

Security and Privacy Solutions associated with NoSQL Data Stores

Gerasimos Vonitsanos*, Elias Dritsas*, Andreas Kanavos*, Phivos Mylonas[†], Spyros Sioutas*

*Computer Engineering and Informatics Department

University of Patras, Patras, Greece

{mvonitsanos, kanavos, sioutas}@ceid.upatras.gr, eldritsas@gmail.com

[†]Department of Informatics

Ionian University, Corfu, Greece

fmylonas@ionio.gr

Abstract—Technologies such as cloud computing and big data management, have lately made significant progress creating an urgent need for specific databases that can safely store extensive data along with high availability. Specifically, an evergrowing amount of companies have put to use a multitude of non-relational databases, typically known as NoSQL databases. Not utilizing predefined guidelines, these databases largely offer a well structured mechanism for the retrieval and storage of numerous data. This fact renders them superior to RDBMS platforms. When on the other hand, big data and parallel processing is encountered in particular, and therefore the use of relational modeling is rendered obsolete, NoSQL databases largely exceed in superiority. Sensitive data is stored daily in NoSQL Databases, making the privacy problem more serious while raising essential security issues. In our paper, security and privacy issues when dealing with NoSQL databases are introduced, and in following, security mechanisms and privacy solutions are thoroughly examined.

Index Terms—NoSQL Databases, NoSQL Security, Secure Databases, Policy Update, Access Control

I. INTRODUCTION

The advances in cloud computing technology and distributed web applications, along with the ever-increasing enormous volume of data for storage and further processing, have necessitated the adoption of non-relational databases, commonly called NoSQL or "Not only SQL" [22]. It is widely known that the traditional SQL database is not able to cope with Big Data [28] as NoSQL systems, nowadays, are experimenting with an escalated popularity [20]. In recent years, many NoSQL databases have made their appearance; for example, Cassandra and MongoDB are two popular ones, to name a few. Some useful features of NoSQL databases are the high availability, scalability, better performance, plus the advantage to store and process large-scale semi-structured and/or unstructured data faster than traditional RDBMS [20], [28].

However, due to the ever increasing use of NoSQL databases, a significant amount of sensitive data is exposed to a number of security vulnerabilities, threats, and risks. Absence of encryption support and a not so strong authentication status

between servers and clients are some of the leading security issues in NoSQL Databases. Also, we should bear in mind that simple authorization is provided without support for role-based access control (RBAC); therefore is minimum protection for injections and denial of service attacks [14].

Brewer introduced the CAP (Consistency, Availability, Partition Tolerance leading to a conjecture [7] about the trade-offs having to do with the progress of developing a distributed database system. Brewer's conjecture, therefore, in its more formal version, is officially published as the CAP theorem in [13]. Specifically, the CAP theorem indicates that no shared data system could simultaneously provide above two out of the three characteristics, including consistency, availability, as well as partition tolerance.

Regarding organizations, Amazon developed the Dynamo technology [12], whereas Google produced the distributed storage system Bigtable [9]. These particular technologies have inspired many NoSQL applications installed in companies, like Facebook or Twitter. Modern companies encounter non-relational data and thus require superior databases compared to traditional ones that face scalability and availability problems due to their data size. There are already several authorization models in relational databases where views are usually utilized. In this way, SQL queries are used to display a certain state of a specified part of the database [4]. Some NoSQL Databases managed by Big Data use new authorization models, specifically designed for structure, speed, and a huge amount of data. These models include key-value, wide column, and document-oriented authorization. In addition, the storage and retrieval of these records are achieved through a unique key for each record while providing a swift search [10].

In this work, we present security and privacy issues within NoSQL databases. Additionally, we examine the possibility to propose the most efficient security mechanisms and privacy solutions. More to the point, protecting data and handling control to access are regarded as vital security issues in NoSQL. Several threats concerning NoSQL database security are also considered, ranging from distributed environment, authentication, fine-grained authorization to protection of data at rest and in motion.

The remainder of this paper is categorized in the follow-

ing mode. Section II analyzes a survey of existing related works concerning mechanisms to overcome security issues. Moreover, Section III overviews a comparative study between Relational and NoSQL Databases, while in Section IV, our security and privacy-preserving mechanisms are proposed. Finally, in Section V, the summary of the proposed paper is presented.

II. RELATED WORK

Many early papers that described the relation between Relational and NoSQL databases were studied meticulously leading to an overview of the NoSQL database, along with the related characteristics and types. These surveys were positively predisposed towards NoSQL and the mode it employs to decline [23], [29]. However, in [5], there was clear interest about structured and non-structured databases. The mode in which how using NoSQL databases, like Cassandra, would meliorate the system performance was also considered and fully described. It can also scale the network without changing any hardware or needing to alternate the server infrastructure. This results in improving the network scalability, with low-cost commodity hardware.

In [16], a survey paper regarding relational databases is introduced along with NoSQL features and shortcomings. These shortcomings and issues of the NoSQL databases have been mentioned in [18]; complexity, consistency and limited eco structures are considered serious concerns. Another in depth analysis having to do with security and privacy issues in big data and NoSQL is considered in [27]. Specifically, because of the towering size, rapidity and diversity of big data, safety and privacy matters are diverse in such streaming data infrastructures with several data formats; therefore, conventional security models present complexities in treating such large scale data. In [21], it is stated that the request for a relational database is here to stay for some time. This means that it will exclusively serve in line with applications that will support business operations. However, NoSQL databases will serve the large, public and content centric applications. Another similar work is presented in [22], where an extensive analysis of security issues with NoSQL Database, like Cassandra and MongoDB, is considered.

Furthermore, authors in [2] introduce the basic variations amongst conventional relational databases and NoSQL ones, while focus on the security mechanism that must be implemented at the middleware by the developers so as to surpass safety matters of NoSQL databases. In [8], a variety of SQL and NoSQL data stores formulated to scale plain OLTP-style application loads over a number of servers, is examined. Concretely, the new systems are contrasted against their data model, consistency and storage mechanisms, durability guarantees, availability, query support, as well as other aspects. An assessment criterion consisting of multiple safety characteristics for the analysis of sharded NoSQL databases is proposed in [31]. This analysis helps various organizations select appropriate and reliable databases according to their preferences and security requirements.

Several solutions have also been proposed to improve privacy-preserving in NoSQL databases. More specifically, in Arx, a proxy is employed to rewrite NoSQL queries at the trusted premises. A back-end component, deployed at the untrusted premises, is utilized to handle computation over encrypted data [25]. In terms of BigSecret system, standard encryption is used for protection of the stored data, while the indexes are encoded with the aid of special methods to permit comparisons (pseudo-random functions) and range queries (order-preserving partitioning) [24]. Authors in [30] employ algorithms of searchable encryption to create a privacy-preserving key-value store on top of the Redis database. In this approach, the values are safeguarded using symmetric encryption, while the keys are protected with pseudo-random techniques. In another solution, SafeRegions combine secret sharing and multiparty computation to perform secure NoSQL queries on three independent and untrusted HBase clusters providing thus safe computation over the stored values and security guarantees similar to standard encryption at the same time [26].

III. COMPARISON OF RELATIONAL AND NOSQL DATABASES

During the last decades, relational databases, sub-divided into groups known as tables, have been used to store structural data. The units of data in each table are known as columns, and each unit of the group is known as a row. Also, the columns in a relational database have relationships amongst them. This phenomenon is bound to alteration over the last years because of the escalation of big-sized web applications, which output a large amount of data that traditional relational databases cannot handle any more [11]. NoSQL databases are sometimes referred as “Not only SQL” to give some emphasis on the fact that they may support query languages that are SQL-like. Nowadays, it is stated that NoSQL databases have more to offer than just present solutions to scale problems, while also provide many important advantages [11] like the following:

- The data representation is schema-less, and there is no need to define a certain structure from scratch since new fields can be added at run-time.
- The speed, even with a small amount of data, can be processed in milliseconds instead of hundreds of milliseconds.
- The elasticity of the applications because of the scalability features NoSQL databases provide.
- Reduction in development time, as developers do not undertake the laboring task of facing complex SQL queries and difficult joints so as to collate the data from varied tables into a new view.

Some of the differences between relational and NoSQL databases are listed in the following paragraphs.

A. Reliability of Transactions

The ACID (atomicity, consistency, isolation, durability) model is fully supported by the design of relational databases,

providing high reliability in transactions unlike the NoSQL databases.

B. Scalability Issues and Cloud Support

The primary purpose of cloud technology is to offer end-users an array of services. NoSQL databases are fully compatible with cloud environment requirements as they can analyze not only raw structured data, but also semi-structured or unstructured data from various sources, since they are not compliant with the ACID model. On the other hand, the relational databases do not provide data search on a full basis and their characteristics are now designed for cloud use.

The need for scalability may be one of the most significant problems of relational databases as they rely on vertical scalability to upgrade the performance. More specifically, this upgrade method requires the purchase of expensive equipment such as RAM, processors, SSD hard drives, etc. and in some cases, this is not easily achieved due to each system constraints. Also, the possibility of horizontal scaling is not supported by the addition of extra nodes and therefore, cannot sustain demanding online applications with many users and distributed data. However, NoSQL databases support only horizontal scaling since they do not deal with relational data.

C. Complexity and Big Data Management

The complexity of NoSQL databases is less than that of relational databases as it is not necessary to create tables to record data, but instead the modeling by considering a query method that can be used. The development of a database structure on a relational database is always considered a complicated task compared to the abstract model of a NoSQL database, where data can be stored regardless of whether they are structured, unstructured, or semi-structured.

NoSQL databases have a valuable role in Big Data management since they are well-suited for storing or retrieving data in high speed across distributed nodes, thus taking advantage of multi-core GPU architectures. In relational databases, where accuracy is more important than speed, the data should be stored in tables' rows and columns, while the scalability is always considered a big issue. In the case of supporting conventional applications with small datasets, they are the most reasonable choice, but slitting the data across different servers increases the arduousness requiring complex SQL queries for joining the data again.

D. Data Model

Sets in mathematics are the driving force for a relational database; all the data are represented as mathematical n -ary relations, where an n -ary relation is a subset of the Cartesian product of N domains. The data are represented as tuples inside the database and are further grouped into relations. The relation (represented by table) contains a set of tuples (represented by rows); where the column in the relation table utilizes the sequence of attributes, the type of an attribute is identified by the domain, which is the set of values that have a common meaning. This data model is precise to the point and

well structured, while the columns and the rows are described by a well-defined schema.

NoSQL databases can employ many modeling methods such as graph, key-value stores and document data model. In terms of classification procedure, NoSQL is named after their data model. Still in some cases, NoSQL database system can be identified using two or more of the data models that represent their data. NoSQL data model does not utilize the table as storage structure of the data and this is considered the main feature that distinguishes the NoSQL from relational databases. Furthermore, it is schema-less and as a result, can handle the unstructured data like word, pdf, images, as well as video files, in a very efficient method.

E. Data Warehouse and Crash Recovery

Regarding data warehousing, relational databases gather data from an array of sources and the oversize of stored data results in big data problems. To name a few, some problems are the performance degradation when utilizing an OLAP (Online Analytical Processing), statistical process or Data Mining. On the other hand, NoSQL databases are not designed when considering data warehouse applications, because designers are focused on issues, like scalability, availability and high performance.

Crash recovery is implemented in relational databases via the recovery manager, which holds the responsibility to ensure durability and transaction atomicity by using log files and ARIES algorithm. The crash recovery in NoSQL databases depends on replication to recover from the crash.

F. Privacy and Security

Most relational databases do not provide a single feature regarding embedding security in the database itself. As a result, this requires developers to directly impose security systems in the middleware. Classic cryptography mechanisms and encryption protocols, like asymmetric key encryption schemes, digital signature schemes, zero-knowledge Proof of Knowledge, as well as commitment schemes, which are based on SRSA (Strong RSA), bilinear maps [3], discrete logarithm, homomorphic encryption, fully or not [1], have been widely considered for securing communication and ensuring data confidentiality in relational databases.

Nonetheless, a very vital shortcoming of NoSQL databases is considered the fact that data files are not by default encrypted, but such a process takes place in the application layer before sending data to the database server. Although there are solutions that provide encryption services, these lack horizontal scaling and transparency required in the NoSQL environment.

Furthermore, only a few NoSQL databases provide encryption mechanisms to protect user-related sensitive data. By default in NoSQL databases, the inter-node communication is not encrypted and does not support SSL (Secure Sockets Layer) client-node communication (as in relational databases), breaking the network security [28]. Also, there is no integration of authentication or authorization mechanisms.

The distributed environments increase attack surface across several distributed nodes and enforcing integrity constraints is much complex in NoSQL databases. In general, only a few categories of NoSQL databases provide mechanisms to employ encryption techniques protecting data at rest.

IV. PROPOSED SECURITY AND PRIVACY SOLUTIONS

Below are presented our proposed security and privacy solutions for NoSQL data stores.

A. Pseudonyms-based Communication Network

In the context of this system, users can have access to multiple services by inserting their credentials only once, that is when they are initially connected to the system. Such a system is called anonymous because users can be known only through their pseudonyms, and the transactions demonstrated by the same user cannot be linked, as their identity is disclosed. For this reason, it is considered the best means in terms of user protection. Furthermore, it is based on two vital protocols, the RSA (RivestShamirAdleman) and the Diffie Hellman. Its structure and operations extend Brand's credential system [6], whereas it consists of four parties: the users U , a central Identity Provider denoted as IP , the Service Providers SPs , and the organization for issuing and validating credentials. Users are entities that receive credentials and are known to Service Providers only through their pseudonyms.

The central Identity Provider creates its own public and secret key, denoted as (P, S) respectively, and uses its secret key to digitally sign its sensitive data. Each credential is encoded with $m + 1$ attributes, denoted as y_1, y_2, \dots, y_m, t (where t is the credential issuing time). The IP decides on the \mathbb{G}_q , a finite cyclic group of prime order q , to which the random generators $g_1, g_2, \dots, g_m, g_{m+1}, h_0$, involved in keys generation, belong to. Specifically,

$$S = (y_1, y_2, \dots, y_m, t, s) \quad (1)$$

$$P = g_1^{y_1} g_2^{y_2} \dots g_m^{y_m} g_{m+1}^t h_0^s, \quad (2)$$

where $s \in \mathbb{Z}_q$ is secret.

Under the Discrete Logarithm assumption in \mathbb{G}_q , these keys are unique. The IP is responsible for the distribution of the digital pseudonyms p_1, p_2, \dots, p_m to any user. An organization can issue a credential to a pseudonym, and the corresponding user can prove its ownership to another organization (who knows it by a different pseudonym), by just revealing the ownership of the credential.

Additionally, the Credential Authority CA prevents the sharing of credentials or pseudonyms and guarantees that users who enter the system have a public and secret key that makes them unique to the system. Another entity in the system is the Verifier V , whose role is to certify the validity of the user credentials and to communicate with either the Issuing Authority or the Credential Authority and inform that the user is not the owner of the credential that is presenting. Users, in terms of a digital credential, transmit the public key and

the CA s digital signature derived from a Proof of Knowledge, through which they prove that they know the secret key and the attributes in the digital credential that satisfies the particular attribute property they are revealing.

Each pseudonym and credential belong to well-defined users. More in detail, collaboration amongst different users, to show part of their credentials in a Service Provider as well as to obtain a credential for a user that they could not obtain (coherent credentials), is not possible. As organizations are autonomous and separable entities, they can select their public and secret key independently of the other entities, so as to ensure the security of these keys and facilitate the keys' management system.

The pseudonyms' system can protect user privacy and provide security, as in such a system, an organization can not find out anything about a user other than the ownership of a set of credentials. Specifically, two pseudonyms that belong to the same user cannot be linked (unlinkability) and identified as in the Brands system, except for specific conditions. In order to be efficient, any communication in the system involves as few entities as possible along with the minimum amount of information. If a user holds a credential, this can be shown multiple times without the need to reissue (and consequently resign) it.

When a user has access to a service, they are validated by proving that they know the secret key of their pseudonym, without revealing it, thus preventing pseudonym repetition. Also, for each pseudonym that a Service Provider associates with a user, it requires the user to unveil a different encoded random number of their pseudonym each time and thus, ensures the unconditional unlinkability of their pseudonyms. Although the Identity Provider blindly encodes the random numbers in all of a user's pseudonyms, that are uniquely related with them, if a user makes abuse of the service, the SP can blacklist and reveal numbers. In following, it is able to globally revoke their pseudonyms and abolish their access to any of the services they have previously had.

Finally, users, under Discrete Logarithm, can conclusively prove that their encoded numbers do not belong to the SP 's blacklist, while using this one as input on a zero-knowledge proof and without revealing any information about their identity. Hence, this technique does not impact users' privacy and does not strengthen the SP and IP .

B. Monitoring, Filtering and Blocking

As mentioned above, available applications designed to monitor NoSQL databases cannot detect and then disable malicious jobs and queries. The Kerberos central authentication system can be easily bypassed via advanced scripts, and in general, the level of monitoring is limited to data processing mainly in the API [17].

In a cloud environment, no information regarding the communication of nodes in the cluster or user connection details or data altering information (even editing or deleting), is recorded. In general, since there are no log files, a challenging

problem is to identify incidents of data breach or malicious data loss in the cluster [15].

Real-time security mechanisms exist in big data technologies, resulting in high-speed data analysis. Therefore, the detection of anomalies is real-time implemented and the recording of security analytics can be frequently updated [14]. Some monitoring tools are available but are limited to controlling user requests at the API level. In general, neither the characteristics of a malicious query in big data technologies are defined, nor complete monitoring tools to disable these malicious queries, exist. One potential technique could be an initial authentication via Kerberos and in following, a second level authentication for accessing MapReduce [19].

V. CONCLUSIONS

In this paper, we have addressed major security concerns regarding NoSQL databases. Data protection and access control can be considered some of the key issues of security in NoSQL technology. Reasons for security threats in various NoSQL databases have also been thoroughly discussed in the current work, like privacy of user data, distributed environment, authentication, fine-grained authorization and access control, securing integrity as well as protection of data at rest and in motion.

In NoSQL databases, Kerberos is used to authenticate the clients and data nodes. Specifically, to ensure fine-grained authorization, data are grouped according to their security level. On the other hand, Cassandra uses TDE technique to protect data at rest, whereas administrators must implement controls for ensuring that application and users have only access to the data they need in order to maintain a secure MongoDB deployment. Various techniques for mitigating the attacks on NoSQL databases have also been discussed along with proposed security and privacy solutions regarding NoSQL databases.

REFERENCES

- [1] M. Ahmadian, F. Plochan, Z. Roessler, and D. C. Marinescu. Securenosql: An approach for secure search of encrypted nosql databases in the public cloud. *International Journal of Information Management*, 37(2):63–74, 2017.
- [2] J. Ahmed and R. Gulmeher. Nosql databases: New trend of databases, emerging reasons, classification and security issues. *International Journal of Engineering Sciences and Research Technology (IJESRT)*, 2015.
- [3] E. Bangerter, J. Camenisch, and A. Lysyanskaya. A cryptographic framework for the controlled release of certified data. In *12th International Workshop on Security Protocols*, volume 3957, pages 20–42, 2004.
- [4] E. Bertino and R. S. Sandhu. Database security-concepts, approaches, and challenges. *IEEE Transactions on Dependable and Secure Computing*, 2(1):2–19, 2005.
- [5] A. Bhatewara and K. Waghmare. Improving network scalability using nosql database. *International Journal of Advanced Computer Research*, 2(4):488, 2012.
- [6] S. Brands, L. Demuyneck, and B. D. Decker. A practical system for globally revoking the unlinkable pseudonyms of unknown users. In *12th Australasian Conference on Information Security and Privacy (ACISP)*, volume 4586, pages 400–415, 2007.
- [7] E. A. Brewer. Towards robust distributed systems. In *19th Annual ACM Symposium on Principles of Distributed Computing*, page 7, 2000.
- [8] R. Cattell. Scalable SQL and nosql data stores. *SIGMOD Record*, 39(4):12–27, 2010.
- [9] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Transactions on Computer Systems (TOCS)*, 26(2):4:1–4:26, 2008.
- [10] P. Colombo and E. Ferrari. Fine-grained access control within nosql document-oriented datastores. *Data Science and Engineering*, 1(3):127–138, 2016.
- [11] A. Davoudian, L. Chen, and M. Liu. A survey on nosql stores. *ACM Computing Surveys*, 51(2):40:1–40:43, 2018.
- [12] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon’s highly available key-value store. In *21st ACM Symposium on Operating Systems Principles (SOSP)*, pages 205–220, 2007.
- [13] S. Gilbert and N. A. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, 2002.
- [14] N. Gupta and R. Agrawal. Chapter four - nosql security. *Advances in Computers*, 109:101–132, 2018.
- [15] V. N. Inukollu, S. Arsi, and S. R. Ravuri. Security issues associated with big data in cloud computing. *International Journal of Network Security and Its Applications (IJNSA)*, 6(3):45, 2014.
- [16] N. Jatana, S. Puri, M. Ahuja, I. Kathuria, and D. Gosain. A survey and comparison of relational and non-relational database. *International Journal of Engineering Research and Technology (IJERT)*, 1(6):1–5, 2012.
- [17] P. Kadebu and I. Mapanga. A security requirements perspective towards a secured nosql database environment. In *International Conference of Advance Research and Innovation (ICARI)*, 2014.
- [18] N. Leavitt. Will nosql databases live up to their promise? *IEEE Computer*, 43(2):12–14, 2010.
- [19] S. LLC. Securing big data: Security recommendations for hadoop and nosql environments. 2012.
- [20] M. Mohamed, O. G. Altrafi, and O. Ismail. Relational vs. nosql databases: A survey. *International Journal of Computer and Information Technology*, 3(3):598–601, 2014.
- [21] C. Nance, T. Losser, R. Iype, and G. Harmon. Nosql vs rdbms - why there is room for both. In *Southern Association for Information Systems Conference*, 2013.
- [22] L. Okman, N. Gal-Oz, Y. Gonen, E. Gudes, and J. Abramov. Security issues in nosql databases. In *IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 541–547, 2011.
- [23] R. P. Padhy, M. R. Patra, and S. C. Satapathy. Rdbms to nosql: Reviewing some next-generation non-relational database’s. *International Journal of Advances in Engineering, Science and Technology (IJAEST)*, 11(1):15–30, 2011.
- [24] E. Pattuk, M. Kantarcioglu, V. Khadilkar, H. Ulusoy, and S. Mehrotra. Bigsecret: A secure data management framework for key-value stores. In *6th IEEE International Conference on Cloud Computing*, pages 147–154, 2013.
- [25] R. Poddar, T. Boelter, and R. A. Popa. Arx: A strongly encrypted database system. *IACR Cryptology ePrint Archive*, 2016:591, 2016.
- [26] R. Pontes, F. Maia, J. Paulo, and R. M. P. Vilaça. Saferegions: Performance evaluation of multi-party protocols on hbase. In *35th IEEE Symposium on Reliable Distributed Systems Workshops (SRDS)*, pages 31–36, 2016.
- [27] E. Sahafzadeh and M. A. Nematbakhsh. A survey on security issues in big data and nosql. *Advances in Computer Science: An International Journal*, 4(4):68–72, 2015.
- [28] H. Shahriar and H. M. Haddad. Security vulnerabilities of nosql and sql databases for mooc applications. *International Journal of Digital Society (IJDS)*, 8(1), 2017.
- [29] V. Sharma and M. Dave. Sql and nosql databases. *International Journal of Advanced Research in Computer Science and Software Engineering*, 2(8), 2012.
- [30] X. Yuan, X. Wang, C. Wang, C. Qian, and J. Lin. Building an encrypted, distributed, and searchable key-value store. In *11th ACM on Asia Conference on Computer and Communications Security (AsiaCCS)*, pages 547–558, 2016.
- [31] A. Zahid, R. Masood, and M. A. Shibli. Security of sharded nosql databases: A comparative analysis. In *Conference on Information Assurance and Cyber Security (CIACS)*, pages 1–8, 2014.