Ammune DDoSM tool for Anomaly Detection & DDoS Mitigation

Whitepaper

2022

Authored by:
Yisrael Gross,
Rajarajan Muttukrishnan

ENERGY SHIELD

Ammune DDoSM Tool for Anomaly Detection & DDoS Mitigation

IN A NUTSHELL

L7 Defense's AmmuneTM provides effective DDoS mitigation for the smart grid as part of a package of tools to address smart meter botnets and attacks using the AMI (Amazon Machine Image) as a vector. In parallel, the "FC-DDoS" analytical model has been developed to understand attack parameters and explore mitigation dynamics unique to the smart grid context. Both were explored using expanded network simulations for testing and validation, and the results have verified the effectiveness of this approach. Meanwhile, the DDoSM (Denial of Service mitigation) has been successfully integrated with the security information and event management (SIEM) and other tools for collaborative protection. Following the integration and enhancements made to Ammune, the DDoSM tool is ready for deployment in the field test environments.

> ***"Our Ammune™ solution automatically protects the APIs of Smart Grids from the growing risk of DDoS attacks that can be launched from Smart Meters and can cause a critical cascading impact on the energy sector"***
>
> ***Yisrael Gross, Co-Founder & VP of Business Development at L7 Defense***

CONTEXT

The smart grid consists of information technology (IT) and operational technology (OT) networks that work together for dynamic monitoring and control. This entails that those failures in the IT network and disruption to IT data flows can lead to widespread disruption.

DDoS attacks are aimed to degrade availability. As IT component operations rely on Application Programming Interfaces (APIs) for communication, the latter need to be protected from this type of threat with the DDoSM tool.

This DDoSM tool consists of Ammune, which is placed in front of vulnerable IT components to protect from incoming cyber-

attacks. Ammune's AI-driven system is bolstered by analytical models to capture the key components of a DDoS attack. It also captures the dynamics of a spreading wave of threat vectors that is triggered by the attack throughout the grid networks. Ammune AI builds a dynamic protection profile and policy for all active API endpoints through in/out traffic inspection.

The botnet DDoS activity may simulate an external botnet attack targeting the SCADA or a DDoS attack created within the AMI from corrupted smart meters. These two scenarios are depicted in Figure 1. As shown, Ammune detects and stops such DDoS attacks in near real-time at their entry point

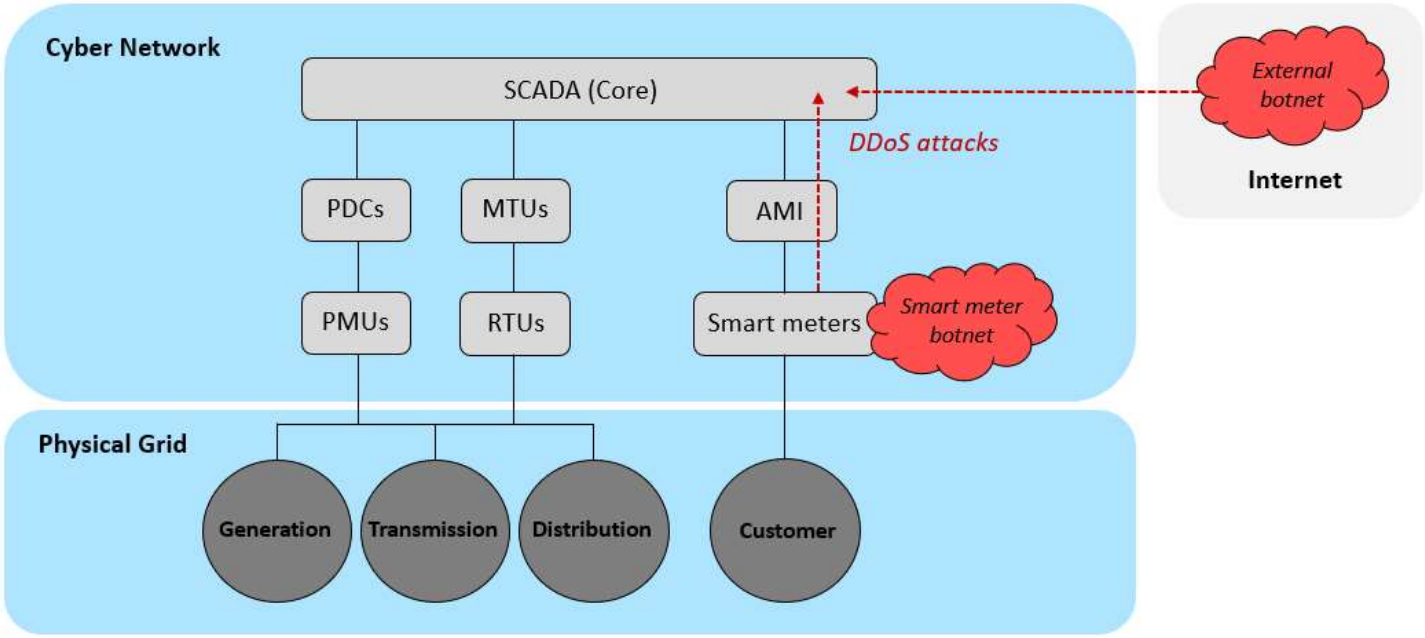

TECHNICAL DETAILS

Methodology used. City, University of London based its analytical modeling approach on epidemiological principles. This methodology was used to model the smart grid components activity alongside applied DDoS attack dynamics. These models provide insight into the impact of timings, durations, and many more. It also looks at the influence of simulated DDoS attacks on the settings, deployment, and analysis capabilities of the Ammune tool.

Features. L7 Defense's Ammune™ API security solution is an advanced AI solution based on unsupervised machine learning (ML). It protects corporate APIs from advanced attack types, with a minimal impact on the legitimate traffic. It automatically discovers and protects each API separately. Ammune continuously builds a specific profile (an AI baseline) for each API, which is used to spot and stop emerging threats that otherwise go unnoticed, in real-time, and without any prior knowledge or signatures of the attack characteristics. It is inspired by the

"innate immune system" model, designed for accuracy, minimizes the damage from both erroneous detections (false positives) and from incoming attack penetration (false negatives).

Ammune contains the following functional modules:

- API-WAF protects from content injection threats targeting remote command execution, data exfiltration, denial of service, and more.
- API-BOT protects from advanced automated threats that implement data exfiltration, fraud actions, account takeover, functionality abuse, and more.
- API-DDoS protects from DDoS attacks on API business logic that attempt to overload computation and memory resources of the application servers.
- API-BL protects APIs from business-level exploits, such as authorization and authentication bypasses.

The Ammune architecture consists of a real-time traffic enforcement unit and an analytics unit that provides near real-time analytics of the traffic flows. Ammune supports various embedding architectures, including:

- Network test access point (TAP) - Copy of the traffic is received by Ammune Real-Time module directly from network TAP without reverse proxy). The module blocks command that could be sent to other enforcing devices.
- Log feed - Copy of traffic that could be received from other sources, such as log feeds in the security information and SIEM.
- Integration with Kubernetes ingress - Ammune integrates with ingress (reverse proxy-based) instead of a standard reverse proxy.
- Inline integration - Ammune integrates in front of the customer's applicative architecture or between two hops in the flow.

As far as integration goes, alerts and security incidents generated by Ammune can be forwarded to other tools such as SIEM, security orchestration, automation and response (SOAR), or ticketing systems. Furthermore, a special integration pattern was added to support the Energy Shield project. Ammune can be integrated with API gateways to receive traffic log feed for analysis. Log feeds could also be fetched from other sources such as SIEM or weblogs. Ammune can also integrate with packet brokers as a traffic feed source, where a special adapter module will extract logs and forward these to Ammune's main engine. Ammune also contains a rich UI interface to control network flow, configurations, and enforcement policies by the user.

ENERGYSHIELD DEMONSTRATOR

For the project, Ammune algorithms were retrained and recalibrated to handle smart grid scenarios. For instance, Ammune DDoS detection was adjusted to better handle extra heavy-duty calls. It can now detect DDoS detection successfully at rates above one request per second. Also, a significant weight to calls returning error was added to increase the sensitivity for capturing attacks that overload the system with calls to non-existent objects to bypass cache. The Ammune analysis time slot was reduced from 5 to 1 second to reduce the attack mitigation time to under 1 minute.

About the simulation

DDoS attacks were initiated and generated by sending API calls via randomly selected proxies at a relatively short timeout to maximize the attack rate at around 400 requests per second. Three types of attack scenarios were conducted. The first one was aimed at the "meter update" API endpoints where the smart meter id and meter reading parameters were randomly selected to add extra load on the server. The second one consisted of attacks aimed at the "read region power consumption" API endpoints. The third one combined the two scenarios.

A simulated botnet was launched to mimic a DDoS attack. When the server is overwhelmed, aggregators are impacted, and subsequently, smart meters are impacted. The simulation environment was used for testing and validation activities. The IA-DDoS model was deployed to capture population fluctuations in a DDoS-enabled botnet. The IA-DDoS model successfully passed testing and validation against L7 Defense's simulations.

Also, the Secured Authentication Communication (SAC) model was deployed to test the possibility of using population-level observation of the smart grid and its component networks and/or systems in the context of DDoS impact propagation. Unlike IA-DDoS, the SAC model also splits the overall smart grid into subpopulations to check the impact that one subpopulation could have on another, given any dependencies existing between them. The SAC model was successfully validated using parametric testing and numerical simulation.

The FC-DDoS model was analyzed using numerical simulations and tested under varying conditions. For the DDoS module, it was observed that the malicious stream consumed a rising number of resources as the arrival rate increased. Also, the duration of the attack increased the period of disruption. For later attack end times, a downward recovery slope indicated that more damage occurred. The arrival rate was mitigated by blocking the attack as soon as possible.

Ammune DDoS attack discovery was made by analyzing the "distance" of incoming request flow from the generic and the business logic (BL) profiles. Both profiles are implemented at a

single API endpoint, entity, and multi-entity (campaign) levels. For distance regarding the general profile, the API-based anomaly detection consisted of weighing the anomaly of the specific API endpoint call rate by its reply complexity. For the API-based anomaly detection, the anomalies in the context of the API endpoint calls, such as time intervals between calls and unexpected calls sequence, were detected.

For the distance regarding the BL profile, the API-based anomaly and session detections consisted of weighing each anomaly at specific API call content by its campaign discovery consisted of a group of vectors or sessions that share similar anomalies.

Results . Regarding attack scenario 1, flooding the server with bogus smart meter update requests, where smart meter ID and reading is selected randomly, Ammune was able to perform efficient mitigation in 30 seconds from the attack initiation. As the botnet sources were rotating, Ammune was able to update its mitigation policy on-the-fly, without any further degradation of service.

Regarding attack scenario 2, flooding the server with read region power consumption (heavy requests), Ammune started its efficient mitigation within 30 seconds from the start of the attack. As the botnet sources were rotating, Ammune was able to update its mitigation policy on-the-fly, without any further degradation of service.

Regarding attack scenario 3, the combination of attack scenarios 1 and 2, Ammune started to mitigate the attack after 30 seconds from its initiation, which is the experience “set-up time” for a

visible impact of the attack on an API activity, under the simulation conditions. As the botnet sources were rotating, Ammune was able to update its mitigation policy on-the-fly, without any further degradation of service. Figure 2 shows a screenshot of the Ammune dashboard. It displays the summary, live activity, detected malicious IPs (in this case, 986), domain statistics, and the latest notifications of malicious bot activity with details, date, and time.

Using generic Ammune capabilities and the novel Smart Meter business logic implementation model, an efficient anti-DDoS solution against realistic DDoS attack simulation was provided. Ammune response restored the service activity within 30 seconds from the start of the attack, thus preventing long-term damages. The applied IP rotation attacking tactics did not affect the results as Ammune captured new source IPs and blocked them almost immediately. Apart from a few short service degradations, normal traffic was not affected with false positive (false blocks) kept at zero in these simulations. Although a few "fresh" bots were not immediately identified during the simulation, this would not happen in reality, where new bots accumulate incriminatory evidence much faster.

BEST PRACTICES & LESSONS LEARNED

DDoS remains an effective attack vector for threat actors that smart grid networks are susceptible to. Disruptions triggered by successful DDoS attacks can disturb smart grid processes that can subsequently cause imbalance and desynchronization and where the impact is allowed to accumulate and develop. This is because of the tight interconnection of both IT and operation devices, which function in tandem to serve common goals. To protect the smart grid, incoming DDoS attacks must be blocked or mitigated within the shortest amount of time.

The DDoSM tool achieves this through Ammune's immediate dynamic response which adapts itself to the attack dynamic, which is highly stochastic in nature. The Ammune AI engine was shown to be powerful in adapting itself to these attack conditions, even at very low traffic rates, while protecting sensitive API endpoints without causing damages to normal traffic during the attack mitigation or the learning period.

DDoSM is informed by the FC-DDoS model, which captures the dynamics of population compromise given behavioral assessments of a DDoS attack. This is achieved by combining epidemiological modeling methods with dynamic modeling to analyze the grid networks and the attack itself, respectively. This contributes a novel approach alongside traditional graph-based approaches and provides validation for enhancements made to Ammune to fit the smart grid context.

DISSEMINATION & COMMUNICATION

Published papers

- Acarali, Dilara, et al. "Modelling smart grid IT-OT dependencies for DDoS impact propagation." Computers & Security 112 (2022): 102528. (Journal)
- Zarpelão, Bruno Bogaz, et al. "How Machine Learning Can Support Cyberattack Detection in Smart Grids." Artificial Intelligence Techniques for a Scalable Energy Transition. Springer, Cham, 2020. 225-258. (Book Chapter)
- Dixit, Akanksha, et al. "FAST DATA: A Fair, Secure and Trusted Decentralized IIoT Data Marketplace enabled by Blockchain." IEEE Internet of Things Journal (2021). (Journal)
- Acarali, Dilara, et al. "A characterisation of smart grid dos attacks." International Conference on Security and Privacy in New Computing Environments. Springer, Cham, 2020. (Conference)
- Acarali, Dilara, et al. "Modelling DoS Attacks & Interoperability in the Smart Grid." 2020 29th International Conference on Computer Communications and Networks (ICCCN). IEEE, 2020. (Conference)
- Dixit, Akanksha, Waqar Asif, and Muttukrishnan Rajarajan. "Smart-Contract Enabled Decentralized Identity Management Framework for Industry 4.0." IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society. IEEE, 2020. (Conference)

ABOUT THE COMPANIES

L7 Defense helps organizations to protect their infrastructure, applications, customers, employees, and partners against the growing risk of API-borne attacks. APIs have become critical for data sharing and applications integration and are therefore an attractive path for malicious attacks that expose organizations to new, continuously evolving threats.

With a team of experienced leaders and innovators, L7 Defense revolutionizes the way organizations protect their APIs using its advanced AI-based technology.

Ammune™, L7 Defense's platform technology, received in 2020 a Product Leadership Award by Frost & Sullivan for protecting APIs thanks to its novel unsupervised learning AI approach of protecting APIs.

City University of London's Institute of cyber security leads research in the areas of data privacy and security of the critical national infrastructure. The group has worked in network security and SCADA security for well over 20 years and has an international reputation in this area of research. As part of the EnergyShield project they developed new theoretical models to understand the impact on the smart grid due to the interdependency between the IO and OT layers. The developed models were validated using experimental smart meter data and showed a very good correlation. In addition, they also explored the DDoS attack scenarios in a smart grid jointly with L7Defense.

This project has received funding from the European Union's H2020 research and innovation programme under the Grant Agreement No. 832907

Website: www.energy-shield.eu **Twitter**: @EnergyShield_
LinkedIn Group; https://www.linkedin.com/groups/8831159/
Youtube channel: https://www.youtube.com/channel/UCtNRIf0uXvDsxVCSo1NfF_Q
E-mail: EnergyShield@siveco.ro