Ovum
TMT intelligence | informa

How Modern Database Security Complements Other Tools to Mitigate Risk

Database security complements and enhances perimeter security tools, IAM platforms, DLP, and log management/SIEM

Publication Date: 06 Mar 2019

Author: Rik Turner

Summary

Catalyst

In a previous white paper we highlighted the importance of data security, of which database security is a component part. We noted that securing data is often overlooked as companies focus on perimeter security technologies and identity and access management (IAM) tools.

In this white paper we consider why it is important to protect databases. We also define some of the core capabilities of data and database security. We then look at a range of technologies designed to detect and or/prevent breaches, outline their various shortcomings, and discuss how database security complements each of them.

Ovum view

Enterprise security is not only a combination of people, processes, and technology. It is also a mixture of the right technologies, correctly deployed, employed, and integrated into an enterprise-wide infrastructure.

Ovum recognizes the need for up-to-date defenses in the area of perimeter security, bearing in mind that the perimeter now extends to end-user devices outside the corporate infrastructure as well as to applications that reside in cloud environments. We also preach the value of contemporary IAM systems to handle the access requirements not only of employees, but also of contractors/partners, and increasingly too of customers.

However, as valuable as perimeter security tools and core security (Ovum's term for all forms of identity management technology) may be, there is no doubt that data security, including security arrangements around information held in databases, is of paramount importance, both as a standalone capability and as a complement to the other two areas of security infrastructure.

The case for database security almost makes itself. Databases typically hold between 50% and 60% of a company's most sensitive data, but it is also worth pointing out that, as revealed by the 2018 edition of the respected Verizon Data Breach Investigations Report (DBIR), the database server topped the list of assets involved in the 2,216 breaches studied for the report (see Figure 1).

Modern database security, which comprises both data activity monitoring (DAM) and security analytics, is thus a necessary part of an enterprise's security investments in its own right, serving to protect sensitive data in the most highly targeted of corporate assets. Above and beyond that, however, Ovum sees database security as an essential tool for mitigating cyber risk, as well as for improving breach prevention by complementing other types of security technologies, such as encryption and data loss prevention (DLP).

Figure 1: The database is the most targeted asset in breaches

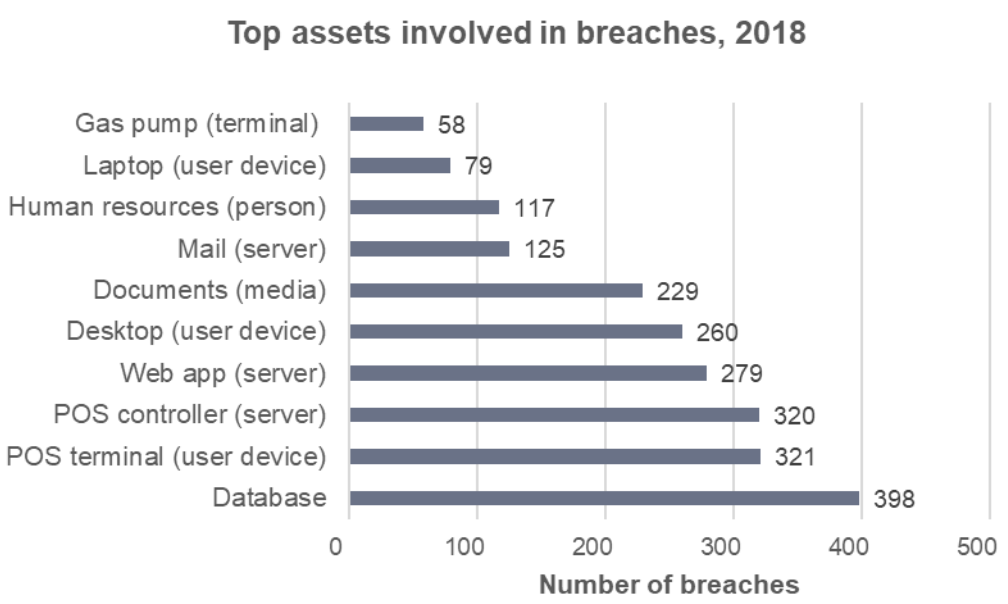

Source: Verizon DBIR

Key messages

- Perimeter defenses are not enough.
- Data security should be used to complement perimeter tools.
- Obfuscation renders data useless to thieves.
- IAM lists all employees and their access rights.
- DLP combats insider threat.
- Native logging/security incident and event management (SIEM) collect and correlate logs.
- Database security can complement and enhance all these capabilities.

Why data security matters

Perimeter defenses are not enough

Full-year figures are not yet published, but in the first half of 2018, a staggering 4.5 billion data records were compromised worldwide, or 291 every second. And of course, it was also the year that we learned of Cambridge Analytica's abuse of data on 87 million Facebook users, because 270,000 of them had installed its "Thisisyourdigitallife" application, which meant that the offending company could access data on all their friends' likes and dislikes, etc.

It is clear that perimeter security technologies such as anti-malware, intrusion detection and prevention systems (IDS/IPS), and firewalls are insufficient to combat breaches. This is particularly true now, as over a third of attacks currently use fileless techniques that leverage legitimate utilities such as PowerShell, making them far more difficult to detect.

Data security should be used to complement perimeter tools

In this scenario, securing the actual data becomes a vital component of companies' defenses, complementing both their perimeter security systems and whatever IAM arrangements they have for employees, partners, and customers. This focus on the data can be thought of as the backstop security measure, designed to thwart attackers when all else has failed.

In Ovum's categorization, the term "data security" covers two main types of technology:

- data obfuscation techniques, such as encryption (including format-preserving encryption), tokenization and masking
- database security, which spans the capabilities of data discovery and classification, DAM, database firewalling, and security analytics that enhances data breach prevention.

We will now look at a range of technologies across the areas of data, perimeter, and core security and consider how database security can complement and enhance them. Those technologies are as follows:

- obfuscation
- identity and access management (IAM)
- data loss prevention (DLP)
- native logging/security incident and event management (SIEM).

Obfuscation renders data useless to thieves

Encryption protects data at rest and in motion

Data obfuscation techniques, of which encryption is the best known, are designed to protect data when it is at rest, which may be on the hard drive of a laptop or in a dedicated storage device in a data center, as well as when it is motion – that is, when it is traversing a network connection.

Encryption is the process of translating plain text data (plaintext) into something that appears to be random and meaningless (ciphertext). An encryption key is used during both the encryption and decryption processes, and to decrypt a piece of ciphertext, the same key that was used to encrypt the data must be used.

This process, and the scenarios in which it is applied, means that if an attacker has penetrated a corporate infrastructure and gains access to – and even steals – data, it will be useless to them unless they have the appropriate key to decipher it. Equally, if a laptop is stolen or mislaid, the data held on the device will not be at risk if it is in encrypted form.

Meanwhile for data in motion on a LAN, WAN, mobile, or broadband connection, encryption can be carried out using the Transport Layer Security (TLS) protocol to guarantee security and integrity.

Other forms of obfuscation exist for different purposes

Other forms of obfuscation include the following:

- tokenization, which is widely used in the payment card industry and, rather than using a cipher like encryption, works by generating random strings of non-sensitive data to replace sensitive ones
- format-preserving encryption (FPE), which is a special subset of encryption in which the output (i.e., the ciphertext) is in the same format as the input (the plaintext), rendering it useful in certain scenarios
- data masking, which is the process of hiding original data with modified content (characters or other data) in such a way that it can still be used in test environments for purposes of application development.

Obfuscation has shortcomings

All forms of obfuscation have their shortcomings in terms of security. Since encryption (including FPE) relies on keys, it is only as secure as its key management system. Keys are stored in a vault, so if an attacker can gain access to the vault, they effectively have the "keys to the kingdom." Equally, hackers can brute force individual passwords, and with the advent of commercially available quantum computers over the next decade they will also be able to reverse engineer current encryption algorithms.

By contrast, tokenization has no underlying cipher; that is, there is no code determining what the non-sensitive text that replaces the sensitive one contains. Rather the non-sensitive text is randomly generated, which makes it impossible to reverse engineer, since there is no mathematical logic behind it. That said, it is obviously as critical to protect the vault in which the tokens are stored as it is to keep encryption keys securely locked away: if someone can steal the tokens, it's game over for the protection system.

Meanwhile some forms of encryption, such as Transparent Data Encryption (TDE), only provide physical security, and of course all forms of obfuscation are useless against ransomware attacks.

How database security can help

Database security comprises various capabilities, which collectively can provide additional protection beyond data obfuscation techniques. First a company's sensitive data must be discovered (i.e., all its locations across the infrastructure are identified and catalogued) and classified according to a standard, enterprise-wide taxonomy defining its level of sensitivity.

Once this phase is completed, DAM can be implemented. Continuously and in real time, it monitors access requests and database queries to the data to provide visibility into the context (who, what, where, when, how) of what is occurring in a database and, by integrating with security analytics, reaches decisions in real time about what should be blocked and what can be allowed, using database firewalling to enforce that decision.

Security analytics can also help identify when a compromised user account is starting to act differently from its normal activity. Thus, if a privileged user's account has been compromised, since such users typically access a database with the keys in hand, analytics can pick up on changes to that user's behavior.

IAM lists all employees and their access rights

Identities enable employees and partners to do their jobs

IAM is a well-established technology in which the identities of a company's entire staff are held in a directory, together with their role in the organization and the corporate assets that they are therefore entitled to access to do their jobs. As an increasing number of business processes have been outsourced to other companies, employees of those partner organizations can also be included in an IAM system for the same purpose.

Beyond these B2E and B2B environments, the expansion of e-commerce has meant that in recent years IAM has been applied to the end customer (the B2C scenario). In online banking, for instance, an account holder returning to the bank's website for some level of interaction, from checking their balance to making a payment, can be recognized and authorized for the appropriate level of access. Institutions can even stipulate different security requirements, such that checking one's balance can be allowed with only basic credentials, while transferring funds may require additional authentication via a one-time password sent to a mobile phone.

IAM's shortcomings for security

IAM systems were not designed to detect breaches, but merely to make decisions on whether to enable access to a given asset and to determine what level of access is appropriate (e.g., read only, read and write, download data, print data, copy to another file, etc.). Thus, if an attacker has successfully stolen a legitimate user's credentials, and particularly if they are the credentials of a privileged user such as a systems administrator or a C-level executive, an IAM system will see their log-on activity as legitimate and enable access accordingly. It is no coincidence that the more advanced cyberattacks will often take up residence within a victim's infrastructure and perform reconnaissance, seeking privileged users' credentials to gain higher levels of access rights before attempting any actual data theft.

Privileged access management (PAM) systems have evolved to provide greater security around the credentials of privileged users. However, many of them rely on a password vault in the back end, which if breached will again enable an attacker to gain the "keys to the castle," with nothing inherent in the PAM system to detect that breach.

How database security can help

Unlike IAM systems, modern database security looks not at a user's credentials but at their behavior (think of it as user behavior analysis for databases, if you will). As such, it can complement an IAM and/or PAM system, helping security teams to detect anomalies and identify which user account may have been abused/compromised. In other words, instead of looking at what the users say they are, database security looks at what they do.

It also monitors how privileged users access data and what type of data they touch, and it triggers a real-time alert if a security policy has been violated or if suspicious data activity has been detected by security analytics. Once an alert has been triggered around anomalous behavior or a policy violation, a security team can block or quarantine suspicious users and inappropriate data access, containing potential data breaches more effectively.

DLP combats insider threat

The three categories of insider threat

Data loss prevention (DLP) technology is designed to help companies combat insider threat, which itself falls into three different categories:

- Firstly, DLP helps them defend against threats from insiders who have gone rogue, by inspecting outbound traffic and determining when something contravenes preset rules.
- Secondly, there are threats from insiders who have simply made a mistake and are about to unwittingly send sensitive data outside the corporate perimeter. Again, by inspecting their outbound traffic, a system can avert any loss before the data has left the company's infrastructure.
- Thirdly, DLP helps defend against apparent insider threats that, in reality, are the result of account hijack by an external threat actor. Again, outbound traffic inspection can detect what may therefore be an attempt to exfiltrate data to a command and control server.

DLP's shortcomings

For all its undoubted value, DLP technology also has its limitations. For instance, it does not, per se, detect breaches, though it does help a company defend against their effects, namely the exfiltration of data. However, if an attacker has been able to hoodwink a DLP system and has exported sensitive data, say to a command-and-control server somewhere on the public internet, the company has no further control over it; it cannot, using DLP technology, remotely "kill" a file or render it unreadable once it is outside the corporate environment.

On the operational front, traditional DLP is rules based, which is a time- and resource-intensive process requiring coordination between different teams and implementation of appropriate policies and standards. And of course, it is always behind the curve, chasing the latest attack methodologies.

Finally, DLP relies on the company having a complete understanding of where all its data (both classified and unclassified) resides, who its users are, and what they are entitled to access. A proper data discovery and classification project (itself a time-consuming activity) is a sine qua non of a good DLP outcome.

How database security can help

The first component of database security is also a discovery and classification process. However, unlike an enterprise-wide project covering all aspects of a company database, security discovery and classification is designed specifically to detect sensitive data in databases. In other words, it takes a more focused approach, which can help kick-start a DLP deployment by getting it going more quickly, leaving the rest of an all-encompassing discovery and classification project till after the DLP is up and running, with the database discovery and classification serving as the basis for the DLP systems to start operating.

This process also helps make decisions about whether data that is being sent beyond the corporate perimeter should be able to leave, who is entitled to send it, and who are legitimate recipients. Discovery can also help by locating rogue and unknown databases, enabling the security team to shut them down and thereby reduce the scope of what the DLP system monitors.

In addition, database security can ease the pressure on DLP systems by detecting and blocking anomalous behavior within the database before the exfiltration attempt is even made, thus reducing the DLP platform's workload. It can also get ahead of the DLP system, detecting anomalies even before a company's DLP rules have been updated.

Native logging/SIEM collect and correlate logs

Log management alone is not enough

Log management, or native logging, provides a central data store for all event logs generated by a company's infrastructure. While it is an essential part of security and compliance functions in many companies, it is clearly not enough in its own right. There is the issue of event overload, where a company's infrastructure produces simply too many logs for it to carry out any meaningful analysis in real time.

A recent study suggests that a SIEM operated by the average Fortune 500 company typically receives anywhere from 12TB to 15TB of plain text data every month. The job of sifting through mountains of data to detect indications of compromise, not to mention reaching any meaningful conclusions, is a truly Herculean task in such circumstances.

Beyond that, there is the problem that native logging on its own does not provide companies with any actionable insights into what is happening in their infrastructure.

SIEMs were the next step after log management

This is why SIEM platforms emerged in the 2000s. While their primary focus is still logging, they bring the following additional key functionality into play:

- security event management (SEM), which focuses on real-time monitoring, correlating events, providing overarching console views, and customizing notifications
- security information management (SIM), which provides long-term storage, analysis, manipulation, and reporting on logs and security records
- security event correlation (SEC), which tracks and alerts designated administrators when a peculiar sequence of events occurs, such as three failed login attempts under the same user name on different machines.

The shortcomings of SIEM

While SIEMs are a well-established part of the security market, in recent years they have been shown to be lacking in a number of areas. There is the obvious problem of volume, in that the average corporate network throws up hundreds of thousands, if not millions of events every day, such that finding meaningful incidents within them becomes a bit like looking for a needle in a haystack. There are customers who complain at the storage cost associated with holding so much data, when they then face the challenge of actually finding anything within it.

In addition, there is the fact that SIEMs are designed to work as central repositories for logs, and even with their SEM, SIM, and SEC capabilities, they don't have the necessary data access information to provide context. Neither do they address other types of data besides logs, and thus must be

supplemented with other tools such as user and entity behavior analysis (UEBA) to broaden their scope. Indeed, SIEM vendors have acquired UEBA platforms to integrate the two, while some UEBA vendors now position their products as "next generation" SIEMs.

There are problems even with the combination of SIEM and UEBA, however. Firstly, UEBA's acknowledged focus on the detection of lateral movement within a corporate infrastructure, such as user logins and logouts, can throw up a lot of false positives, a feature which can frequently lead security teams to turn it off in frustration.

For instance, if a database administrator typically works from 9am to 5pm, logging on between those hours will constitute normality, while an attempt to log on outside of them would tend to trigger an alert. That might be a legitimate cause for concern, in that it might indicate that the admin is doing something untoward and potentially malicious. It might even mean that their account had been hijacked and that an attacker was masquerading as them. However, it could simply mean that something urgent had come up after hours and the DBA genuinely needed access to the system. Supplementary information about the criticality of the information accessed could help determine whether an alert really did need to be raised, or indeed whether an access request should be blocked.

In addition, UEBA typically analyzes logs spanning a variety of endpoints and network security products but lacks knowledge of databases.

How database security can help

Modern database security applies advanced security analytics to the database events captured by DAM and enables them to be processed without additional scripting and with no additional infrastructure in the form of hardware, software, or storage. Security analytics, meanwhile, can enable anyone in the security operations center (SOC) to understand the security posture of their database in minutes, which is clearly valuable, given the shortage of expertise and the volume of database traffic to be sifted through.

Database security can also complement the SIEM + UEBA combination, firstly by improving detection accuracy and reducing the logs sent to SIEM, and secondly by adding the database awareness that is lacking in UEBA. For instance, it can distinguish between sensitive and non-sensitive data access, which a UEBA platform cannot. It can further distinguish between machine and human access, which again is not a capability of UEBA.

A further feature of database security that complements UEBA is its ability to report on the number of records accessed and on whether any of the access was to sensitive data. Finally, it can enable comparisons between the database activity of individuals within a peer group, which again is not something a UEBA platform can do.

As a result, database security with analytics can make a SIEM more effective and reduce costs by sending only meaningful incidents to SIEM.

Imperva's approach to database security

Imperva Data Security helps mitigate data breach risks by pinpointing dangerous data activity using security analytics and continuous data activity assessment. It utilizes machine learning and data access analytics to distill millions of alerts and to prioritize the most critical incidents, accurately

detecting suspicious data activity while reducing the alerts that get sent to a SIEM. It provides security analysts the risk context in plain language to quickly investigate and contain a potential breach. The solution also offers visibility into a company's data and actionable insights, determining and providing context around the following:

- where a potential breach comes from
- how critical the threat is to data
- the users/database/types of data involved, including privileged users
- whether it is a suspicious/dangerous data access and why
- what that user's typical behavior is, performing so-called peer group analysis
- what data has been touched and by whom, and how it is being used.

Once inappropriate data access is identified, Imperva Data Security allows the security team to either block or quarantine that user/data access.

Imperva Data Security also helps organizations reduce data breach risks by uncovering what data they process or possess. As data grows over time, the solution automates data discovery and classification, so companies can keep track of where sensitive data resides while complying with data-privacy requirements. The data-masking capability within the solution helps companies reduce their attack surface by shrinking their sensitive data landscape, mitigating risk by preventing data exposure to third-party contractors or users beyond need-to-know.

Recommendations

Securing information in a database helps mitigate risk

With databases being the number one target for attackers, there is a clear and obvious reason to deploy database security. There is a direct correlation between the use of modern database security – DAM together with security analytics – and a company's ability to mitigate the risk of data loss.

Database security complements other security tools

If you are planning to improve your company's security infrastructure with data obfuscation, identity and access management (IAM), data loss prevention (DLP), or security incident and event management (SIEM), bear in mind that all these technologies have their weaknesses. Database security can address their shortcomings and thereby enhance their functionality. It should be deployed as a means of reducing your company's overall risks and improving its security posture.

Appendix

Author

Rik Turner, Principal Analyst, Infrastructure Solutions

reik.turner@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

CONTACT US

ovum.informa.com

askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing

Boston

Chicago

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

Paris

San Francisco

Sao Paulo

Shanghai

Singapore

Sydney

Tokyo