hexatier

Secure, Comply, Go!

Database as a Service (DBaaS) Security Research 2016

Security is the number one concern preventing companies from moving their most sensitive business data to the cloud.

Database as a Service (DBaaS) Security Research 2016* was conducted by HexaTier, under the leadership of David Maman, CTO and founder of HexaTier, and with the contrbutors Hadar Eshel, independent consultant, and Ziv Gadot, external security research analyst.

The views expressed in this report, as well as the information included in it, do not necessarily reflect the opinion or position of HexaTier.

*This report is dated in the year of publication rather than the fieldwork completion date. Please note that the research was conducted in 2015.

Introduction

Over the past three years, companies have been moving IT resources to the cloud at a rapidly increasing rate. SMBs and large enterprises alike are embracing the cloud with the goals of reducing costs, improving scalability, increasing productivity and otherwise serving the organization's business goals.

As part of this major trend, organizations are migrating business-critical and sensitive data stored in databases to cloud-hosted and Database as a Service (DBaaS) solutions. DBaaS significantly reduces costs and management overhead while offering rapid provisioning and dramatically improved scalability planning. DBaaS helps organizations better focus on their core business.

Meanwhile, databases represent the number one source of data breaches, whether located on premises or in the cloud. Research by Verizon indicates that 92% of stolen records come from compromised databases, leading to the inescapable conclusion that protecting their databases must be a top priority for all organizations.

It seems that there is a never-ending stream of damaging data breaches among some of the world's largest organizations, in nearly every industry: Sony, Subway, Target, Anthem, Ashley Madison and even the US federal government's Office of Personnel Management are examples from the recent past.

This reality has changed the way today's organizations think about securing their databases in general, and it highlights the critical nature of DBaaS security in particular.

About this DBaaS Security Research

We conducted market research to help organizations focus on the most important aspects of DBaaS security, whether in their current or future DBaaS deployments, by determining the top concerns facing organizations considering moving their databases to DBaaS providers. This paper is structured around the top 10 concerns indicated by the survey's 574 respondents.

Target Audience

HexaTier reached out to possible participants for this survey in two ways:

- The company contacted its database of contacts directly via email.
- The company publicized the survey via social media (Twitter, LinkedIn and Facebook).

Participants

Figure 1. presents the job roles of the survey respondents and found that database administrator was the most common job role among participants in this survey.

73% said that security and compliance are the #1 concern when migrating databases to the cloud.

92% **of stolen records come from compromised database servers.**
Source: Verizon Data Breach Investigations Report

Figure 1 lists respondents job roles.

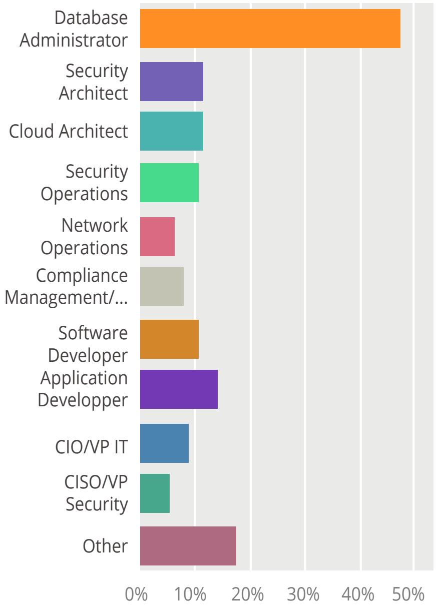

DBaaS Defined

Database as a Service (DBaaS), also known as a cloud-hosted database, is a database system that runs on a cloud computing platform managed outside of an organization's own premises.

There are two common database in the cloud deployment models: self-managed and as a service. In the former arrangement, customers install a virtual machine image as well as a database server on this image and manage it themselves. In the latter, the cloud hosting provider provides database as a service and manages the server and database on behalf of the customer.

Some cloud databases are SQL-based (Relational) and some use a NoSQL data model.

The DBaaS Security Research Structure
For each concern, we have mapped the risks, the solutions provided by the major DBaaS providers, the solutions available from the DBaaS provider and specific recommendations that an organization should consider. Particularly by understanding the gaps between the threats and what the DBaaS providers offer, organizations can best satisfy the organization's business needs while addressing security concerns.

For this paper, we examined three of the largest DBaaS providers:

1. Amazon Relational Database Service (Amazon RDS)
2. Microsoft Azure SQL Database
3. Google Cloud SQL

However, the concerns and solutions described are applicable to any DBaaS provider.

Why Organizations Choose DBaaS
The two primary reasons that organizations choose to run their databases in the cloud, instead of using on-premises servers, are lower costs and greater agility/scalability. Figure 2. shows other reasons for choosing DBaaS.
Note that these motivations are not unique to DBaaS, but rather are the primary motivations to move to cloud-based services in general.

Author's Note

While organizations and security consultants have accumulated years of experience and best practices regarding securing on-premises database infrastructure, DBaaS security mostly remains a mystery. Organizations have unwittingly begun storing their most sensitive data in the cloud without fully understanding the risks and implications. This paper intends to shed some light on this mystery.

Understanding Database Security

While most organizations address IT security in terms of policies of who has what level of access to what resources, the fact is that the "attack surface" in cyber security is much wider than this.

For a brief overview of database security in general, before delving into the specific issues of DBaaS, read the appendix, Understanding Database Security.

DBaaS Defined

Figure 2 relates to the question: Why DBaaS service was chosen as opposed to a traditional database.

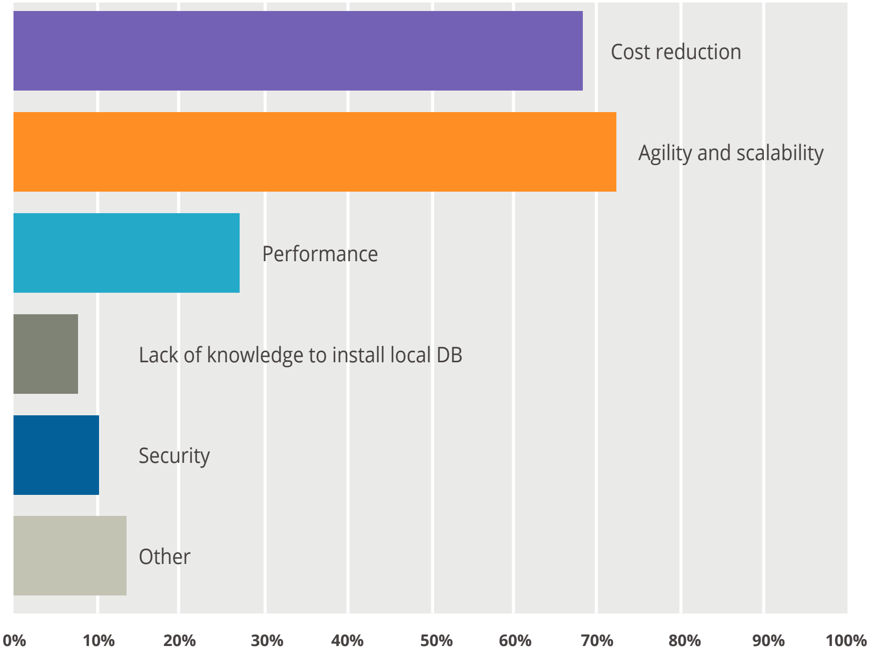

"Cost reduction" and "agility and scalability" are the two main reasons cited for organizations to move their databases to the cloud. These are also the main reasons to move to the cloud in general.

47%
of the respondents are using or evaluating DBaaS providers.

DBaaS Security

Placing a database in the cloud significantly changes its security threat landscape. While many of the traditional on-premises risks remain-data leakage risk from privileged users with access to the data, the presence of unidentified sensitive data and SQL injection attacks are some examples-the cloud introduces its own additional risks. On the other hand, there are ways to leverage the cloud by outsourcing some of the risk mitigation to the cloud provider. For example, physical access security and OS security is always the responsibility of the DBaaS provider.

Key Findings

1. Security is the number one concern preventing organizations from using DBaaS

Security risks were found to be the number one factor with 44% preventing organizations from using DBaaS instead of traditional on-premises databases or Virtual Private Cloud (VPC) databases. This particular finding illustrates how important it was for us to tackle DBaaS security concerns and provide clarity about DBaaS security, thus, enabling organizations to make smarter security -related decisions.

Figure 3 states respondents main reason that prevents their organizations from moving to DBaaS.

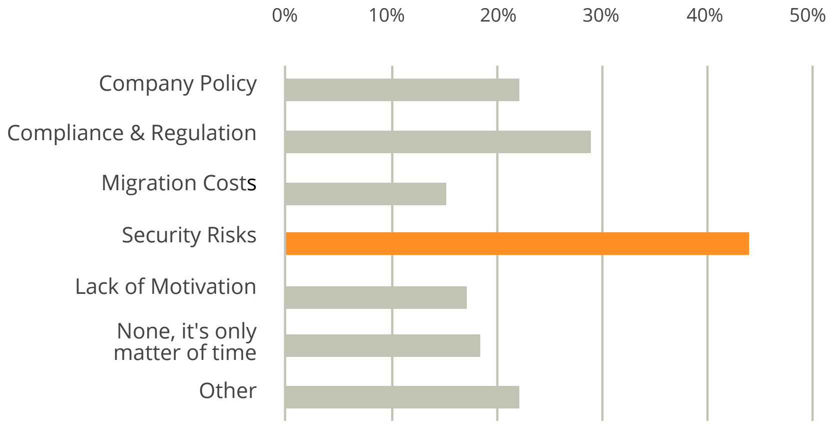

Security risks are, by far, the primary reasons given preventing organizations from moving to the cloud (44%).

2. Compliance and regulations are the second-most cited concerns

Concerns related to compliance and regulations came in second place, with 29% of respondents indicating that these concerns are preventing them from moving to DBaaS.

When we counted individual responses that indicated 'Security Risks' or 'Compliance & Regulation' or 'Company Policy' or all three, we see that 61% of respondents consider these reasons a barrier to DBaaS.

Key Findings

Our research indicates that this concern is not merely due to fear of the unknown, but that the tools and procedures available for DBaaS compliance are not as mature as those available for traditional databases. To comply with common regulations, including PCI, HIPAA and SOX, it is obvious that access to sensitive data must be monitored, sometimes blocked and even masked when accessed by people in specific roles (such as DBA, developer or QA engineer).

Cloud vendors provide their customers with an overall compliance posture, covering their entire cloud infrastructure, without focusing on which components comply with which regulations. This situation is confusing, to say the least. As time passes, this will have to improve, enabling organizations to get a clear view of which cloud components can truly comply with regulations relevant to them.

61%
of respondents consider regulatory reasons as a barrier preventing them from moving to DBaaS.

Figure 4 states the four major DBaaS concerns indicated in the survey results.

CONCERN TYPE	CONCERNED ORGANIZATIONS
Safety of data storage	39.2%
DBaaS provider can access your data	35.3%
Location of data (storage, backup or DR)	35.3%
External users are able to reach the DB	31.4%

Our survey contained concerns that are genuinely related to cloud-hosted services (e.g., that a cloud provider can access an organization's data that it stores), as well as concerns that also exist in traditional on-premises databases (e.g., data theft by authorized users).

When examining the top concerns indicated by the survey results, it is clear that they focus on concerns that exist only in cloud environments. When moving data to the cloud, organizations are rightly concerned that cloud providers receive a degree of control over their data and where it is stored, responsibility for its security and actual access to the data itself. Any of these aspects can impact on data security, company policies and compliance with data-related regulations.

Key Findings

4. Some concerns are real, while others are more frightening than they actually are

Some of the concerns expressed by the survey participants are very real and indicate problematic or immature aspects of cloud-hosted services. For example, the issue of regulatory compliance is indeed more problematic in the cloud, simply because the cloud services currently available for the enforcement and monitoring of compliance are much less mature than those for traditional databases.

On the other hand, other issues that initially appear as great concerns in the cloud are not really significant problems. For example, the concern that external users can easily gain access to the database now that it's outside the organization's firewall is not a big issue. All major cloud providers offer basic firewall capabilities. Organizations simply need to configure these firewalls appropriately. For organizations that require additional protections not provided by native cloud firewalls, third-party solutions are available.

5. Encryption is only a partial solution to data security concerns

Many of the concerns address the security of sensitive data in the cloud, and the risk that the cloud provider or another tenant can access or manipulate the data. This is a valid concern. The only way to fully secure data is to encrypt it. There are actually several types of encryption (all are addressed in this paper), but the only way to fully secure the data is to encrypt it before it is transferred to the cloud. This is known as client-side encryption.

The problem with client-side encryption is that this type of encryption disables many core database functions, such as stored procedures and fast search/sort functions. Because many organizations cannot pay this price, cloud-hosted data security remains at risk-and must be trusted to the DBaaS provider.

Figure 5 illustrates the respondents employment by industry - the most common was finance and banking.

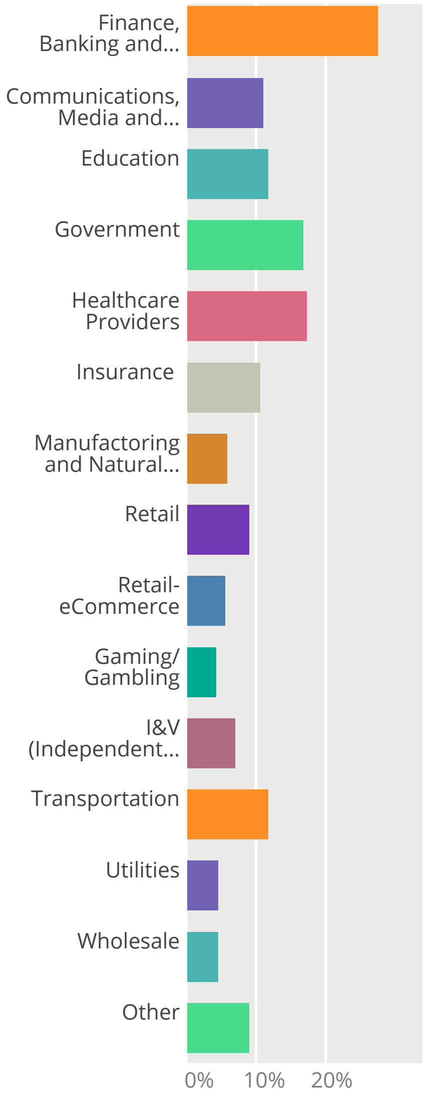

Survey Results

This survey was conducted among existing database community members as well as top thought leaders. 574 IT professionals participated in the survey, including market analysts, product managers, large enterprise infrastructure managers, analysts, DBAs, security officers, compliance professionals, cloud architects, and dev ops. Most survey participants work at large enterprises that are interested in moving portions of their infrastructure to cloud environments.

One of the main goals of the survey was to determine which risk factors are considered critical enough to actually prevent an organization from moving certain databases to the cloud.

Figure 6 presents the list of security concerns that prevented respondents from moving to a DBaaS service.

ANSWER OPTIONS	RESPONSE PERCENT
Safety of data storage	39.2%
DBaaS provider can access your data	35.3%
Location of data (storage, backup or DR)	35.3%
External users are able to reach the DB	31.4%
Compliance	29.4%
Performance attacks on the DB	23.5%
Hidden sensitive data (sensitive data exists in the DB which no one realizes)	21.6%
SQL injection attacks	17.6%
Data theft by authorized users	17.6%
A neighbor tenant has indirect access to data (e.g., by shared, non-sanitized memory)	17.6%
A neighbor tenant infiltrates the database	15.7%
BCDR (Business Continuity Data Recovery)	11.8%
Other	11.8%
DBaaS provider has been compromised	5.9%

Survey Results

Based on these responses, we chose to focus this research on the top 10 security concerns around Database as a Service:

- 🟣 Concerns also for on-premises databases
- ⚪ Genuine DBasS-specific concerns

TOP 10 SECURITY CONCERNS

01	Safety of data storage	🟣
02	DBaaS provider can access your data	🟣
03	Physical location of the production data and the backup data	🟣
04	External users can reach your DBaaS	🟣
05	Difficulties with meeting regulatory compliance	⚪
06	DDoS and performance attacks	⚪
07	Hidden sensitive data	⚪
08	SQL injection attacks	⚪
09	Theft by authorized users	🟣
10	Neighbor tenant access	🟣

It's important to note that some of these concerns relate to risks that are truly part of the DBaaS environment, however, some should be of equal concern for on-premises databases as well.

Shared Responsibility Analysis

The following chart maps the security concerns indicated in our survey to the divisions of responsibility between provider and customer:

Figure 7 lists the DBaaS security concerns in respect to shared responsibility that prevent from moving to DBaaS.

PROVIDER RESPONSIBILITY	
Safety of data storage	39%
DBaaS provider can access your data	35.2%
Location of data (storage or backup)	35.2%
Performance attacks on the DB	28%
A neighbor tenant has indirect access to data (e.g. by shared non-sanitized memory)	17%
A neighbor tenant can infiltrate the database	15.7%
BCDR (Business Continuity Data Recovery)	13%
DBaaS provider has been compromised	6%
SHARED RESPONSIBILITY	
Compliance	29%
DBaaS does not meet regulations or compliance	19.6%
CUSTOMER RESPONSIBILITY	
External users are able to reach the DB	32%
Hidden sensitive data (sensitive data exists in DB which no one realizes)	22%
SQL injection attacks	17.4%
Data theft by authorized users	17.4%

Firstly, it is important to note that participants expressed concerns regarding many risks that are the cloud providers responsibility. This indicates that the shared responsibility concept should not be taken as a "good will" agreement made between the customer and the cloud provider, or as a line drawn between the two where both parties (or at least the customer) feel comfortable.

Instead, it is line that is formed by tearing the responsibility into two. The DBaaS provider is responsible for physical security because it's a necessary part of the arrangement, not because the customer intentionally wants to outsource this responsibility.

Furthermore, the fact that cloud providers offer features such as Transparent Data Encryption (TDE)-a protection of the customer's data at rest-is proof that, among other concerns, there is some lack of trust of the provider's employees who, necessarily, have access to customers' data.

Here, we see a problem with shared responsibilities as, neither side is happy with the arrangement and don't fully trust one another.

> Most of the concerns indicated by survey respondents regard concerns that fall under the cloud provider's responsibility, indicating that respondents trust their own abilities to secure their databases more than those of cloud providers.

Top 10 DBaaS Security Concerns

01 Safety of data storage

Introduction to the Data Storage Threat

As seen in the survey results, the single biggest concern indicated by 39.2% of participants is "safety of the data".

What does "safety of data storage" mean? While this is probably the broadest term we provided in the available responses to the questions, it's a valid concern. Safety of the data includes issues such as:

- Where is the data stored?
- Who has access to the data?
- What actions can those with access perform on the data?
- What access is logged and reported?

When referring to "safety of data storage" the high-level concern relates to more specific issues, like:

- DBaaS provider can access your data
- A neighbor tenant has indirect access to data (e.g., by shared, non-sanitized memory)
- A neighbor tenant can infiltrate the database

We will address this concern, that the DBaaS provider can access your data, in the following chapter. Concerns about a neighbor tenant having indirect access to data and infiltrating the database, will be discussed in chapter 10. However, the topic of encryption is a primary means to address all of these specific concerns, and this will be discussed in detail in the following chapter.

02 DBaaS provider can access your data

Introduction to the Threat

An obvious side effect of hosting your data in the cloud is that you cannot control who has physical or logical access to the servers on which your data is stored and processed. More specifically, certain employees of the cloud provider will certainly have access to the servers, and therefore, your data. This introduces a possible avenue of data theft over which you have no control at all.

To understand the degree of risk inherent in DBaaS, it is worth comparing it to a database running on a virtual machine in a VPC. In a VPC environment, a cloud provider employee can physically steal a hard drive or eavesdrop on network traffic, but you still maintain your own software, OS, encryption and system hardening. Thus, you can take measures to make your VPC-hosted database extremely difficult to breach. However, in DBaaS environments, where the provider provisions the software, you are forced to trust the provider and its security practices to a much greater degree.

Of course, every reputable cloud provider will have in place procedures, security protocols, protection mechanisms and auditing facilities to ensure that your data remains safe. And it is important to consider that, in many situations, the cloud provider's physical security infrastructure will likely be superior to your company's own in-house procedures and defenses. Nevertheless, the perceived threat of strangers having access to your servers is quite real. In other words, this aspect of your data's safety is one price you pay when moving to DBaaS.

Encryption as a Solution

Encryption is meant to ensure that a third party cannot access the sensitive data and when the data is sufficiently encrypted it's only accessible when someone has the key required to decrypt it. "Encryption as a Solution" only works well when the third party cannot access the key.

However, there are downsides to some encryption scenarios which may make their use problematic. These scenarios depend on the state of the data that we are interested in protecting.

02 DBaaS provider can access your data

The Three Data States

Data can exist in any of the following three states; the advantages and disadvantages of employing encryption to mitigate the data breach risk differs among them. The three states are:

Data at rest - Data stored in persistent storage media, such as hard disks and backup tapes.

Data in motion - Data in transit from one computer to another via any kind of network.

Data in use - Data stored in a non-persistent digital state, typically random access memory (RAM), ready for processing by software.

Figure 8 diagram shows available encryption technologies in DBaaS environments.

	PHYSYCAL BREACH	EAVESDROPPING/ MAN IN THE MIDDLE	FULL SYSTEM COMPROMISED
Data at rest	✓		
Data in motion		✓	
Data in use*	✓	✓	✓

*It is assumed that if the data is encrypted in use, it is also encrypted in motion and at rest.

This image illustrates the encryption technologies available to address these data states in the realm of DBaaS. The following sections describe each of the technologies mentioned here:

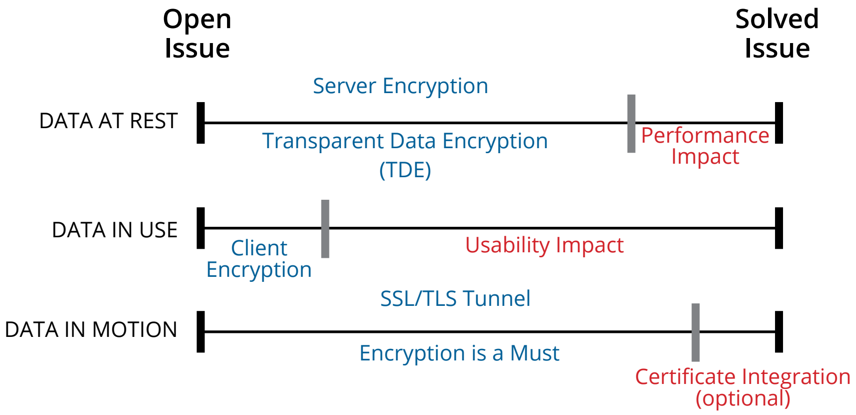

The diagram shows how each of the encryption types (at rest, in motion, in use) is compatible and feasible with DBaaS. Data in motion is the most compatible and is a de-facto standard for any cloud service communication. Data at rest is also feasible; even though it introduces a performance hit, it is highly recommended. Data in use is generally unfeasible as the encryption disables many of the database's higher functions, such as stored procedures and fast search.

Available Solutions for Each Data State

Data at Rest: Transparent Data Encryption (TDE)

TDE is a server-side encryption technology offered by database management systems (DBMS) that encrypt data at rest, and thus provide protection against a physical data breach or logical breach to the file system. Examples include stealing a hard disk, or an employee of a cloud provider with console access to a server copying data off the server.

In TDE, the DBMS encrypts data as soon as it is written into the hard-drive. This is based on file-level hardware encryption, which persists when the data is backed up or archived. TDE ensures that even a hosting service employee with direct access to the file system will not be able to copy meaningful data from the database.

While the "transparent" in TDE refers to the fact that the DBMS handles the encryption automatically, without developers having to incorporate encryption into their software code, TDE does introduce some performance degradation. This fact must be considered and evaluated for each particular database usage scenario, although it is usually recommended to enable TDE when available.

Data in Motion: Transport Layer Security (TLS)

TLS, which supplanted the earlier Secure Sockets Layer (SSL), is an encryption system that secures data in motion. In other words, TLS allows clients and servers to exchange data across a network without risking data exposure via eavesdropping or tampering.

TLS is a commonly-used security "tunnel" for client-server communications and is not specific to databases. While there is the possibility of some impact on performance, TLS should always be used to prevent the possibility of man-in-the-middle and other types of eavesdropping attacks.

Data in Use: Client-side Encryption

Client-side encryption means that all data is encrypted before it is sent to the DBaaS provider for storage. Neither the unencrypted data nor the encryption keys are ever transmitted to the DBaaS provider. This approach ensures that the hosting provider's employees never have access to unencrypted data because when the data is stored in the database-and also when it is being processed by the cloud provider's servers-the data always remains encrypted. Only when the data is sent back from the database to the client software at a later time is it decrypted, on the client side. This is the only available method to employ encryption for data in use.

While, at first glance, this sounds like an ideal solution to the concern being addressed by this chapter, client-side encryption introduces a major drawback: some core database functionality is disabled when client-side encryption is employed. This is because a DBMS is not simply a system to store and retrieve data on demand. Rather, the DBMS provides many high-level processes and functions-such as stored procedures and fast, key-based searching and sorting-that become completely inoperable when the data itself is encrypted. This is because these functions need to operate on the actual data, not an encrypted representation of it.

Most existing applications will not be able to employ client-side encryption. A hybrid approach might be to have the client side encrypt only the most sensitive data (and forgo the DBMS functionality for this data) and leave the rest unencrypted while in use.

Cloud Provider Encryption Services

The following table shows what some large DBaaS providers offer in terms of encryption. What they do not provide, organizations need to deploy themselves.

	AMAZON RDS	MICROSOFT SQL AZURE	GOOGLE CLOUD SQL
Client-side Encryption (protects data in use)	X	X	X
Server-side Encryption (protects data at rest)	✓ (TDE)	✓ (TDE)	X

Conclusion

Encryption solutions exist to protect DBaaS data at rest and data in motion. These solutions should be used in most situations. Data in use is a more challenging data state, and employing encryption may not always be feasible. Consider using client-side encryption for the most sensitive data.

Top 10 DBaaS Security Concerns

03 Physical location of the production data and the backup data

Introduction to the Threat

Before the advent of the cloud, an organization's data was almost always stored on the organization's own premises or private data centers. Security was wholly in the hands of the organization.

When data is stored in the cloud, however, the location of the data becomes a big question. Cloud providers maintain physical datacenters in multiple geographical locations for reasons of performance and redundancy. This can introduce problems relating to regulatory compliance and internal company policies. For example, some regulations prohibit organizations from exporting certain types of information outside the country in which the organization is located.

Addressing Compliance Challenges

For cloud-hosted data that must remain within a certain geographic region, there is only one solution: the organization must select a cloud provider with a datacenter within the required region. The following chart lists where the major providers maintain their data centers.

Figure 9 depicts the available DBaaS regions of major cloud providers.

AMAZON RDS	MICROSOFT SQL AZURE DATABASE	GOOGLE CLOUD SQL
US East (N. Virginia) US West (N. California) US West (Oregon) EU (Ireland) EU (Frankfurt) Asia Pacific (Tokyo) Asia Pacific (Singapore) Asia Pacific (Sydney) South America (Sao Paulo) China (Beijing) AWS GovCloud (US)	Central US East US East US 2 US Gov Iowa US Gov Virginia North Central US South Central US West US North Europe West Europe East Asia Southeast Asia Japan East Japan West Brazil South Australia East Australia Southeast	Central US (Iowa) Western Europe (Belgium) East Asia (Taiwan)

Microsoft leads with 17 regions, Amazon has 11 and Google offers only three regions.

03 Physical location of the production data and the backup data

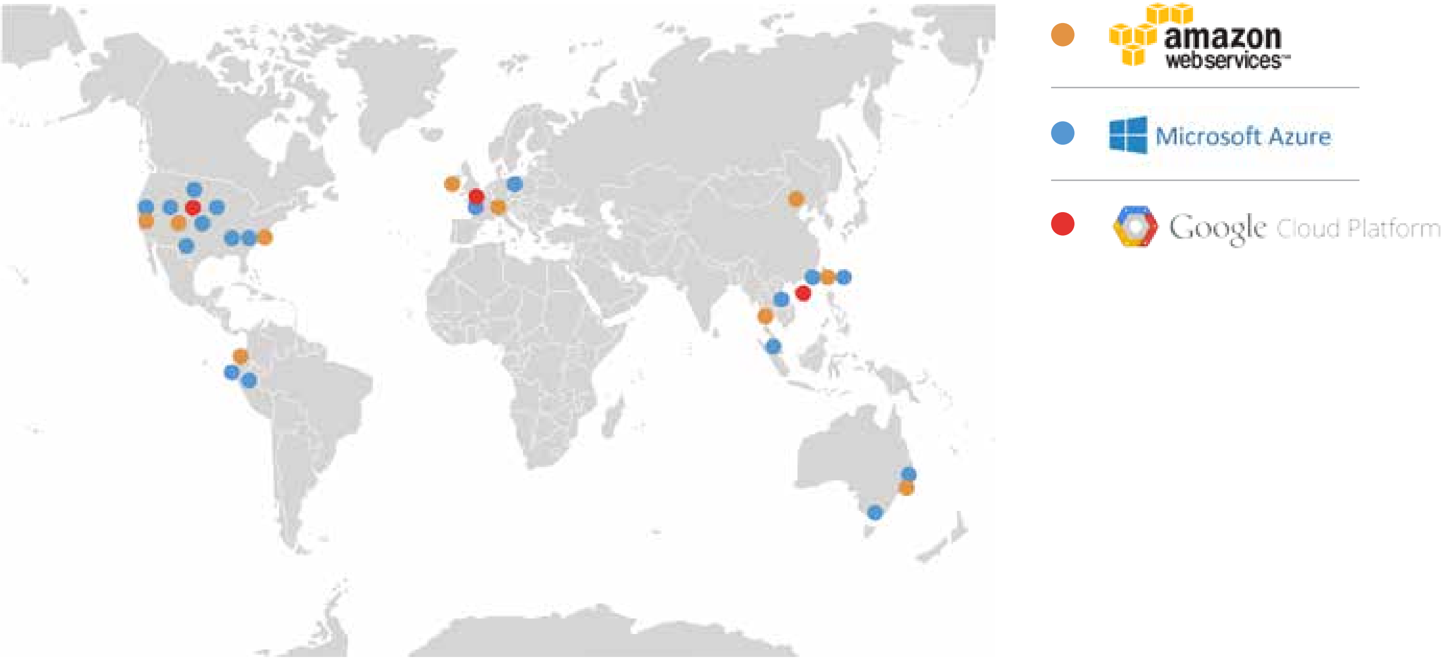

Location of DBaaS Backups

So far, we have explored the question of where the DBaaS data itself is stored, and if the selected location satisfies the organization's security and compliance requirements. However, it is equally important to ensure that these requirements are also met regarding where the DBaaS data is backed up.

DBaaS providers pride themselves on offering multiple backup and replication options. Firstly, they back up their clients' data for the purposes of their own business-continuity/disaster-recovery (BCDR) needs; the client does not have access to these backups. Secondly, they offer automatic backups that allow customers to restore data from their own backups if such a need arises. Thirdly, they offer data replication services that mirror their clients' data, sometimes to other regions, for load-balancing and performance purposes.

In light of these realities, it is critical that organizations understand that it is not enough to address only the location of the primary DBaaS data storage. Where the various types of backup data are stored is an equally important aspect to address.

Unfortunately, getting clear answers to these questions is not always easy. We researched the topic-by carefully reviewing the major DBaaS providers' end user license agreements (EULAs) and based on our own knowledge and experience- and present here our main findings:

DBaaS customers often cannot know for certain where their providers are backing up their data. Providers do not provide logs, validation processes or anything else that assures their customers that 100% of their DBaaS data is stored only in desired regions. Thus, one must assume that backup data, even if only BCDR data, is stored outside the organization's desired region, and that the provider will not notify its clients of this situation. This is one price we pay by using public cloud services.

> **Amazon RDS** – It is our understanding that none of the default options in Amazon RDS will automatically copy the database outside the customer's selected region.
>
> **Microsoft Azure** – The default in Microsoft Azure SQL Database is that the database will be copied to another region. Both BCDR backups and automatic backups provided to customers are stored in neighboring regions (examples include West US to East US, West Europe to North Europe, and Southeast Asia to East Asia). Microsoft customers must make themselves aware of any security and compliance ramifications which may result from this situation and address them accordingly.

Conclusion

Organizations do have reasonable options regarding where their cloud data will be physically located. To comply with relevant regulations, the organization must select a provider that can meet its location-related requirements.

04 External users can access the DBaaS

Introduction to the Threat

In traditional, on-premises databases, it was nearly impossible for people external to the organization to gain access to the database. This was because the data was physically stored in the company's data-centers offices and external access was blocked by network firewalls. Corporate databases would not even be located in the network DMZ.

When an organization's database is in the cloud, however, it becomes theoretically possible for people anywhere in the world to gain access. All it would take is to crack the database password or exploit some other weakness in the security configuration.

Addressing Remote Access Security Challenges

While this concern is a valid one, there are a number of ways to secure cloud-hosted databases from unauthorized external user access.

DBaaS deployments can be protected by network firewalls that function similarly to on-premises firewalls. They can be configured to allow access only from certain networks or IP addresses. This will prevent anyone else from being able to gain remote access. The big DBaaS providers all provide some level of firewall protection.

An additional level of protection is called a "database firewall." These solutions secure databases with additional layers of protection. For example, they can selectively allow or block access from particular applications or users to particular databases, tables or fields. Even in the event that a hacker manages to circumvent the cloud provider's network firewall, a database firewall will make it nearly impossible for the hacker to compromise the database itself.

04 External users can reach your DBaaS

Cloud Provider Solutions

Each of the large DBaaS providers offers its own means of securing cloud databases from unauthorized external access. Each organization should consider which offerings best suit its needs.

Amazon EC2 Security Groups	This is similar to a standard basic firewall, providing rule-based access controls. Security Groups can allow/block access based on IP address ranges, ports and Amazon EC2 security groups.
Microsoft SQL Azure Firewall	Microsoft provides basic firewall functionality based only on IP address ranges. Of the three, this is the only database firewall that goes beyond generic firewalls. As such, it can restrict access to specific databases (not only based on IP addresses). Note that Microsoft provides a database firewall because it does not provide a standard network firewall. Unlike the others, Microsoft DBaaS is a genuine SaaS offering and not a mere instance that comes with a firewall.
Google Cloud SQL	Google allows database access to be allowed/blocked based on either IP address or on specific Google App Engine applications (identified by their application IDs).

Conclusion

The concern that a cloud-based database is more accessible to unauthorized external users is a valid security issue as compared with traditional on-premises databases. Fortunately, the available methods of protecting DBaaS from external users are straightforward and effective.

All major cloud providers provide basic firewalls that, with a few simple rules, can ensure that access to the database will only be granted from the organization's own network. Furthermore, third-party database firewalls provide even more extensive and granular security beyond what native network firewalls deliver.

Therefore, with proper attention to the matter, this concern can be addressed, keeping a cloud-hosted database safe from unauthorized external users.

Top 10 DBaaS Security Concerns

05 DBaaS regulatory compliance

Introduction to the Threat

DBaaS regulatory compliance refers to the actions taken by organizations to comply with the laws and regulations that apply to databases hosted in the cloud. Commonly applicable regulations include PCI-DSS, SOX, Basel II and HIPAA.

Regulatory compliance in the DBaaS environment is more complex than for traditional, on-premises databases. This is because most regulations address aspects including where the sensitive data is stored, who can access it, what access logging is performed and so forth. The fact is that the tools for addressing these kinds of issues in the cloud are not yet as mature today as are their traditional counterparts.

Figure 10 presents the regulations organizations need to comply with.

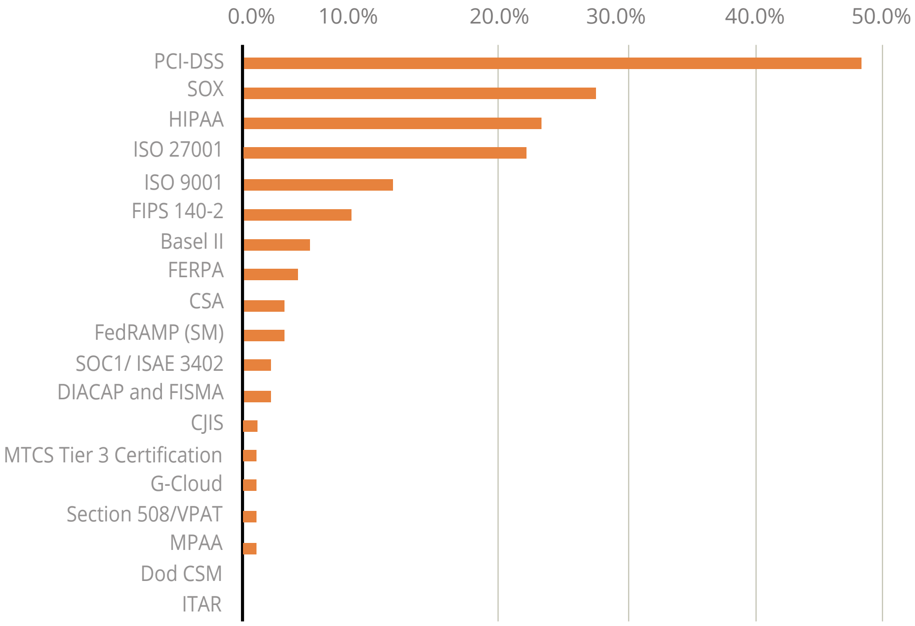

Almost 50% of the participants need to comply with PCI-DSS. SOX, HIPAA and ISO 27001 are also common requirements.

05 Difficulties with meeting regulatory compliance

Addressing Compliance Challenges

Because there are so many different laws and standards with which organizations must comply, achieving compliance can be a very complex matter-in any environment. On the other hand, there is extensive "common ground" among most regulations. The types of requirements relevant to databases include:

- The location of sensitive data
- The people who have access to sensitive data
- Control of administrative operations on sensitive data
- Logging of all events related to sensitive data

The main difference between regulations is the type of data considered sensitive. For example, PCI is focused on credit card account information, while HIPAA is focused on patient and medical information.

There are two main reasons why the above issues are more difficult to address in DBaaS environments than in on-premises databases. The first is that questions of location and access control immediately become harder to answer when the database is outsourced, and thus outside the organization's domain.

The second is that meeting all applicable regulations is challenging even with an on-premises database. What makes it more difficult in the DBaaS realm is the dearth of available tools and the immaturity of the tools that are available. It is common knowledge among IT security professionals that, while there is ongoing progress on this front, cloud-based and virtual security technologies do not yet deliver the protections of the appliances and software offered for on-premises environments, databases included. In other words, if an organization is currently using a certain set of technologies to meet compliance for its on-premises databases, it is possible that solutions with the equivalent functionality, stability and performance are not even available yet for cloud environments.

Another source of difficulty is the lack of documentation made available by cloud providers. This is another example of the fact that cloud-based compliance remains relatively immature: while cloud providers do take many measures to comply with relevant regulations, they account for it only generally (for their cloud services as a whole), and do not provide sufficient compliance-related documentation for DBaaS specifically. Organizations thus only have access to generic documentation which relates to all cloud services in general.

For example, as seen in the following screen shots, Amazon describes their compliance measures in terms of AWS only; the RDS section of their documentation does not address database compliance specifically.

05 Difficulties with meeting regulatory compliance

Figure 11 Illustrates Amazon's compliance described for AWS, but not RDS.

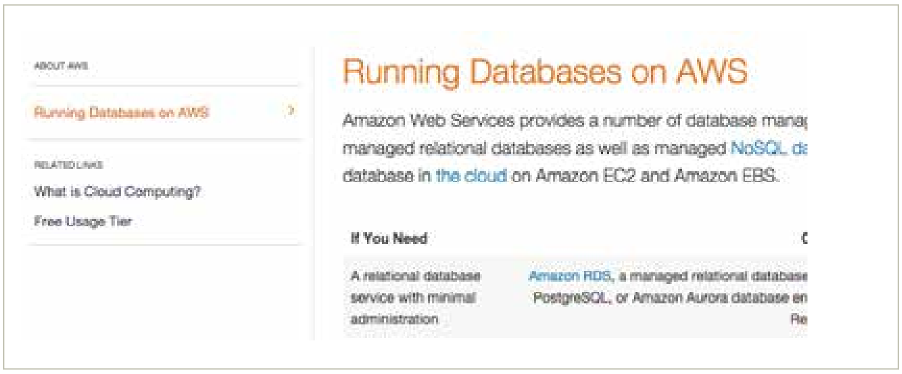

ABOUT AWS

Running Databases on AWS ›

RELATED LINKS

What is Cloud Computing?

Free Usage Tier

Running Databases on AWS

Amazon Web Services provides a number of database manag
managed relational databases as well as managed NoSQL da
database in the cloud on Amazon EC2 and Amazon EBS.

If You Need	C
A relational database service with minimal administration	Amazon RDS, a managed relational database PostgreSQL, or Amazon Aurora database en Re

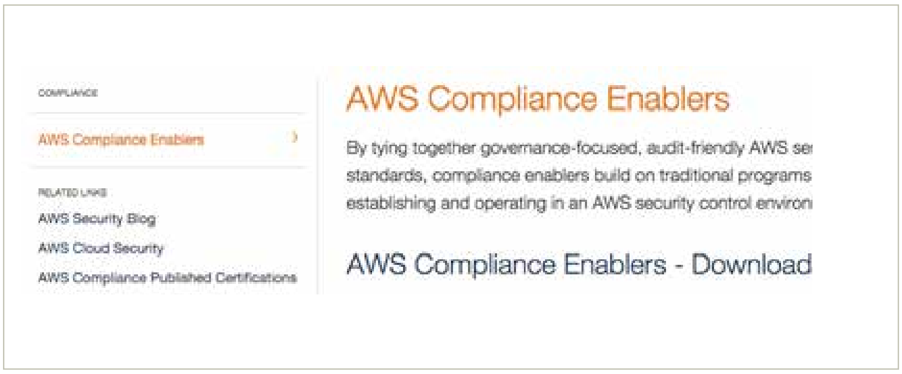

COMPLIANCE

AWS Compliance Enablers ›

RELATED LINKS

AWS Security Blog

AWS Cloud Security

AWS Compliance Published Certifications

AWS Compliance Enablers

By tying together governance-focused, audit-friendly AWS se
standards, compliance enablers build on traditional programs
establishing and operating in an AWS security control environ

AWS Compliance Enablers - Download

Amazon describes their compliance measures in terms of AWS only, while the RDS section of their documentation does not address database compliance specifically.

Relevant Technologies Offered by DBaaS Providers

The major cloud providers offer different levels of compliance:

- Amazon Web Services – Amazon's compliance center appears very mature. In addition to the available white papers and instructions, Amazon offers compliance enablers such as books, guides and a compliances toolkit.

- Microsoft Azure – Azure's compliance center is also mature. There are detailed documents that describe compliance with specific regulations.

- Google Cloud Platform – As compared with Amazon Web Services and Microsoft Azure, Google's compliance center appears much less developed, although they claim to support compliance with most of the common regulations.

05 Difficulties with meeting regulatory compliance

The following table summarizes how the major cloud providers address compliance with some common regulations.

COMPLIANCE	AMAZON WEB SERVICES	MICROSOFT AZURE	GOOGLE CLOUD PLATFORM
PCI DSS Level 1	✓	✓	✓
HIPAA	✓	✓	
SOC1	✓	✓	✓
SOC2	✓	✓	✓
SOC3	✓		✓
FISMA	✓	✓	✓
EU Data Protection	✓	✓	

Conclusion

Regulatory compliance in the DBaaS environment is more complex than in traditional, on-premises databases, and currently available solutions are less mature.

It is very likely that, in the future, regulatory compliance of DBaaS will be as easy-or even easier-than for on-premises databases, because the tools and techniques for cloud-based compliance are continuously developing and improving. In fact, because cloud providers are responsible for the physical security aspects, DBaaS compliance may even become easier for individual organizations.

Presently, any organization that has to comply with regulations in a DBaaS environment has to use third-party tools to achieve compliance.

Top 10 DBaaS Security Concerns

06 DDoS and performance attacks on the database

Introduction to the Threat

A performance attack is one that causes a database to become slow or unresponsive. This can be caused by flooding the database with artificial requests (DDoS) or by targeted attacks that impact a sensitive asset in the database (for example, queries that activate resource-intensive functions). These are indirect attacks that reach the database by way of a website that sends data to the database.

Both cloud-hosted and traditional on-premises databases can be targeted by performance attacks.

Relevant Technologies Offered by DBaaS Providers

In actuality, cloud-hosted databases are more resistant to performance attacks because of their inherent agility. The flexibility and scalability of DBaaS environments-and thus their resilience in the face of a performance attack-will always be greater than any standard on-premises database. More specifically, in the self-service DBaaS model, the customer can increase the instance's resources to handle the extra load of a DDoS attack. On the other hand, pure DBaaS deployments, such as Microsoft Azure SQL Database and Aurora, should not have performance limitations, in theory.

Even in DBaaS deployments from Amazon RDS and Google Cloud SQL, which are instance based, performance scalability is available by increasing the power of the DBaaS instance or spanning more instances.

Conclusion

It is fair to say that DBaaS offers an advantage in respect to performance attacks. While their on-premises counter parts are constrained by hardware limitations, DBaaS environments provide the scalability and agility benefits of the cloud. Nevertheless, cloud-hosted databases can still be affected by targeted attacks that abuse a specific resource or method.

A unique DoS/DDoS attack against databases

One example of a DDoS attack unique to databases is to fill the database with "garbage" data until there is no more room left in the database, or until the database becomes extremely slow. An attacker can do this by finding any website form that does not require authentication (e.g., a 'contact us' form) and re-submitting the form thousands of times a minute with lengthy garbage data. This will eventually lead to a slow database or out-of available space errors resulting in denial-of-service. It may be very difficult to diagnose this kind of attack.

Top 10 DBaaS Security Concerns

07 Hidden sensitive data

Introduction to the Threat

One of the biggest challenges facing IT security teams in large organizations is discovering all the sensitive and regulated data within the organization, such as Personally Identifiable Information (PII). The challenge of discovering sensitive data falls into two categories: finding it in structured data and in unstructured data (databases are structured data). It is obvious that before one can protect sensitive data, one must know where it is. Likewise, one cannot implement the necessary regulation compliance mechanisms-such as access control, reporting, auditing and masking-without identifying data covered by regulations.

In a DBaaS context, discovering sensitive data is especially important once an organization decides to move data to the cloud: before migrating regulated data to the cloud, each table needs to be assessed to make sure that moving it to the cloud satisfies all relevant regulations.

Relevant Technologies

Manual attempts to identify and locate all the sensitive and regulated data in an organization are very difficult. Typically, the individuals who introduce sensitive data into the database are not the same people responsible for data security and for complying with data-related regulations. Therefore, those in charge of security and/or compliance are frequently unaware of the existence and location of much of the organization's sensitive and regulated data.

Fortunately, third-party vendors offer technology that is capable of scanning all of an enterprise's databases and automatically discovering sensitive data. These tools can identify specific data fields (e.g., SSN, credit card numbers, e-mails, passwords), based on particular security concerns and/or specific regulations (e.g., PCI, HIPAA, SOX).

Once discovered, tools are available to provide one-click activation of access control, access auditing and data-masking features for those fields, fully addressing most or all security and compliance requirements.

Where Does Hidden Sensitive Data Come From?

How is it possible that sensitive data, such as PII, can suddenly appear in unprotected database fields? Here are two common scenarios:

• An organization deploys an off-the-shelf CRM application that includes standard forms to be filled out by customers. Even if employees are well-trained and careful to only input sensitive data into specific fields, customers frequently enter PII (even social security and account numbers) into unprotected fields via these forms.

• An employee is sometimes faced with the need to record sensitive data into a form that cannot contain all of it in protected fields. For example, a customer representative speaking with a customer might have to record two credit card numbers or two identification numbers, while the application only contains one designated field. In these situations, the employee may enter the sensitive data into an unprotected "comments" field.

Provided by DBaaS

To date, none of the major DBaaS providers offer tools for the automatic discovery of sensitive and regulated data.

The third-party tools available can address both on-premises and cloud-hosted databases. This essentially allows one-click activation of polices for database monitoring, auditing and data-masking for every sensitive data field in an organization's cloud-hosted databases.

Conclusion

Finding sensitive and regulated data within enterprise databases is one of the key challenges to implementing data security and regulatory compliance. While the major DBaaS providers do not offer tools for the automatic discovery of sensitive and regulated data, third-party tools are available to address this challenge.

Top 10 DBaaS Security Concerns

08 SQL injection attacks

Introduction to the Threat

SQL injection attacks remain the most prevalent database breach method in use today. In an SQL injection attack, malicious SQL statements intended for execution by a database are inserted into an entry field on a website. One possible result is that the server will expose data from the database that it should never return. These types of attacks exploit a security vulnerability in a Web application's software and sometimes also in a closed source application.
The first line of defense against SQL injection attacks is one that must be implemented by developers on the code level: never blindly trust user input. This means that all input to the SQL engine must be validated before it is passed on to the database. The three drawbacks of relying on this approach, however, are that it is difficult to enforce and verify, it adds time and cost to software development and cleverly-constructed queries may still slip through the validation routines. For these reasons, many applications remain vulnerable to SQL injection attacks.

Relevant Technologies

As user input that eventually reaches the database arrives mostly from a Web server, a second line of defense is a Web application firewall (WAF) that can “sanitize” most dangerous requests at the application layer.
A deeper line of defense, not less important, that protects the database directly, is a database firewall. Database firewalls are deployed, as proxy servers or in a bridge model, between the database and all applications which access it, and filter all data moving in and out of the database. This allows the system to identify and prevent malicious attacks by comparing every query's structure with a signature bank of known attacks. Suspicious or dangerous queries never reach the database, preventing SQL injection attacks from succeeding.

Provided by DBaaS

To date, only Amazon Web Services provides WAFs that specifically filter out SQL injection attacks. The third-party WAFs and database firewalls available can protect both on-premises and cloud-hosted databases from SQL injection attacks.

Conclusion

Protecting cloud-hosted databases from SQL injection attacks is critical. The major cloud providers do not currently offer off-the-shelf solutions to this threat; instead, it is expected that third-party solutions be used.
While there are mature solutions provided by third parties, organizations may find that their preferred vendor does not have a mature solution for DBaaS and need to look elsewhere. It is strongly recommended that organizations select and deploy a third-party database firewall solution designed to identify and block SQL injection attacks.

Top 10 DBaaS Security Concerns

09 Data theft by authorized users

Introduction to the Threat

DBAs, software developers, quality assurance personnel and others (whether employees or external partners) frequently require extensive access to databases in order to perform their roles. This essentially means that the organization cannot block their access, even to sensitive databases, without preventing them from doing their jobs. On the other hand, it is dangerous and unnecessary for these types of roles to have total and unfettered access to all of a company's sensitive data.
This is a very real concern, equally relevant to cloud-hosted and on-premises databases.

Relevant Technologies

There are two ways to address this threat.
The first, already described in this report, is a database firewall. A database firewall can selectively allow or block access from particular users to particular databases, tables and even fields. This solution will eliminate the threat in situations where authorized users can do their jobs without access to the particular databases, tables or fields containing sensitive information (this benefit of database firewalls is known as separation of duties, or segregation of duties).

The second solution available is dynamic data masking. Also known as on-the-fly data obfuscation and real-time data scrambling, dynamic data masking is a technology that sits between databases and client applications and obfuscates any sensitive data sent to application screens, reports, development environments, DBA tools, etc. In other words, dynamic data masking removes or changes only the sensitive data on its way out of the database so that the client application receives meaningless data. Typical examples of data masking include credit card and account numbers (all digits are shown as zeroes except the last four digits), email addresses (everything before the @ sign is masked or randomized) and currency amounts (all amounts are shown as $12.34 or some other selected figure).

Dynamic data masking can be based on masking policies which remove or obscure particular fields for specific users/roles. For example, data masking policies can be defined per data table column, user/role, client IP address and/or client application. Based on these policies, protected data fields are masked or randomized in real-time using a context-relevant template, allowing users and applications to function normally, without ever gaining access to sensitive data. It is also important, of course, that the masking be applied during every kind of access, including stored procedures and views.

09 Data theft by authorized users

Dynamic data masking is a powerful security tool which allows users to work with databases to fulfil their job functions without ever gaining access to sensitive data fields-and is thus a tool which dramatically reduces the risk of insider data theft.

Provided by DBaaS

None of the major DBaaS providers currently offer a database firewall that provides a complete and granular separation-of-duties solution that can selectively allow or block access from particular users to particular databases, tables or fields. To implement this level of protection, a third-party solution is required. While DBaaS providers do provide generic firewall capabilities (see the External Users can Access the Database chapter, above), these solutions do not provide table or field level granularity. The most granular option currently available is from Microsoft: their Azure SQL database firewall can selectively allow/block access to specific databases.

Among the major DBaaS providers, only Microsoft's Azure SQL Database provides dynamic data masking functionality.

Conclusion

Data theft by authorized users will always be a threat in any organization that handles sensitive or regulated information. On a technological level, the best an organization can do is to mitigate the threat by limiting access to sensitive information to only those employees and contractors who absolutely require that access.

The two technologies available to address this threat are both policy-based: database firewall and dynamic data masking. Among the major DBaaS providers today, the former requires the deployment of a third-party solution. Microsoft provides dynamic data masking functionality as a Preview feature in Azure SQL Database, while the other providers do not currently offer it. Third-party dynamic data masking solutions, which are available for cloud databases, are a recommended approach to limiting unnecessary access to sensitive data.

Top 10 DBaaS Security Concerns

10 Neighbor tenants can access your data

Introduction to the Threat

Cloud-hosted applications and databases typically utilize resources that are shared by multiple customers, or tenants. Two of these shared resources-memory and disk-present the danger that one tenant may be able to access sensitive data belonging to a different tenant.

This chapter discusses two concerns related to the fact that, in a DBaaS environment, your data might become accessible to another tenant using the same hosting service:

1. A neighbor tenant gains indirect access to your data. This can occur, for example, if an area of memory or disk was recently used by your database without encryption and was not properly erased before being offered for use by another tenant. In this situation, it might be possible for the other tenant to access remnants of your data there. This was concern #10 in our survey.

2. A neighbor tenant infiltrates your database. Hypervisor technology aims to ensure that one tenant can never access data belonging to another tenant sharing the same resources. However, if another tenant sharing your resources is able to maliciously compromise the hypervisor in a deliberate attempt to access your data, your data may be at risk. This was concern #11 in our survey.

Relevant Technologies

One of the key responsibilities of a hypervisor is to sanitize all resources before provisioning them to a tenant, specifically to avoid the possibility of the first concern occurring. However, the fact is that there are occasionally reports of data remnant vulnerabilities occurring, indicating that this concern remains a valid one.

The second concern, that a neighbor tenant might be able to compromise the hypervisor and gain access to your data, is no worse than the similar concern, previously discussed, that your DBaaS provider can access your data. The only difference is that the neighbor tenant has to hack its way to your data while the provider has more direct access.

Therefore, a similar solution is in order: to encrypt your sensitive data. As discussed in detail in the chapter about the risk of your provider accessing your data, our recommendations are to employ data-at-rest encryption and, where comprehensive data-in-use encryption is not practical, to at least encrypt very sensitive fields privately (on the client side).

10 Neighbor tenants can access your data

Provided by DBaaS

See the discussion in the earlier chapter:
Your DBaaS Provider Can Access Your Data.

Conclusion

As mentioned regarding risks stemming from your DBaaS provider's ability to access your data, there is no way to fully avoid these tenant-based threats. These risks are part of the price you pay for enjoying the benefits of sharing resources with other tenants.

Summary

Main Findings

The following table summarizes how each of the concerns is addressed today by the three leading DBaaS providers.

	AMAZON RDS	MICROSOFT SQL AZURE DATABASE	GOOGLE CLOUD SQL
	Your DBaaS Provider can Access your Data		
Client-side encryption (protects data in use)	✗	✗	✗
Server-side encryption (protects data at rest)	✓ **(TDE)**	✓ **(TDE)**	?
	Location of the Data		
Coverage	✓	✓+	✓
Number of sites	11	17	3
	External Users can Access the Database		
Basic firewall (infrastructure only)	✓ Amazon EC2 Security Groups	✓ Microsoft SQL Azure Firewall	✓ Google Cloud SQL
	Compliance		
Generic cloud compliance support	✓ Mature	✓ Mature	✓ Basic
DBaaS-specific compliance instructions	✗ None provide compliance instructions specific for DBaaS.		
	DDoS and Performance Attacks on the Database		
Elasticity to absorb DDoS	✓ Instant elasticity is a benefit of cloud services		
	Hidden Sensitive Data		
Hidden sensitive data discovery tools	✗	✗	✗
	SQL Injection Attacks		
WAF	✓	✗	✗
Database firewall (with relevant SQL mitigation technology)	✗	✗	✗
	Data Theft by Authorized Users		
Database-level sepatation of duties	✗	(only on a database level)	✗
Dynamic data masking	✗	SQL Database Dynamic Data Masking	✗

Summary

Conclusions for One Evaluating DBaaS

IDC reports that security is the #1 concern regarding the adoption of the cloud. This statement is supported by the results of HexaTier's survey which also found that security is the #1 concern preventing companies moving their most sensitive business data to the cloud. On the other hand, IDC says that the cloud can be more secure than organizations' own on-premises environments.

Today's organizations understand the cost, scalability and elasticity benefits that DBaaS offers. They also understand that they must deal seriously with IT security in order to continue driving their business forward and winning the "hearts" of their customers.

This research attempted to close the gap between the market's top DBaaS security concerns and the ability of the cloud provider to address them. This gap is real and primarily exists for two particular reasons:

- The major cloud providers do not yet deliver some of the most common security products/capabilities available for on-premises data centers.
- Most security vendors have not yet adapted their products to fit within the cloud's network infrastructure.

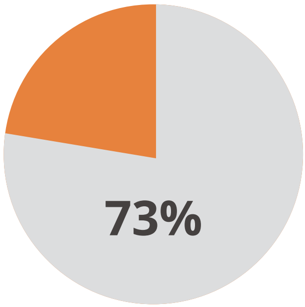

73% of the survey participants indicate that data security and compliance is the #1 concern when moving to DBaaS.

There is no simple yes-or-no answer for an organization trying to decide whether to move to DBaaS or not. The answer is organization-specific and centers around the ideal balance between the cloud's functional advantages and its security concerns. It is important to note that that the term "one evaluating DBaaS" can be misleading in certain scenarios because the decision of whether to move to the cloud or not is a strategic decision within which DBaaS is only one component.

The following section addresses what you actually need to do once you have decided to adopt a DBaaS solution.

Conclusions for One Currently Using DBaaS

Here are our recommendations for organizations already using DBaaS, summarized in one table. The reasoning behind these recommendations can be found in the above chapters. The links in the table point to external websites and provide starting points for the included recommendations.

Summary

CONCERN	AMAZON RDS	MICROSOFT SQL AZURE DATABASE	GOOGLE CLOUD SQL
	Your DBaaS Provider can Access your Data		
Protect data at rest by enabling server-side encryption	Instructions	Instructions	✗ Not supported
Privately encrypt your most sensitive data	Identify the most sensitive fields and encrypt them on the client side using your own libraries.		
	Location of the Data		
Find out where your data is stored	1. Determine which regulations apply to your organization. 2. Work with your cloud provider to ensure that your data is located where it should be. 3. Ensure that your backup locations also comply with all relevant regulations.		
	External Users can Access the Database		
Make sure your cloud firewall is properly configured	Amazon EC2 Security Groups	Microsoft SQL Azure Firewall	Google Cloud SQL
Deploy a third-party firewall	Consider a third-party firewall that has more extensive features than those offered by your cloud provider.		
	Compliance		
Utilize all the available materials and compliance kits	- AWS Compliance - Compliance Enablers	- Microsoft Azure Trust Center: Compliance	- Google Cloud Platform Security - Compliance
Deploy third-party technology	Utilize third-party technologies to meet compliance. Check if your preferred vendor has cloud-supporting technology. If not, find alternatives.		
	DDoS and Performance Attacks on the Database		
Measure your capacity	Utilize third-party technologies to meet compliance. Check if your preferred vendor has cloud-supporting technology. If not, find alternatives.		
Retain the mitigation services of a third party	A cloud-hosted database will not be attacked by DDoS directly, but by an indirect attack coming from an application (typically, a Web application). There are various services (and even virtual appliances) that can protect your application and thus your DBaaS instance.		
	Hidden Sensitive Data		
Deploy third-party tools	Use third-party tools to detect hidden sensitive data.		
	SQL Injection Attacks		
Deploy a third-party WAF	The first line of defense is a Web Application Firewall (WAF), available either as a virtual appliance or in the SaaS model.		
Deploy a database firewall	The last line of defense is a database firewall with relevant SQL threat-mitigation technology.		
	Data Theft by Authorized Users		
Implement a Separation of Duties (SoD) solution			
From your DBaaS provider	✗	✓ Microsoft SQL Azure Firewall (only on a database level)	✗
From a third party	Use a third-party database firewall that has field-by-field level SoD technology.		
Implement dynamic data masking			
From your DBaaS provider	✗	SQL Database Dynamic Data Masking	✗
From a third party	Use a third-party database firewall that has dynamic data masking.		

Summary

Can the Cloud Actually be More Secure?

Finally, we cannot conclude this discussion of DBaaS security without presenting the claim of many security experts and researchers that the cloud is actually more secure than an organization's on-premises data center. This position is held by professionals who are well-aware of the many valid concerns, but attribute the fears more to the newness and unfamiliarity of the cloud instead of a logical comparison of the realities of the two environments.

After all, can a typical organization compete with the overall security that giants like Amazon, Microsoft and Google are capable of implementing? It is important to consider both the resources and expertise that these companies can bring to bear, and the fact that these companies' entire cloud businesses are dependent on rock-solid security. Yes, you hand over a great degree of control over your data, but chances are that they are providing you with physical security, network security, software patching and so forth that are far better than your team would be able to do itself.

This issue is a topic for a paper of its own, given the strong opinions on both sides, and the varying definitions of, and requirements for, IT security.

Understanding Database Security

Introduction

While most organizations address IT security in terms of policies of who has what level of access to what resources, the fact is that the "attack surface" in cyber security is much wider than this.

From our experience working with enterprises that have been breached, and conducting forensic analyses following breaches, we see a clear pattern: Even though in more than 90% of the cases the information was stolen from a database, the database management system (DBMS) itself was not under attack. Instead, the attacker was able to compromise other components of the IT ecosystem in order to steal information from the database. The primary targets attacked were:

1. Application servers that have access to the database (most commonly via SQL injection attacks)
2. The database server's operating system
3. DBA workstations (because they have stored credentials to directly access the database)
4. IT staff workstations

In fewer than 10% of the breaches we researched, the DBMS itself was under direct attack using specific database exploits.

The Scope

In order to truly understand how to protect sensitive databases, it is critical to map all the attack surfaces which have access to the database, and to systematically secure each one. Because of the current trend of migrating databases to Database as a Service (DBaaS), the relevant scope is changing and evolving as well. This section presents a high-level mapping of the database security scope.

1. Automated and user connections

There are basically two types of connections accessing most databases:

> **a. Automated connections** – These are connections made by applications that have a stored connection with access credentials, enabling on-demand access to the database. In most cases, the types of access performed via these connections are predictable (i.e., specific parts of the database are accessed via specific queries performing specific operations).
>
> **b. User-based connections –** These are connections made by people, such as DBAs, developers, QA teams and IT staff. These people are entrusted with the credentials required to access the database directly, and can typically perform a wide range of unpredictable actions.

The first step in controlling database access and mitigating unauthorized activity is mapping and understanding all the authorized channels to the database, and what each one should and should not be doing.

Beyond user name and password credentials-which provide only the most elemental form of access security-it is important to implement and map other restrictions to minimize the possibility of unauthorized database access. Examples include IP address and port filtering, and two-factor authentication.

2. Applications connecting to the database

Mapping and enforcing all the particular applications that require access to the database is the next step to increasing database security. Most database protocols contain the application name used to connect to the database; there are solutions available to restrict access to specific applications. So, even if an application server is breached, connections from that same server or operating system will be refused by the database, if the same application name is not used. This is not an iron-clad defense, but it makes it harder for a compromised server to become a conduit to your most sensitive information.

3. Network layer and connecting drivers

Each application uses a specific type of driver in order to connect to the database. Organizations sometimes unknowingly permit older versions of drivers to connect, even though they utilize more vulnerable connection protocols. This often occurs when compatibility support is enabled for older applications, without knowledge of the security implications.

Awareness and enforcement of which protocol versions are permitted to access the database, and blocking unneeded legacy protocols, provides an additional level of database security.

4. Database server operating system

Hardening the database server operating system is crucial. Reducing the attack surface of the database server significantly increases the level of database security. The more functions a system fulfils, the wider the attack surface. Removal of unnecessary third-party applications, disabling unused operating system services and limiting operating system access as much as possible are all crucial steps.

Oftentimes, attackers accessing the database via the server operating system gain the ability to steal entire database files, allowing them to easily open them on a different server. This attack vector is particularly difficult to detect as it will not generate any log events in the database or in a third-party Database Activity Monitoring solution.

5. Database Management System (DBMS)

It is important to keep the DBMS system updated with the latest vendor patches and updates. Besides increasing security, this also increases the reliability and availability of the database. Because attackers are also constantly looking for weaknesses in the latest updates, this is a never-ending process.

6. Stored Data Files

Each database stores data in the form of files in order to maintain persistence. Because these files are often stored, mirrored or backed up on remote storage systems, organizations also need to control who has access to these additional storage locations.

7. The data itself and the different way to retrieve it (stored procedures, views, etc.)

It is critical to understand which columns contain sensitive data and which types of access should be permitted to retrieve or modify this data. Armed with this knowledge, organizations are in a much better position to create policies that can actually enforce who, when and how sensitive data can be accessed/modified. On the other hand, without mapping the locations of the organization's sensitive data, much security infrastructure will be rendered nearly useless.

Furthermore, many regulations define sensitive information differently. Understanding and mapping sensitive data according to the applicable regulations is also critical to achieving regulatory compliance.

Encryption used to be the most reliable way to protect data, but, as described elsewhere in this document, encrypting data at rest in a DBaaS environment prevents using some of the core database functions (e.g., sort, search). Additionally, any application that contains the keys to the encrypted data poses a weak link in the organization's security posture: if the application is breached, the encrypted data may be breached as well.

Controlling the exposure of sensitive information is an important task. Even if the data at rest remains unencrypted, there are still multiple solutions to help control the exposure of the organization's sensitive information.

Traditional database security scope summary

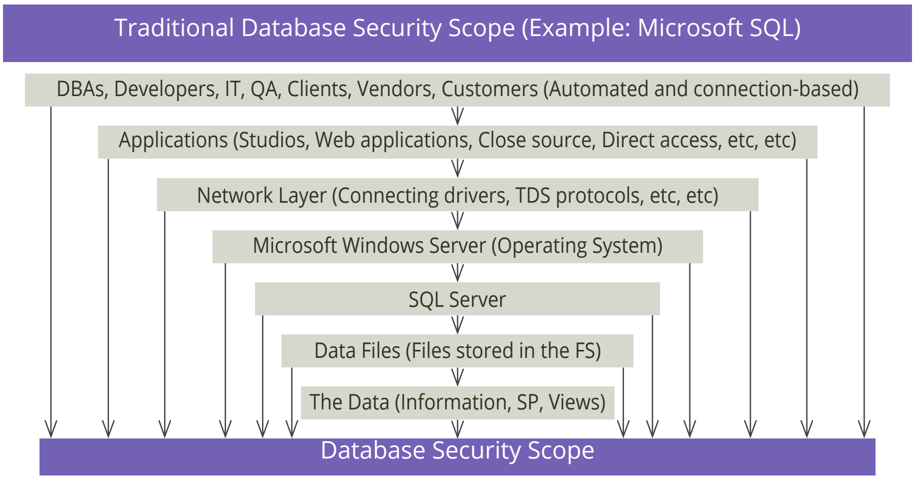

Traditional Database Security Scope (Example: Microsoft SQL)

DBAs, Developers, IT, QA, Clients, Vendors, Customers (Automated and connection-based)

Applications (Studios, Web applications, Close source, Direct access, etc, etc)

Network Layer (Connecting drivers, TDS protocols, etc, etc)

Microsoft Windows Server (Operating System)

SQL Server

Data Files (Files stored in the FS)

The Data (Information, SP, Views)

Database Security Scope

Moving to Database as a Server (DBaaS)

When moving to, or considering moving to, a DBaaS solution, much of the attack surface changes. This forces organizations to focus on fewer points, namely those that are still under the organization's control. This does not mean that the organization's data is immediately safer, because DBaaS also introduces numerous additional questions that can't be easily answered, as addressed throughout this paper.
In DBaaS, the layers that are fully under the control of the DBaaS provider are:

- Network layer – connecting drivers
- Database server operating system
- Database Management System (DBMS)
- Stored data files

Without control of these layers, the organization's focus shifts to the data itself, including access control. From the organization's perspective, the data security scope in a DBaaS environment is:

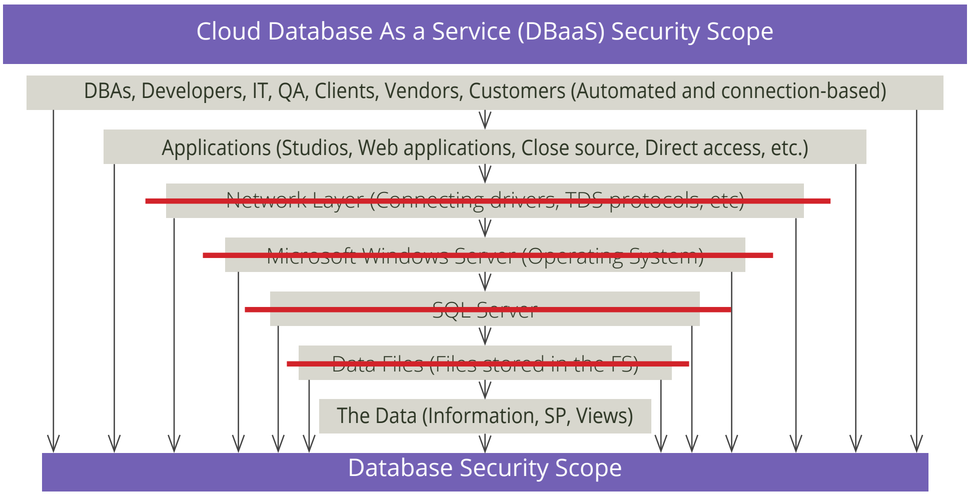

Cloud Database As a Service (DBaaS) Security Scope

DBAs, Developers, IT, QA, Clients, Vendors, Customers (Automated and connection-based)

Applications (Studios, Web applications, Close source, Direct access, etc.)

~~Network Layer (Connecting drivers, TDS protocols, etc)~~

~~Microsoft Windows Server (Operating System)~~

~~SQL Server~~

~~Data Files (Files stored in the FS)~~

The Data (Information, SP, Views)

Database Security Scope

Virtual Private Cloud

As mentioned in the first chapter, this paper focuses on comparing traditional on-premises databases to DBaaS, but does not consider a third model, the Virtual Private Cloud (VPC). In the VPC model, the hosting provider gives you complete access to infrastructure upon which you build your data center. To run a database in VPC, you need to install the operating system, the database application and load in the data.

Although VPC was not a topic of this paper, a short overview is in order.

A VPC-hosted database is, perhaps, best suited for organizations that want to gain the cost reduction and agility benefits that the cloud offers, but are not willing to pay the full security price of DBaaS. However, not paying the full price also means not gaining all the benefits.

The main advantage of VPC, as compared with DBaaS, is much better control over your data. Since, in VPC, you run your own application on top of your own operating system and, assuming you encrypt the data at rest, then you essentially prevent the cloud provider from having any access to your data. This is something that DBaaS cannot achieve (see the chapter, Your DBaaS Provider Can Access Your Data).

Also, standard practice with VPC is that access to the database is via VPN tunnel or other encrypted means. This ensures that no unauthorized external users can access your database.

Another advantage of VPC is that you have more control over where your data is located. VPC is a very popular service which you can almost certainly obtain in your own country/region.

Regulatory compliance is simpler with VPC, but still not as simple as for an on-premises database. Like in DBaaS, many applications required to validate compliance are still not mature enough for VPC environment.

Perhaps the most important aspect to understand with VPC are the security risks you accept by taking complete responsibility over the operating system and database application. With VPC, it is your sole responsibility to secure and harden these critical layers, and to continuously patch them as new vulnerabilities are addressed by software vendors. It is unlikely that you will be able to do this as efficiently and reliably as the major DBaaS providers do.

Figure 12 presents the shared responsibility – on-premises vs. VPC vs. DBaaS.

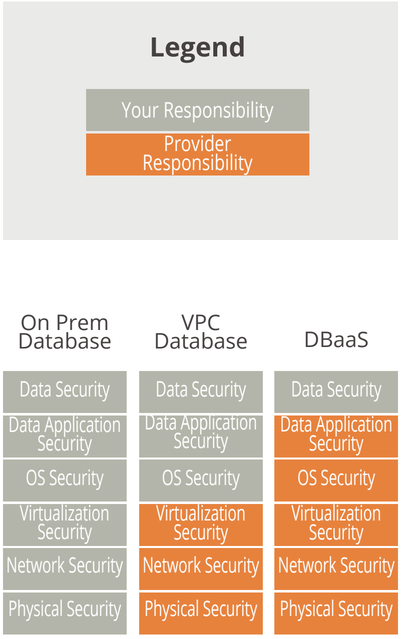

Legend

⬜ Your Responsibility
🟧 Provider Responsibility

On Prem Database	VPC Database	DBaaS
⬜ Data Security	⬜ Data Security	⬜ Data Security
⬜ Data Application Security	⬜ Data Application Security	🟧 Data Application Security
⬜ OS Security	⬜ OS Security	🟧 OS Security
⬜ Virtualization Security	🟧 Virtualization Security	🟧 Virtualization Security
⬜ Network Security	🟧 Network Security	🟧 Network Security
⬜ Physical Security	🟧 Physical Security	🟧 Physical Security

hexatier
Secure, Comply, Go!

About HexaTier

Established in 2009, HexaTier sets the industry standard for database security and compliance in the cloud with its software-based unified solution that provides Database Security, Discovery of Sensitive Data, Dynamic Data Masking and Database Activity Monitoring. Utilizing purpose-built, patented Database Reverse Proxy technology, the company protects against both internal and external security threats.

Backed by leading investors such as JVP, Magma VC and Rhodium, HexaTier is the first company worldwide to provide security for cloud-hosted databases and DBaaS platforms through a streamlined and simple solution. Headquartered in Tel Aviv with offices in Irvine, CA and Boston, MA. HexaTier secures databases for nearly 200 organizations globally.