F5 White Paper

Application and Data Security with F5 BIG-IP ASM and Oracle Database Firewall

Organizations need an end-to-end web application and database security solution to protect data, customers, and their businesses. The integrated solution from F5 and Oracle provides improved protection against SQL injection attacks and correlated reporting for richer contextual information.

by Peter Silva
Technical Marketing Manager—Security

Contents

Introduction

Defense in depth is a military technique that uses multiple layers of defense to prevent enemy infiltration and protect multiple locations simultaneously. With this strategy, the defender spreads its resources rather than protecting only a single location, causing the attacker to have to breach several layers of protection before reaching the intended target. One important facet of defense in depth is redundancy: if a defense component fails or is compromised, others are ready to step in and keep the protection intact.

Information technology also recognizes this technique as a best practice in system protection. Defense in depth fortifies infrastructure and systems with a layered security approach. Firewalls are stationed at the edge of the network and security mechanisms are usually at every segment. If an attacker circumvents the first layer, the next one should net them.

In IT, employing a defense in depth strategy involves redundancy by placing multiple iterations of a defensive mechanism in the path of an attacker. The muscle of a firewall is a critical defense component; but to achieve a truly secure system, fortification must also be based on context. A system with context takes into account the environment or conditions surrounding an event to make an informed decision about how to apply security. This is an especially important part of protecting a database. Database firewalls can stop a database attack in an instant; however, they lack the ability to decipher contextual data like user ID, session, cookie, browser type, IP address, location, and other metadata about the attack.

F5 and Oracle have partnered to offer enhanced security for web-based database applications. The integration between F5® BIG-IP® Application Security Manager™ (ASM) and Oracle Database Firewall provides richer forensic information about SQL injection attacks through correlated reporting.

Layered, Contextual Data Security

A database is the primary repository and retrieval mechanism for an enterprise's critical data—and protecting that database is crucial. As more application traffic moves over the web, sensitive data is exposed to new security vulnerabilities and attacks. Standalone technologies that protect against web or database attacks are

available, but their disconnection from one another means they lack context. Organizations need an end-to-end web application and database security solution to protect their data, their customers, and ultimately their businesses.

When a database firewall receives a request for access, typically it comes from the web server tier in front of the database. Because it is coming from the web server, the request looks legitimate; therefore every access attempt appears to be a trusted web-tier user request. Malicious hackers take advantage of this trusted status by attempting a SQL injection attack to fool the database into divulging sensitive information. SQL injection is an attack in which the user inserts malicious code into a string that is then passed on to the database for execution. This is usually accomplished by entering a SQL query or command script in a user input field, such as the password field. The hacker is essentially trying to bypass the application servers in order to manipulate the database directly. A successful attack could result in just unauthorized entry, or it could return the entire database containing user names, passwords, and other sensitive information. A database firewall, which protects against such an attack, does not have the visibility to gather information such as host name, user name, client IP, and browser. While it can see that a particular SQL query is invalid, it cannot decipher who made the request.

A web application firewall (WAF) on the other hand gathers user-side information so it can base policy decisions on the user's context. A WAF monitors every request and response from the browser to the web application and consults a policy to determine whether to allow the action and data. It uses information like user, session, cookie, and other contextual data to decide if the request is valid. WAFs are primarily focused on HTTP attacks and do a great job of thwarting that type of malicious traffic. They can also block most database-targeted attacks launched through a browser. However, given the complexity of detecting SQL injection attacks in the web application tier (lack of SQL-related context, understanding of SQL protocol) WAFs are not a fool-proof SQL injection prevention solution—there is a chance of false positives or overlooked attacks.

The answer is a joint solution from F5 and Oracle. The integration of F5 BIG-IP ASM and Oracle Database Firewall offers the database protection that Oracle is known for with the contextual intelligence that is baked into every F5 solution.

The power of BIG-IP ASM and Oracle Database Firewall working together is in the consolidated reporting of attacks, and the ability to set policy at the web application layer, which is coordinated at the database layer. With F5 and Oracle, an enterprise's database is protected by a layered, defense-in-depth architecture, backed with the contextual information required to make informed, intelligent decisions about

database security incidents. It's a comprehensive approach that enables enterprises to adapt quickly to changing threats and provides the logging and reporting capabilities needed to meet auditing and compliance regulations.

Two-Tier, End-to-End Protection

The F5 Networks® and Oracle partnership has a long history of producing integrated solutions. The BIG-IP Application Security Manager and Oracle Database Firewall solution links a web application firewall with a database firewall. BIG-IP ASM is an advanced WAF that provides comprehensive edge-of-network protection against a wide range of web-based attacks. It analyzes each HTTP/HTTPS request, and blocks potential attacks before they reach the web application server. Oracle Database Firewall is the first line of defense for databases, providing real-time monitoring of database activity on the network. Highly accurate, SQL grammar–based technology blocks unauthorized transactions, which helps prevent attacks from reaching the database. Oracle Database Firewall is deployed between the web application server and the database. It provides protection against attacks originating from inside or outside the network and works by analyzing the intent of the SQL statements sent to the database. It is not dependent on recognizing the syntax of known security threats, and can therefore block previously unseen attacks, including those targeted at an organization. It is easy to deploy, as it requires no changes to existing applications or databases.

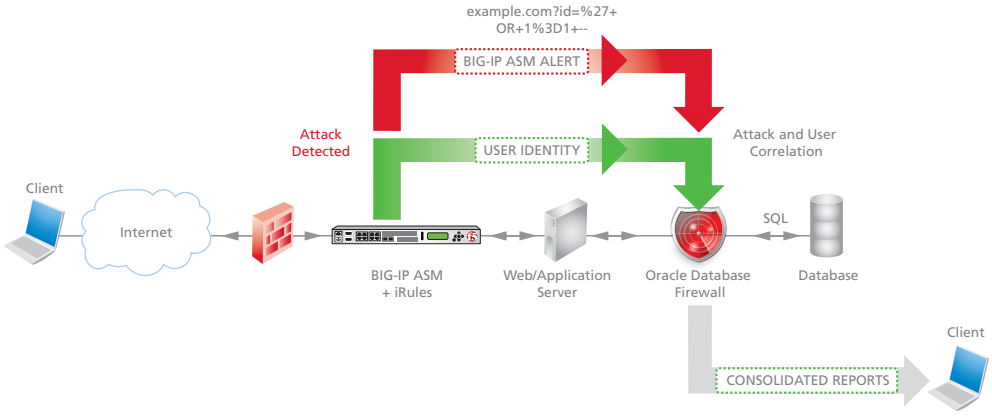

Figure 1: BIG-IP ASM and Oracle Database Firewall correlate and report on security events

BIG-IP ASM secures web traffic, and Oracle Database Firewall secures database traffic.

When a user logs into an application, BIG-IP ASM passes their identity to Oracle Database Firewall.

If a SQL injection takes place, BIG-IP ASM sends all context of the attack to Oracle Database Firewall, and the user's identity can now be associated with the attack in reports, based on session and the BIG-IP ASM session cookie.

The two products share common reporting on web-based attempts to gain access to sensitive data, subvert the database, or execute denial of service (DoS) attacks. Unified reporting for both the application firewall and database firewall provides convenient and comprehensive security monitoring. This integration between the two security solutions offers a comprehensive and holistic approach to protecting web and database tiers from SQL injection attacks.

How BIG-IP ASM and Oracle Database Firewall Work Together

When threats to data are detected, the F5 and Oracle solution monitors, alerts, or blocks the threat, and the identity of the user is shared between BIG-IP ASM and Oracle Database Firewall via iRules®. In the case of a malicious SQL injection, Oracle Database Firewall would block it instantly and log the action, but it can't determine who attempted the breach. BIG-IP ASM gathers the user name, client, browser, session information, time, cookies, URL, SQL statement and so on. Oracle's reporting engine then merges the two products' reports, and administrators can not only see that there was an attempted breach, but also the critical data needed to determine who caused the trigger. The triggered alerts and accompanying detailed reports provide immediate notification on the type and severity of a threat. With two information sources, BIG-IP ASM and Oracle Database Firewall, the resulting correlated data is richer, making policy creation more accurate and more granularly refined. With this level of detail, malicious or compromised users can be isolated, forced to re-authenticate, or prevented from accessing the application in real time. Subsequent attacks from the same user can be prevented, diverted, or rendered inert by the F5 and Oracle solution.

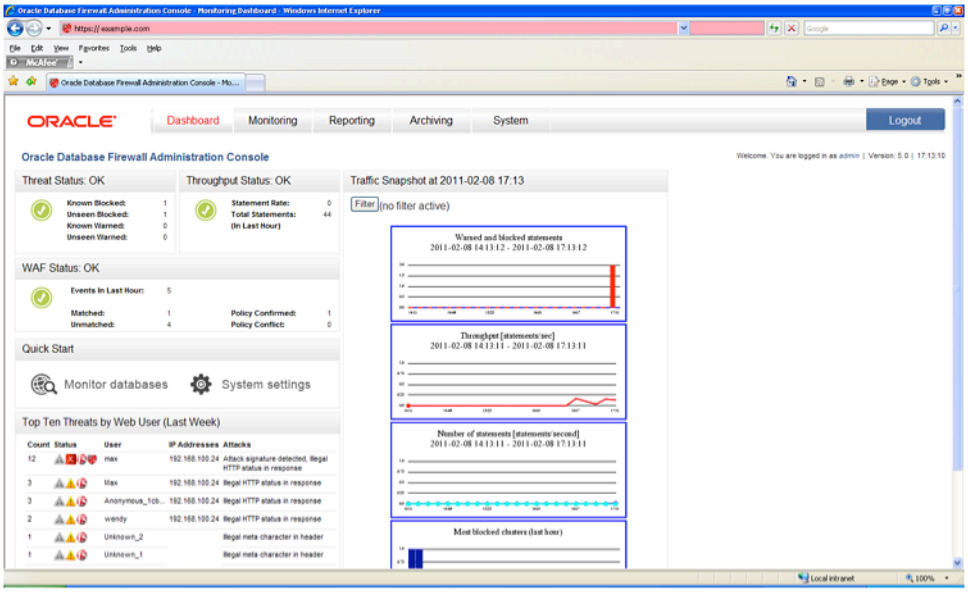

Figure 2: Oracle Database Firewall reporting GUI

There are many types of attacks that a WAF will detect and block that are not pointed at a database, so the database firewall will not address them. Conversely, a database firewall will detect and block many types of attacks that do not originate over the web, and that a WAF will not address. This gives enterprises a high degree of confidence that all SQL injections will be identified and blocked and that sensitive data is protected. SQL injection attacks can be very sophisticated—the combination of BIG-IP ASM and Oracle Database Firewall is an enterprise's best defense.

The database firewall benefits from receiving web application and user context from the WAF, as it generally has no visibility into the web application tier; and WAFs specialize in web application technologies and the threats related to them (XSS, session hijacking, HTTP protocol evasion, parameter tampering). They complement each other's security technology.

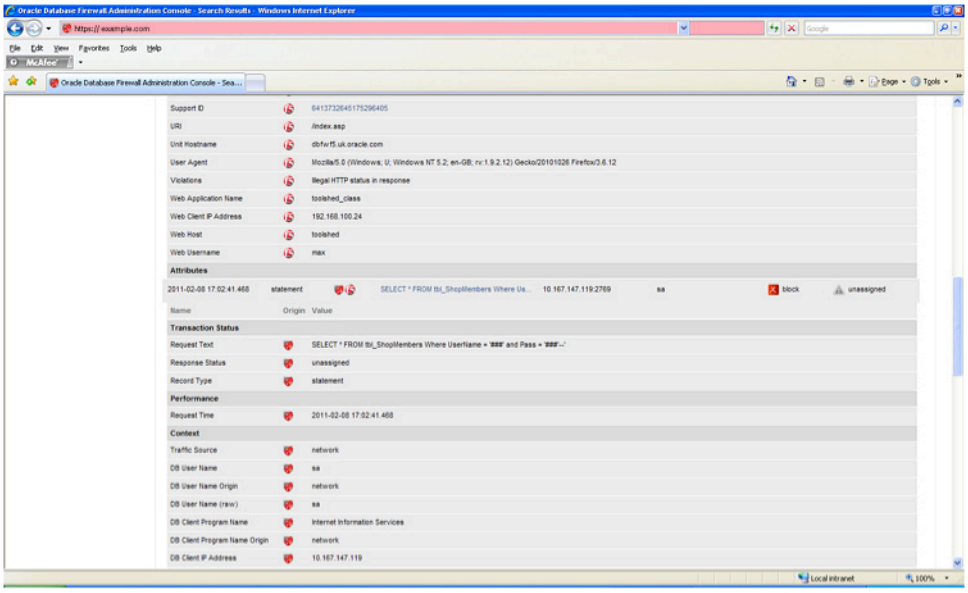

Figure 3: Oracle Database Firewall report screen

Interaction of BIG-IP ASM and Oracle Database Firewall

While both security solutions can detect and block SQL injection attacks, in this solution, BIG-IP ASM is configured to block all threats except SQL injection, which Oracle Database Firewall recognizes and blocks. BIG-IP ASM is configured to actively block malicious web application traffic, but to only alert on the SQL injection–related attempts. When BIG-IP ASM detects a SQL injection attack, it triggers an alert and sends the contextual information it has gathered to Oracle Database Firewall. This way, BIG-IP ASM provides Oracle Database Firewall with web application and user context, but without potentially generating a false positive and preventing Oracle Database Firewall from performing its primary security functions. This integration reduces SQL injection false positives and false negatives, and gives organizations a high degree of confidence that all SQL injection attacks will be mitigated. If BIG-IP ASM blocks HTTP requests that contain a SQL injection attack, Oracle Database Firewall will never receive the SQL request (via web application) and will have nothing to report.

In the unlikely event that information is compromised, BIG-IP ASM addresses the issue on the response. The Mask Data feature in BIG-IP ASM automatically scrubs sensitive data as it passes through to the user, preventing any data leakage. For instance, if a malicious user circumvented the system and generated a request for credit card information from the database, BIG-IP ASM would either block that request or scrub the output by replacing the credit card number with asterisks.

Together with reporting tools, this feature can be instrumental in gaining or maintaining regulatory compliance. Reporting and auditing are top criteria for many of the regulations in place today (PCI, HIPAA, SOX), and this solution can help ensure companies have the most detailed compliance information. Finally, one of the most significant benefits of this solution is that it can protect any SQL-based database, including Oracle Database; Microsoft SQL Server; IBM DB2 for Linux, Unix, Windows; and Sybase databases.

Conclusion

The integration of F5 BIG-IP ASM and Oracle Database Firewall enhances security for web-based database applications and gives enterprises the layered protection that security professionals recognize as a best practice, plus the contextual information needed to make intelligent decisions about what action to take. The integration between solutions provides improved SQL injection protection to F5 customers and correlated reporting for richer forensic information on SQL injection attacks to Oracle database customers.

The F5 and Oracle partnership has developed solutions that help organizations create agile IT infrastructures that align with their business demands. With F5 and Oracle, enterprises' sensitive database information is always secure, available, and delivered quickly.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

CS01-00048 0411