Improvising Database Intelligence in Cloud Databases

Tinesh Kumar Goyal, Amitesh Kumar Gupta, Asish Srivastava, Piyush Saxena

***Abstract*— A cloud database is a database that is typically deployed on cloud platform and it is introduced so that data in cloud can be managed wither by user using virtual systems or by the cloud service provider. The major advantage of using cloud database is being a virtual storage location resource; the data can be managed irrespective of any location. Migrating from traditional relational databases to cloud database has many advantages of its own in the cloud but to maintain the security intelligence of the database is one major challenge. In this paper we try to state the challenges faced for securing cloud databases and probable mechanisms of how to improvise them.**

If we streamline the approach of security protocols to be followed for effective database intelligence, we have to identify the assets we want to protect, the probable threats and the relative countermeasures for them. The cloud service provider will be responsible to take care if its operations, facilities, network, hosts and other components are secure but ultimately it would be the user who would be responsible if the database is secured or not. Cloud providers offering database hosting need to consider database security as a critical service to their customers. Our proposed protocol of IEEMR comprising of essential five steps to protect the data in the cloud database.

***Index Terms*— Cloud Computing, Cloud database security, Security, IEEMR, security in public cloud, Sensitive database, Security as a Service (SaaS), DBaaS (Database as a Service).**

I. INTRODUCTION

Many organizations are looking to cloud-based IT infrastructures as a means of solving scalability, performance, availability, and cost problems. There are three basic deployment models for cloud infrastructures, public, private and hybrid cloud. The public cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services. The private cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise. The hybrid cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability.

Regardless of deployment models, as companies begin to move their databases into the Cloud, database security programs need to follow. Cloud infrastructures providing database hosting, offer very basic security in the form of network-layer firewalls configured with port-based access control lists. Database platforms are often provided as a service where the security details are abstracted away from the user. Those users are offered nothing more than a contract or Service Level Agreement that makes empty claims about security. This basic layer of security, coupled with the increased exposure database applications receive from being deployed in shared and possibly public infrastructures require organizations to review their database security strategies and implement a comprehensive program of database security process control.

As mentioned earlier, there are other emerging cloud services that fall into one of the SPI models. Database as a service is one the new emerging services that can be a subtype of SaaS or PaaS depending on the way it is delivered by the providers. DBaaS provides consumers an on-demand database services in the cloud that can be accessed by the users through the internet. Database-as-a-service makes an efficient use of cloud computing technology by providing businesses with easy access to scalable, on-demand database resources while avoiding the costs and complexity associated with the purchase, installation and maintenance of a traditional on-premise database system.

An important question in Public and Hybrid Cloud infrastructures is, “Who owns responsibility for securing the database and sensitive data?” Ultimately, the data owners need to take responsibility for securing their sensitive data. Data owners should formulate their database security strategies and then partner with their cloud providers to ensure effective implementation. Cloud providers offering database hosting need to consider database security as a critical service to their customers. Whether in Traditional, Private cloud or Public cloud infrastructures, Application Security, Inc. believes it is important for organizations to protect their sensitive data where it lives – in the database. DbProtect Precision Database Activity Monitoring (DAM) helps organizations to protect their cloud based data assets by providing control over the security processes that impacts their sensitive data.

Since virtualization is now in use by nearly all enterprises— and future plans to move some applications to cloud computing are also in the making—it’s time to

Manuscript received May 04, 2014.

Tinesh Kumar Goyal, M.Tech (Computer Science and Engineering), Amity University, Amity School of Engineering and Technology, Noida, India, +91-9417170800.

Amitesh Kumar Gupta, M.Tech (Computer Science and Engineering), Amity University, Amity School of Engineering and Technology, Noida, India, +91-7532836750.

Asish Srivastava, M.Tech (Computer Science and Engineering), Amity University, Amity School of Engineering and Technology, Noida, India, +91-9555531144.

Piyush Saxena, M.Tech (Computer Science and Engineering), Amity University, Amity School of Engineering and Technology, Noida, India, +91-9451427546.

ask whether we need to update our IT security methodologies or continue to rely on the same tools we have relied on in the past. Many would agree that these computing models require an entirely different approach. Sometimes, it seems that network security models lag behind the technology changes occurring in the systems and software environment they are intended to protect. For example, it took firewall vendors a long time to realize that applications are no longer easily protected by simply opening and blocking ports. As enterprises began deploying VoIP and other complex protocols, they typically had to wait years before their firewall vendor allowed them to do so securely by analyzing the session initiation protocol (SIP) and other protocols. With the recent shift in computing architecture from dedicated servers in the data center to virtualization and cloud computing, we need to rethink our IT security methodologies. And, while we focus on database security here, many of our recommendations are applicable to securing most enterprise applications.

As studied from the above trend, the key trade-off for security between the different service models is that more and more security responsibilities shift towards consumers when the provider stops lower down the stack, meaning that consumers, for example, have greater responsibility for management and security of the services if they use IaaS instead of PaaS or SaaS. There are basically two main approaches to deploying databases in the cloud:

1. "Do-it-Yourself" – In this approach, a consumer purchases IaaS services from a cloud provider of their choice and installs their own database instance on the provided platform. All the responsibilities for maintenance, management and patching of the database instances will be with the customer. However, it does offer the benefits of installing additional monitoring and auditing tools together with the database instance on the virtual platform.
2. Subscribing to a full-fledged Database-as-a-Service (DBaaS) – In contrast to the above, in this approach, consumers get the opportunity of buying whole database package as a service from a provider. All the management and maintenance responsibilities will be with the provider and the only thing the consumer would be required to do is to connect their applications to the databases and they would be charged on per-usage basis. This approach is much more convenient for consumers but the disadvantage is that the providers often do not support installation of additional tools together with the cloud databases on their platforms. This makes it less secure as consumers would be expected to trust on provider's monitoring and auditing tools.

II. II. CHALLENGES & THREATS FOR DATABASE

Migrating databases into cloud environment brings a number of security concerns that organizations have to take into consideration as the ultimate responsibility for data security assurance is with organizations and not with providers. When internal databases with sensitive data are migrated to the cloud, users need to be assured that proper database security measurements are in place that encompasses data confidentiality, integrity, and availability. The main aspects of database security in the cloud are that data must be secured while at rest, in transit, and in use, and access to the data must be controlled. That is to say:

1. In order to assure that data does not get corrupted or hijacked, it is very important to have safe procedures in place that would protect data transfer to and from the databases that reside in the cloud.
2. To ensure high confidentiality, it is important that the outsourced data in the stored in cloud databases be encrypted at all times.
3. To ensure high integrity, the access to the data stored at cloud database provider's platform needs to be controlled and monitored properly for all users including the database administrators at the data center.
4. Visibility into virtual machine-to-virtual machine transactions.
5. Architecture complications arise from the fact that instead of creating a dynamic environment on virtual servers, enterprises need to plan ways in which all access to databases will pass through the virtual appliance.
6. Performance over wide area networks monitored more effectively.

Availability in simple terms means the extent to which system resources are accessible and usable to individual users or organizations. It is one of the critical security aspects that organizations need to take into account when considering cloud database services. In the wake of a failure, availability can be affected temporarily or permanently, and the loss can be partial or complete. There are many threats to availability that include DOS attacks, equipment failures and natural disasters. Therefore, when addressing database availability with the vendor, consumers should always demand for the high availability standard known as the five nines. It equals to an uptime of 99.999%, which is like an outage of about five minutes per year.

Although elasticity and flexibility is considered amongst major benefits of cloud computing but it brings an inherent security issue with it. In order to satisfy consumer needs, cloud databases scale up or down frequently which means that physical servers that host databases gets provisioned and de-provisioned often without prior knowledge of consumers. Moreover, in order to provide high availability and redundancy, customer data is usually replicated across several data centres in multiple locations. All of these factors result in non-static environment where consumers have almost no visibility or accessibility to the physical infrastructures. The question that arises is how does all this impact security? The answer is that majority of the traditional monitoring and protecting methods require knowledge of the complete network topology while others rely on access to physical devices such as hardware-assisted SSL. In all of these cases, the dynamic nature of the cloud makes the traditional approaches impractical, as they would require constant configuration changes. Some approaches that require installation of hardware parts will be impossible to implement unless database services are implemented on a private cloud. Another security risk that is associated with physical security is the removal/deletion of data from storage devices known as data sanitization. Sanitization involves the deleting of data from storage media by overwriting, degaussing (de-magnetizing), or the destruction of the media

itself, to prevent unauthorized disclosure of information.

As organizations deploy applications and the databases that support them on virtual servers and in the cloud, a new complication frequently arises. In the past, an application was typically provisioned on one or more servers, and the databases housing the application were installed on separate networked servers. One of the benefits of virtualization (in a private data centre or in the cloud) is the ability to share resources, resulting in environments where both the application and the databases are migrating to virtual machines, in many cases running on the same physical servers. Clearly, the only solution is to bring the security inspection closer to its target. One solution is dubbed the virtual appliance. In this scenario, a virtual machine that runs the software formerly run by a dedicated appliance is installed on virtual servers and the servers are re-architected to send traffic through the virtual machine. This approach has two severe drawbacks performance and architecture complications.

III. III. EXISTING OPTIMAL SOLUTION

Database monitoring or auditing is basically the ability to constantly (and securely) record and report on all the events occurring within a database system. Audited databases generate reports on how, when and by whom different objects are accessed or altered. A strong database auditing and monitoring tool, that should provide full visibility into database activities regardless of its location, is extremely important for cloud based database services.

To meet the challenges of protecting traditionally on premise databases, IT security professionals initially adopted network based IDS and IPS – an appliance that would be placed somewhere in the network and would inspect the traffic for protocol violations, malicious code, viruses, etc. Although enterprises initially ignored internal risks and threats but they soon realized that internal threats can also be very damaging and monitoring must therefore cover local and intra-database attacks as well. [6] The adoption of local agents thus started to begin together with network based appliances making many of today's solutions. In this solution the host agents send local traffic back to the network appliance for analysis, where each transaction is measured against a pre-set policy. Although this hybrid approach is not ideal (ineffective for local breaches against security policies), but many enterprises still adopted it as a security solution.

Another limitation of network based monitoring solutions in the cloud environment is due the virtualization technology. In the past, applications using a database were usually deployed on separate physical servers while the database itself that hosted the applications was installed on separate networked servers. [6] However, through the use of virtualization technology, physical resources are shared in the cloud that sometimes results in environments where both the application and the databases reside on the same physical servers. For example, in the example below, note that communication between the application and the database occurs entirely within the same physical server. Network monitoring appliances will not be able to detect these transactions as no network traffic will be generated during the communication between the virtual machines.

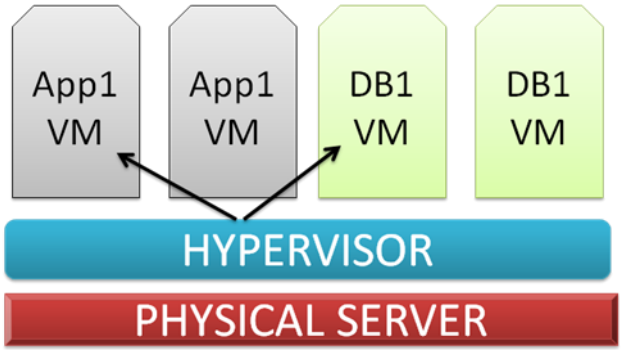

Fig 1:VM-VM communication

In the case of cloud based databases, the "network sniffing" model fails to address several technical challenges as the devices (except for on premise private cloud solutions) are outside the enterprise perimeter. Moreover, for scalability and redundancy purposes, databases residing in the cloud may dynamically appear in new locations over the course of time. This dynamic nature of cloud makes traditional methods impractical and requires that new approaches designed for distributed environments should be considered.

The architecture complications arise because enterprises would need to design in such a way in which all traffic to databases should pass through the virtual appliance first. Taking the dynamic nature of cloud into account, this approach will not be practical for cloud environments where hosts come and go, and adding virtual appliances to the mix would be extremely impractical. A solution with sensor based host agent, that would run alongside the database instances, is considered to be feasible for such environments.

In order for the solution to be effective, the local sensor or agent needs to be capable of reacting swiftly to alerts, implementing the required protections in case of a policy breach and alerting locally. Based on a set of policies and rules that is acquired from central management server, the sensors/agents would audit, send alerts, or suspend sessions that violate preset conditions. For secure and efficient transaction of policies and alerts, traffic between the sensors and remote management console should be encrypted and compressed.

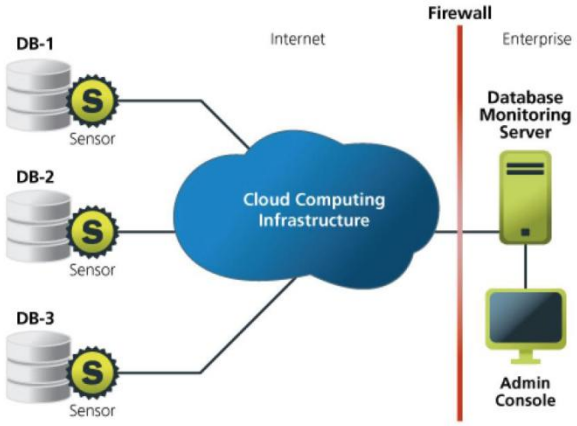

Fig 2: Distributed database monitoring by MacAfee

These agents/sensors have to be designed in such a way as to not suffer from the same weaknesses like intrusive implementation procedures and performance issues that old host-based solutions suffered. The sensors needs to be lightweight software that should act as an add-on and that could be easily added to a virtual machine in parallel with database instances and it must not be based on kernel-level implementations which require machine restarts.

Security software company, McAfee, provides one such solution by providing a software-based sensor. The sensor can be installed on the same virtual machine together with other database instances. The sensor functions by monitoring the database transactions occurring in the memory and thus protecting the system against all types of internal and external attacks. In order to work properly, the only information required by a newly installed sensor is the logical location of the central management server which would enable the host agents to monitor database activities and prevent attacks on the system.

IV. THE IEEMR MODEL

1. Inventory all Databases: The first step to effective database security in the cloud is to inventory all databases. DbProtect's Database Discovery feature generates a complete inventory of all databases deployed cloud-wide. It identifies all production, test, and temporary databases, and more importantly, any unauthorized databases.

2. Eliminate Vulnerabilities: Default and weak passwords, database misconfigurations, and missing security patches provide avenues of attack through standard database security to sensitive data. DbProtect's Vulnerability Management provides unparalleled database vulnerability assessment, allowing organizations to identify and eliminate vulnerabilities and fix misconfigurations that put their sensitive data at risk. Vulnerability Management is driven by a powerful policy development engine that begins with proven database security templates. DbProtect's policy development is fed by the Knowledgebase, the most comprehensive and up-to-date vulnerability and threat knowledge base in the industry. Each check in the Knowledgebase provides clear and detailed remediation instructions to insure that the vulnerabilities exposing sensitive data are fixed in a timely manner. DbProtect's Risk Analysis maps vulnerabilities to risk level and business impact. This helps organizations and Cloud providers to prioritize their remediation plans and ensure the most serious threats to sensitive data are addressed quickly.

3. Enforce least privileges: Over time, users accumulate more privileges than they need to do the job. This can lead to Segregation of Duties (SoD) violations that enable insiders to make fraudulent changes or steal sensitive data. DbProtect's Rights Management provides a detailed view of an organization's data ownership, access controls, and rights to sensitive information. Rights Management enables the organization to enforce the Principle of Least Privileges – grant only the privileges that users need to do their jobs. It allows organizations to restrict database access to a business need to know basis and militate against shared accounts. Rights Management also provides an audit trail on how privileges were granted, to help prevent against future privilege escalation. Backed by the Knowledgebase, DbProtect's DAM reduces risk and offers best-in-class data protection and reporting.

4. Monitor for deviations: Organizations and Cloud providers should track and monitor access to sensitive data and to regularly test database security processes. DbProtect's Database Activity Monitoring (DAM) helps secure sensitive data in the cloud by:

- Validating remediated vulnerabilities.
- Monitoring unremediated vulnerabilities to ensure they are not being exploited.
- Monitoring privileged user activity to ensure they are not engaged in any unauthorized behavior.
- Monitoring for any new avenues of attack.

DBProtect unique approach to Database Activity Monitoring (DAM) is precision monitoring. Precision monitoring employs DbProtect's powerful policy development engine to streamline monitoring operations to focus on any suspicious activity threatening sensitive data. DbProtect's precision DAM solution can be customized to a fine level of granularity – a specific activity, performed by a specific user, accessing specific data, in a specific database.

5. Respond to suspicious behavior: Active Response provides an additional layer of protection around sensitive data in the cloud. Active Response can be configured to take action when an unauthorized and suspicious database activity is detected. Active Response can be customized to a fine level of granularity – a specific activity, performed by a specific user, accessing specific data, in a specific database. Example: A user with excessive privileges attempts unauthorized access to sensitive data. Active Response can:

- Send an alert to IT Security to prompt further investigation.
- Terminate the user session to immediately stop the unauthorized access.
- Lockout the user's account to prevent further attempts to access the cardholder data.

Incorporating a comprehensive and disciplined program of database security process control and managing these five basic steps will help organizations and cloud providers to partner.

V. DEPLOYMENT

One of the best ways to ensure confidentiality of sensitive data in the cloud environment is to use encryption for data in transit as well as data at rest. Encryption support for data in transit is offered by nearly all cloud database providers (using TLS/SSL for transfer of data) but very few offer encryption options for data at rest.

Cloud service providers' main business idea is based on efficient resource utilization by a group of consumers. That is to say, the more customers utilize the same physical resources the more profit the service providers gain. This business model plays an important role for cloud service providers as to whether offer encryptions services or not. Encryption, being a processor intensive process, lowers the total amount of customers per resource and increases overall costs. Therefore, most cloud providers offer only partial encryption on a few database fields, such as passwords and account numbers. Although some providers do offer full database encryption options, but that increases the cost so much that hosting databases in the cloud becomes more expensive than having internal hosting. Alternatives to full database encryption are offered by some providers that have

less impact on the system's performance but it uses an ineffective technique that can be easily bypassed.

VI. CONCLUSION

Data sensitivity is a major concern for every user since it can be confidential or private data. The cloud service provider providing Database-as-a-Service has the sole responsibility of protecting the data but ultimately the user has the prime responsibility of protecting its data. The present model has many benefits of providing a firewall between the sensitive data and the public users in the cloud. The proposed IEEMR model improvises the security protocol of the data in the cloud. The present challenges faced are classified in a broader dimension and addressing for a solution to all of them is a difficult scenario.

Public cloud provides access to all the users that a re using the services of the cloud. Integrating and tightening the security of the data in the database provides many challenges such as authentication, communication among virtual machines where there is a central database. Another important approach followed by the IEEMR model is to provide selective access rights for all the users in the cloud. Data sensitivity

REFERENCES

[1] Irfan Gul, M. Hussain. 2011. Distributed Cloud Intrusion Detection Model. International Journal of Advanced Science and Technology.
[2] Ms. Parag K. Shelke, Ms. Sneha Sontakke, Dr. A. D. Gawande. 2012. Intrusion Detection System for Cloud Computing. International Journal of Scientific & Technology Research Volume 1.
[3] Piyush Saxena, Satyajit Padhy, Praveen Kumar 2013.USE OF STORAGE AS SERVICE FOR ONLINE OPERATING SYSTEM IN CLOUD COMPUTING, International Conference on Telecommunications and Networks (TEL-NET 2013)
[4] Security Guidance for Critical Areas of Focus in Cloud Computing, April 2009. DOI
[5] Miranda Mowbray, Siani Pearson. 2009."A Client-Based Privacy Manager for Cloud Computing." COMSWARE '09: Proceedings of the Fourth International ICST Conference on Communication System software and middleware
[6] Jinpeng Wei, Xiaolan Zhang, Glenn Ammons, Vasanth Bala, Peng Ning. 2009."Managing security of virtual machine images in a cloud environment." CCSW '09: Proceedings of the 2009 ACM workshop on Cloud computing security.
[7] B.Meena, Krishnaveer Abhishek Challa. 2012."Cloud Computing Security Issues with Possible Solutions." in IJCST
[8] Steve Hanna. A security analysis of Cloud Computing. Cloud Computing Journal. DOI
[9] Vahid Ashktorab, Seyed Reza Taghizadeh, 2012. "Security Threats and Countermeasures in Cloud Computing" in IJAIEM.
[10] Frank Doelitzscher, Christoph Reich, Martin Knahl, Alexander Passfal and Nathan Clarke. 2012. An agent based business aware incident detection system for cloud environments. Journal of Cloud Computing: Advances, Systems and Applications
[11] Peter Mell, Timothy Grance. 2011. The NIST Definition of Cloud Computing (Draft). NIST
[12] Ahmed Patel, Mona Taghavi, Kaveh Bakhtiyari, Joaquim Celestino Junior. 2013. An intrusion detection and prevention system in cloud computing: A systematic review. Journal of Network and Computer Applications.
[13] Hassen Mohammed Alsafi, Wafaa Mustafa Abduallah and Al-Sakib khan Pathan. 2012. IPS: An Integrated Intrusion Handling Model for Cloud Computing Environment, International Journal of Computing and Information Technology (IJCIT)
[14] Deris Stiawan, Abdul Hanan Abdullah, Mohd. Yazid Idris. 2011. Characterizing Network Intrusion Prevention System. International Journal of Computer Applications.
[15] Dinesh Sequeira. 2002. Intrusion Prevention Systems-Security Silver Bullet. SANS Institute.

Tinesh Kumar Goyal Pursuing Master of Technology in Computer Science and Engineering from Amity School of Engineering and Technology, Amity University Uttar Pradesh, Noida, India, Area of Interest: Cloud Computing, Computer Networks and Software Engineering.

Amitesh Kumar Gupta Pursuing Masters of Technology in Computer Science and Engineering from Amity School of Engineering and Technology, Amity University Uttar Pradesh, Noida, India, Area of Interest: Computer Networks, Network Security, Cloud Computing, Data Mining and Warehousing.

Asish Srivastava Pursuing Master of Technology in Computer Science and Engineering from Amity School of Engineering and Technology, Amity University Uttar Pradesh, Noida, India. Area of Interest: Cloud Computing, Virtualization, Computer Networks and Network Security.

Piyush Saxena Pursuing Master of Technology in Computer Science and Engineering from Amity School of Engineering and Technology, Amity University Uttar Pradesh, Noida, India, Area of Interest: Cloud Computing, Data Mining and Warehousing and Soft Computing.