

## Overview

Databases continue to be the most commonly used backend storage in enterprises, and are employed in several contexts in the electrical grid. They are often integrated with vulnerable applications, such as web frontends, that allow injection attacks to be performed. The effectiveness of such attacks stems from a *semantic mismatch* between how SQL queries are believed to be executed and the actual way in which databases process them.

SEPTIC is a mechanism for DBMS attack prevention, which can also assist on the identification of the vulnerabilities in the applications. It was implemented in the MySQL database and successfully evaluated experimentally with various applications and alternative protection approaches.

## Approach

Existing DBMS defense mechanisms need to make assumptions about how the DBMS processes SQL queries. This leads to a *semantic mismatch* between the server-side language and DBMS



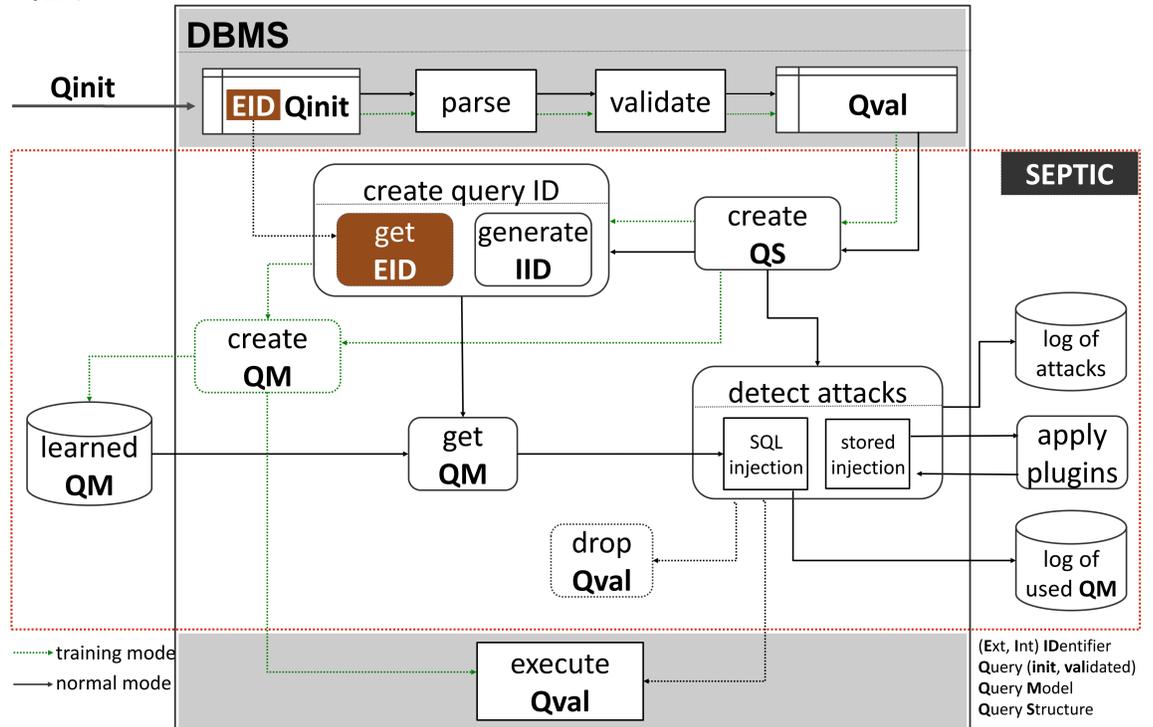
*difference of interpretation between how the queries are believed to be executed by DBMS and how they are actually run*

Our approach is a novel mechanism that **blocks injection attacks inside the DBMS**, protecting any application that uses the database. By running the mechanism inside the DBMS, the defense is provided off-the-shelf (without requiring installation), which leads an automatic protection similarly to what is currently available for binary programs based on techniques like address space layout randomization and canaries.

**Self-Protecting daTabases preventIng attacKs (SEPTIC)** interrupts the attacks in the DBMS at runtime and assists on the identification of the associated vulnerabilities in the applications. The main two categories of database attacks are addressed: *SQL injection attacks*, which continue to be among those with highest risk and for which new variants continue to appear; and *stored injection attacks*, which also involve SQL queries. For SQLI, we propose detecting attacks essentially by comparing queries with query models and circumventing the semantic mismatch problem. For stored injection, we employ plugins to deal with specific attacks before data is inserted in the database.



## Self-Protecting daTabases preventIng attacKs



## Experiments

Research questions:

1. Is SEPTIC able to detect and block attacks against code samples? (see Table 1)
2. Is it more efficient than other tools in the literature? (see Figure 1)
3. Does it solve the semantic mismatch problem better than other tools? (see Figure 1)
4. How does it perform in terms of false positives and false negatives? (see Figure 1)
5. Is SEPTIC able to detect and block attacks against real (open source) software? (see Table 1)
6. Is the performance overhead acceptable? (see Table 2)

Table 1. Two sets of experiments with SEPTIC

Metric	Exper 1	Exper 2	Total
Web applications	5	6	11
Synthetic tests	63	6	69
Attacks blocked	17	85	102
Vulnerabilities	17	25	42
Zero-day	--	1	1

Table 2. SEPTIC performance overhead in MySQL

SQL Inj.	Stored Inj.	SEPTIC performance
off	off	
on	off	0,8 %
off	on	
on	on	2,2 %

Summary of 63 tests

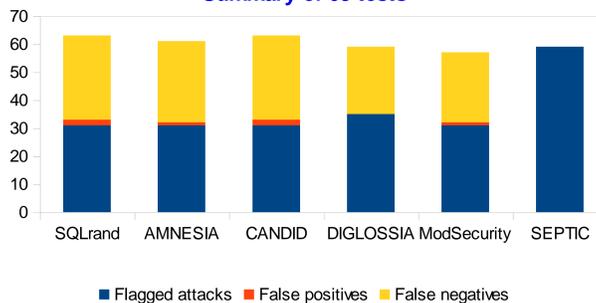


Figure 1. Detection attack comparison

## References

- [1] Ibéria Medeiros, Miguel Beatriz, Nuno Neves, Miguel Correia, *Hacking the DBMS to Prevent Injection Attacks*, In Proceedings of the ACM Conference on Data and Applications Security and Privacy (CODASPY), New Orleans, EUA, March 2016.
- [2] \_\_\_\_, *Demonstrating a Tool for Injection Attack Prevention in MySQL*, In Proceedings of the 47th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Denver, USA, June 2017.