PALANTIR demo: leveraging SecaaS model for managing threats in industrial environments

Carolina Fernández[1], Davide Sanvito[2], Orestis Kompougias[3], Valentino Šafran[4], Maxime Compastié[1], Ignazio Pedone[5], Antonio López Martínez[6], Manuel Gil Pérez[6], Akis Kourtis[7], George Xylouris[7], Izidor Mlakar[8], Stylianos Tsarsitalidis[9], Dimitrios Klonidis[9], Daniele Canavese[5], Vangelis Logothetis[10], Diego Lopez[11], Antonio Pastor[11], Antonio Lioy[5], Ludovic Jacquin[12], Supreshna Gurung[12], Roberto Bifulco[2], Athanasios Priovolos[13], Ilias Balampanis[13], Theodoros Rokkas[10], Nikolaos Papadakis[14], Dimitris Paraschos[14], Primoz Jeran[8], George Athanasiou[15], Dimitris Papadopoulos[3]

[1]i2CAT Foundation, Barcelona, Catalonia, Spain — [2]NEC Laboratories Europe GmbH, Heidelberg, Germany
[3]Infili Technologies SA, Athens, Greece — [4]University of Maribor, Maribor, Slovenia — [5]Politecnico di Torino, Turin, Italy — [6]University of Murcia, Murcia, Spain — [7]ORION Innovations PC, Athens, Greece
[8]Sfera IT, Maribor, Slovenia — [9]UBITECH Ubiquitous Solutions, Athens, Greece — [10]inCITES Consulting SA, Strassen, Luxembourg — [11]Telefonica I+D, Madrid, Spain — [12]Hewlett Packard Enterprise, Bristol, United Kingdom
[13]Space Hellas, Athens, Greece — [14]Stratotiki Sxoli Evelpidon, Vari, Greece — [15]DBC Europe SA, Brussels, Belgium

***Abstract*—PALANTIR, an EU-funded Innovation Action project, delivers tailored and pervasive resource protection through the Security-as-a-Service paradigm. This demo presents the architecture and validates the prototype against three threat scenarios with botnet, data breach and tampering attacks.**

I. Overview

Cybersecurity is increasingly important to the economy of societies around the world. This is evidenced from both the victims' and attacker's side: an estimated global expenditure of $170.4 billion by 2022 to protect organisations [1] and the evolving complexity of exploits and solutions used by cybercriminals to obtain revenue from their malicious activities.

PALANTIR intends to provide affordable Security-as-a-Service (SecaaS) solutions to smaller organisations. SecaaS solutions are special security services, named Security Capabilities (SCs). These are instantiated with containers (Docker, Kubernetes) and orchestrated through frameworks like Open Source MANO (OSM); whilst considering different deployment scenarios depending on the resource constraints.

In addition, PALANTIR offers different centralised and decentralised components, such as i) the *Threat Intelligence* (TI), relying on sophisticated Artificial Intelligence (AI) techniques and processes to efficiently detect specific threats and activities, and ii) the *Fault and Breach Management* (FBM), performing the recovering process once the threats affect clients' assets and procedures. As a whole, PALANTIR can deploy in three delivery modes to approach light, cloud and edge security solutions for adapting the infrastructure and intelligence capacity to the end-user requirements.

II. Innovation

Both Small and Medium-sized Enterprises (SME) and Micro-Enterprises (ME) constitute the backbone of most countries' economies, representing over 90% of the business population, 60-70% of employment, and 55% of the GDP in developed countries [2]. However, they usually lag behind larger organisations in the adoption of more sophisticated security practices, now facing major security challenges due to their low investment in cyberesecurity tools, lack of cyber-skills and relevant best practices aimed at protecting their assets and the value-chain they are connected to. This, along with the uptake of digital tools (especially after the dawn of the work-from-home era), has rendered SMEs/MEs an attractive target for a wide range of cyberattacks. As a result, they usually fall victims to cheap, replicable attacks by cybercriminals leveraging "ransomware-as-a-service" as a means to extort considerable ransoms from SMEs/MEs at a low risk; while sometimes even acting as the weakest nodes in supply chains and providing backdoor routes into larger organisations.

PALANTIR aims to bridge the cybersecurity gap between large enterprises and SMEs/MEs by providing a multi-layered, infrastructure-wide threat monitoring platform offering cyber-resiliency and knowledge sharing in a heterogeneous ecosystem. At the same time, these security services are marketed to third parties in the form of SecaaS [3]. The PALANTIR framework implements a coherent privacy assurance, data protection, incident detection and recovery framework, focusing on highly dynamic service-oriented systems and networks; taking advantage of their inherent programmability features and abstractions. PALANTIR also focuses on cyber-resiliency, leveraging the features of service-oriented systems by i) exploiting the Network Functions Virtualisation (NFV) architecture; ii) employing scalable AI for multi-modal threat detection, standardisation and threat-sharing techniques for risk analysis, monitoring and network operation; and iii) ensuring the SME/ME's compliance with relevant data privacy and protection regulations in the data breach age, implementing the "Privacy by Default" principle to the collection, usage, transfer and storage of personal data between third parties.

 DOI: 10.1109/MeditCom55741.2022.9928648

While the individual PALANTIR assets bring significant innovation potential to threat intelligence, risk assessment, cybersecurity response and service orchestration, the core innovation capacity of the platform stems from the efficient combination of the above, setting it apart from existing commercial offerings. Specifically, PALANTIR combines innovative aspects that distinguish it from the competition, including:

1) an extended value chain among stakeholders (service developers, telcos, SMEs/MEs, GDPR subjects) materialised via three delivery modes (light/cloud/edge) and multiple billing options to accommodate different needs;
2) a service catalogue that allows a commercial ecosystem to be built around the aforementioned value chain;
3) hybrid cybersecurity and active response services, combining approaches based on multi-modal machine learning and signatures, which are tailored to discover cyber threats and complemented with live threat sharing;
4) a service-oriented risk analysis framework that removes the need for expert users to operate complex tools for the identification of attack surfaces and security risks, leading to the reduction of operational costs;
5) the provision of security-based orchestration to simplify the deployment of security capabilities based on end-user feedback and identified cybersecurity needs; and
6) the alignment with ongoing standardisation efforts by IETF [4] & ETSI [5].

III. Relevance for MeditCom

The scope of the demo is directly relevant to several research areas identified in the IEEE MeditCom call for demos, including i) 5G/6G mobile systems; ii) network architectures; iii) Big Data and Machine Learning for communications; iv) network applications and services; and v) cybersecurity, privacy and trust. We believe the audience should be particularly interested in this demo for two main reasons: i) the innovative nature and technical merit of the PALANTIR SecaaS framework integrating NFV, AI and Trusted Computing into a novel enabler platform for SecaaS; bringing a direct socio-economic impact via affordable and easy-to-use cybersecurity for SME/MEs; and ii) the chance to learn about relevant open-source solutions and their integration, e.g., coupling the Elastic stack with Kafka and Spark (for data indexing, analytics and visualisation) and OSM (for SecaaS orchestration).

IV. Related work

The provided demonstration is based on the PALANTIR architecture [3]. Some of the investigated techniques for the storylines are described in previous works on e.g., anomaly detection [6] [7] and botnet detection [8] [9]; where the techniques from the latter are currently in use along with a refined dataset. The SecaaS approach, as implemented by the lightweight delivery mode, is described in depth in [10] along with the internal design of the SCs. It is also applied in the context of 5G cybersecurity [11] and energy awareness [12].

V. Demo Summary

The PALANTIR framework aims to simplify the protection for SMEs and MEs, alleviating to some extent their evident expertise gap on cybersecurity knowledge with respect to medium-sized and large enterprises. PALANTIR follows a SecaaS approach under three delivery modes to better tailor its behaviour to the capacity and needs of each organisation: i) the cloud-based approach, based on a remotely hosted Managed Security Services model (MSS); the ii) lightweight, all-in-a-box approach implementing a virtual Customer Premises Equipment (vCPE) MSS model that is deployed in the client's premises; and the iii) managed edge, adhering to the Multi-access Edge Computing (MEC) paradigm and managed by the network operator serving connectivity to the organisation.

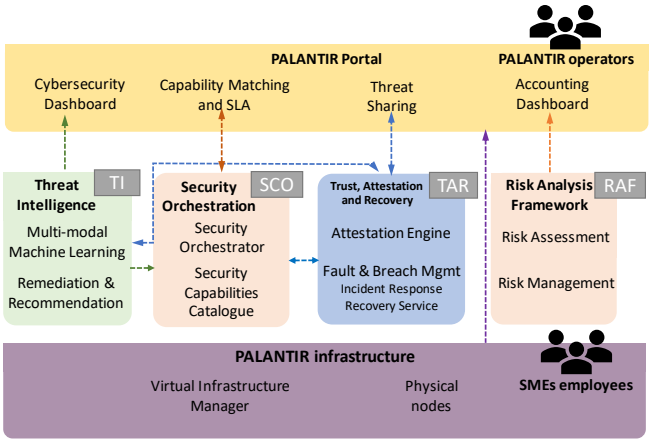

Fig. 1. Building blocks of the PALANTIR architecture

To that end, the PALANTIR architecture considers four building blocks (Figure 1). From top to bottom and left to right, these are the i) Portal; ii) Threat Intelligence; iii) Security Capabilities Orchestration; iv) Trust, Attestation and Recovery; and v) Risk Analysis Framework components. The infrastructure hosts the SCs, key for ad-hoc deployment of security features such as Intrusion Detection Systems (IDS), Network Network Sniffers (NNS), Firewall and Router (FW) functions, among others (e.g., for traffic analysis and backups).

The Portal consists of multiple tailored views of dashboards and provides the user with sections related to security indications and status, sharing Indicators of Compromise (IoC) and monitoring and billing. The TI component implements Machine Learning (ML) and Deep Learning (DL) techniques to identify threat, ingesting data both from traditional signature-based detection and advanced ML/DL techniques. The Security Capabilities Orchestration (SCO) keeps the SCs in a searchable catalogue and manages their life-cycle, configuration and monitoring. The Trust, Attestation and Recovery (TAR) continuously assesses a subset of physical or software nodes via remote attestation [13]. Upon integrity breaches, adequate policies are applied to restore the system to its last trusted state. Finally, the Risk Analysis Framework (RAF) allows an organisation to perform a self-assessment (related

to the ENISA SME framework [14]) to identify the risks in their infrastructure and take adequate measures.

Three storylines are considered for demonstrating the capabilities of the PALANTIR platform. The first one aims at detecting botnet activity in the monitored network and the mitigation of the attack via the isolation of the infected machine(s). The second one attempts to detect a data breach through the inspection and ingestion of system logs. Finally, the third one showcases the remote attestation procedures and the recovery of a previously compromised security service.

a) Storyline 1 (botnet activities): Involves the RAF, SCO, TI and Dashboard. Figure 2 reports all involved subcomponents and their interactions in the different steps.

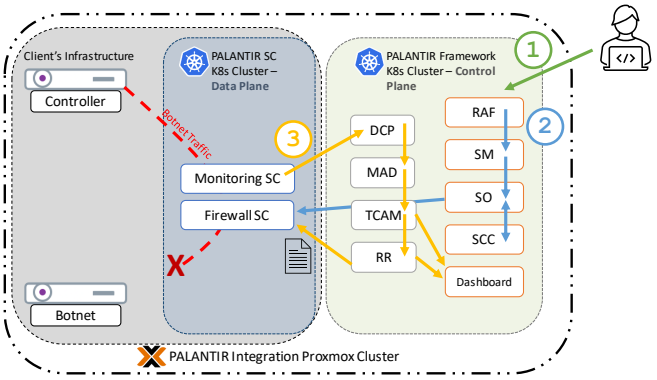

Fig. 2. Workflow for storyline #1

In the first step, the RAF provides the end-user with questionnaires to perform a risk-based assessment, so to analyse the attack surface and identify the assets part of the infrastructure. After this, the user is presented a risk profile and recommended remediation actions (in this demo, deploying a Filtering SC).

In the second step, the Service Matching (SM) subcomponent receives the request for deploying a SC providing a firewalling feature. Throughout a Constraint Satisfaction Problem (CSP), it analyses the properties of the infrastructure, the implementations of the SCs available from the catalogue and their costs to find a compromise between the billing models and the security features to be leveraged. The SM returns a possible solution associating specific SCs to use, the infrastructures to deploy them, and the delivery models to apply along with a price quotation to be reviewed by an operator. Finally, an effective deployment is triggered and the Filtering SC is instantiated by the orchestrator, as confirmed by a notification on the Dashboard.

In the third step, botnet traffic from the USTC-TFC2016 dataset [15] is replayed into the simulated network along with normal traffic. A ping command from the Controller host shows in real time the connectivity towards the victim. The Monitoring SC forwards the traffic to the Data Collection & Pre-processing (DCP) module which pre-processes it (e.g. IP anonymisation) and passes it to the Multi-modal Anomaly Detection (MAD) module. MAD includes a set of Anomaly Detection algorithms running in parallel (Isolation Forest and MIDAS [16]) and analysing traffic's flow-level features. The detected outliers are aggregated and passed to the Threat Classification and Alarm Management (TCAM) module to associate them with a specific threat label using a supervised Random Forest model. TCAM provides an attack report to the Recommendation and Remediation (RR) engine which uses it, together with the current network landscape, to compute a remediation measure for the detected threat. In the case of botnets, it suggests a reconfiguration of the Filtering SC by adding iptables rules to block the malicious traffic. The user is informed in real time through the Dashboard both when the threat is detected and once the generated remediation has been correctly applied. Finally, the interruption of the ping output acknowledges that the attacker can no longer reach the victim.

b) Storyline 2 (data breach detection): Involves the TI, TAR and the Dashboard so to monitor the protected infrastructure and detecting a data breach attempt. Figure 3 presents the involved subcomponents and their interactions.

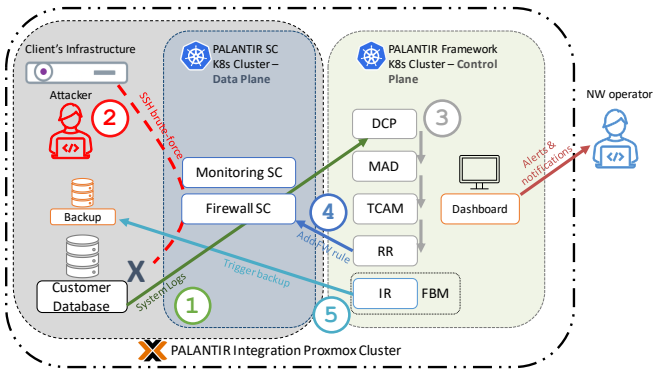

Fig. 3. Workflow for storyline #2

System logs with sensitive data are periodically transferred from a customer server to DCP. A live SSH brute-force attack (simulating a data breach attempt) runs against a database server protected by PALANTIR. Then, TI pre-processes the syslogs in real time, detects the anomalous behaviour (MAD) and classifies it as a brute-force attack (TCAM).

RR proposes two policies based on the threat findings: i) apply a FW rule to isolate the server, and ii) perform an instant backup of the victim server's data. A running SC (firewall) is reconfigured with a newly added rule based on the RR policy to block the brute-force attack.

Shortly after, the Incident Response (IR) subcomponent within FBM successfully backups the SQL dump to a secure node, following the RR policy. In the end, the user is informed of the data breach attack and the remediation steps taken.

c) Storyline 3 (remote attestation and recovery): In this scenario the Dashboard, SCO and TAR interact to deploy, attest and redeploy (upon integrity breach) an IDS-based SC implementing Snort; as shown in Figure 4.

First, the developer implements the logic of the SC (Snort) and proceeds registering (or onboarding) it in the Security Capabilities Catalogue (SCC), within SCO. Here, relevant

metadata is filled by the developer and persisted in the SCC, which is later used to identify the optimal SC to deploy such that it fulfils all security requirements.

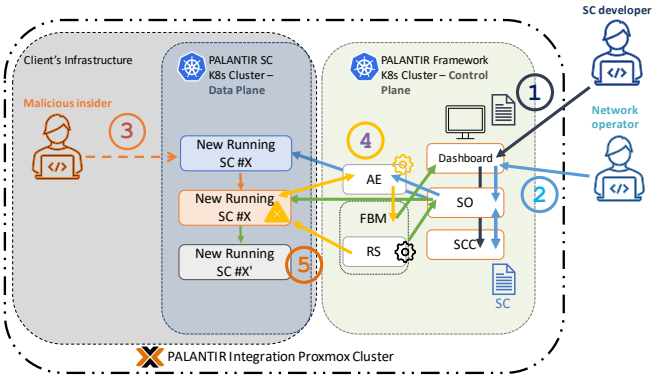

Fig. 4. Workflow for storyline #3

As a second step, and once the SC package(s) are registered into SCC, the Security Orchestrator (SO) onboards them on the NFV orchestrator. The SC is now ready for deployment. At some point, the network operator accesses the Cybersecurity Dashboard and instantiates a given SC. A few seconds later, the SC spins up at the Kubernetes cluster. The SO fetches relevant container runtime data to send to the AE, which records the SC to start the periodic remote attestation.

In the third step a malicious insider accesses the physical infrastructure, the Kubernetes cluster and, finally, the container that runs the IDS SC (Snort). This user tampers a binary in the container (e.g., disguising malware as an OS-provided binary).

Few moments after, in the fourth step, AE runs the periodic remote attestation process and identifies an integrity failure on this SC instance. It marks the container as untrusted, notifies the Dashboard and informs the Recovery Service (RS), which generates a relevant policy that requests the SO to terminate the compromised Snort-related container and to redeploy a new clean instance of the same type.

Finally, in the fifth step, a clean Snort version instance is redeployed, notifying the network operator of the successful recovery and restore to a integrity-compliant state.

VI. Demo processing

The demo can be presented either live or commenting on top of pre-recorded videos, where the latter is expected. It consists of three scenarios, each commented by two contributors. Slides are first presented to introduce the environment and kind of attack. Then, either the video, with embedded subtitles and live comments on top, or the live experiment takes place. Finally, video material could be later uploaded publicly for future consultation by attendees.

VII. Material

A table, a poster, two screens, access to power supply and stable Internet connection are required for the demonstration.

VIII. Acknowledgement

This work received funding by the European Union Horizon 2020 research and innovation programme, supported under Grant Agreement no. 883335. The NFV-related work is also supported by the Spanish Government under Grant PID2020-112675RB-C43. The content of this article does not reflect the official opinion of the European Union or any other institution. Responsibility for the information and views expressed therein lies entirely with the authors.

References

[1] Fortinet, Cybersecurity statistics, https://www.fortinet.com/resources/cyberglossary/cybersecurity-statistics/, [Accessed 09-Jun-2022] (2022).

[2] ENISA, Cybersecurity for SMEs - Challenges and recommendations, https://www.enisa.europa.eu/publications/enisa-report-cybersecurity-for-smes, [Accessed 09-Jun-2022].

[3] E. Mantas, D. Papadopoulos, C. Fernández, N. Ortiz, M. Compastié, A. López Martínez, M. Gil Pérez, A. Kourtis, G. Xylouris, I. Mlakar, S. Tsarsitalidis, D. Klonidis, I. Pedone, D. Canavese, G. Martínez Pérez, D. Sanvito, V. Logothetis, D. Lopez, A. Pastor, A. Lioy, L. Jacquin, R. Bifulco, A. Kapodistria, A. Priovolos, G. Gardikis, I. Neokosmidis, T. Rokkas, N. Papadakis, D. Paraschos, P. Jeran, A. Litke, G. Athanasiou, Practical autonomous cyberhealth for resilient micro, small and medium-sized enterprises, in: 2021 IEEE International Mediterranean Conference on Communications and Networking (MeditCom), 2021, pp. 500–505. doi:10.1109/MeditCom49071.2021.9647609.

[4] Internet Engineering Task Force, https://www.ietf.org/, [Accessed 09-Jun-2022].

[5] European Telecommunications Standards Institute, https://www.etsi.org/, [Accessed 09-Jun-2022].

[6] A. Priovolos, D. Lioprasitis, G. Gardikis, S. Costicoglou, Using anomaly detection techniques for securing 5G infrastructure and applications, in: 2021 IEEE International Mediterranean Conference on Communications and Networking (MeditCom), 2021, pp. 519–524. doi:10.1109/MeditCom49071.2021.9647668.

[7] D. Sanvito, G. Siracusano, S. Santhanam, R. Gonzalez, R. Bifulco, Syslrn: Learning what to monitor for efficient anomaly detection, in: 2nd European Workshop on Machine Learning and Systems (EuroMLSys), ACM, 2022, p. 64–71. doi:10.1145/3517207.3526979.

[8] J. T. Martínez Garre, M. Gil Pérez, A. Ruiz-Martínez, A novel machine learning-based approach for the detection of SSH botnet infection, Future Generation Computer Systems 115 (2021) 387–396. doi:10.1016/j.future.2020.09.004.

[9] O. Kompougias, D. Papadopoulos, E. Mantas, A. Litke, N. Papadakis, D. Paraschos, A. Kourtis, G. Xylouris, Iot botnet detection on flow data using autoencoders, in: 2021 IEEE International Mediterranean Conference on Communications and Networking (MeditCom), 2021, pp. 506–511. doi:10.1109/MeditCom49071.2021.9647639.

[10] A. Martínez López, M. Zago, M. Gil Pérez, Provision of Security-as-a-Service (SecaaS) in lightweight scenarios, in: 2022 Spanish Cybersecurity Research Conference (JNIC), 2022, pp. 327–330, https://2022.jnic.es/Actas_JNIC_2022_v11.pdf, [Accessed 14-Jul-2022].

[11] C.-M. Mathas, C. Vassilakis, N. Kolokotronis, C. C. Zarakovitis, M.-A. Kourtis, On the design of iot security: Analysis of software vulnerabilities for smart grids, Energies 14 (10) (2021). doi:10.3390/en14102818.

[12] I. P. Chochliouros, M.-A. Kourtis, A. S. Spiliopoulou, P. Lazaridis, Z. Zaharis, C. Zarakovitis, A. Kourtis, Energy efficiency concerns and trends in future 5g network infrastructures, Energies 14 (17) (2021). doi:10.3390/en14175392.

[13] M. D. Benedictis, A. Lioy, Integrity verification of Docker containers for a lightweight cloud environment, Elsevier Future Generation Computer Systems 97 (2019) 236–246. doi:10.1016/j.future.2019.02.026.

[14] ENISA, ENISA risk management approach for SME/MEs. (2009).

[15] W. Wang, M. Zhu, X. Zeng, X. Ye, Y. Sheng, Malware traffic classification using convolutional neural network for representation learning, in: 2017 International Conference on Information Networking (ICOIN), 2017, pp. 712–717.

[16] S. Bhatia, B. Hooi, M. Yoon, K. Shin, C. Faloutsos, Midas: Microcluster-based detector of anomalies in edge streams, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34, 2020, pp. 3242–3249.