires to be deployed in a system with **Docker** installed and VT-x enabled, within a dedicated network interface where real-time network traffic is mirrored to.