A Framework for Efficient Storage Security in RDBMS

Bala Iyer[2], Sharad Mehrotra[1], Einar Mykletun[1],
Gene Tsudik[1] and Yonghua Wu[1]

[1] University of California, Irvine, Irvine CA 92697, USA
{yonghuaw, mykletun, sharad, gts}@ics.uci.edu
[2] IBM Silicon Valley Lab
balaiyer@us.ibm.com

Abstract. With the widespread use of e-business coupled with the public's awareness of data privacy issues and recent database security related legislations, incorporating security features into modern database products has become an increasingly important topic. Several database vendors already offer integrated solutions that aim to provide data privacy within existing products. However, treating security and privacy issues as an afterthought often results in inefficient implementations. Some notable RDBMS storage models (such as the N-ary Storage Model) suffer from this problem. In this work, we analyze issues in storage security and discuss a number of trade-offs between security and efficiency. We then propose a secure storage model and a key management architecture which enable efficient cryptographic operations while maintaining a very high level of security. Finally we analyze the performance of our secure storage model implementation by running various experiments based on the TPC-H data set.

1 Introduction

Recently intensified concerns about security and privacy of data have prompted new legislation and fueled the developed of new industry standards. These include, notably, the Gramm-Leach-Bliley Act (also known as the Financial Modernization Act) [3] that protects personal financial information, and the Health Insurance Portability and Accountability Act (HIPAA) [4] that regulates the privacy of personal health care information.

Basically, this new legislation requires anyone storing sensitive data to do so in encrypted fashion. As a result, database vendors are working towards adding security- and privacy-preserving solutions in their product offerings. Two prominent examples are Oracle [2] and IBM DB2 [5]. Although this problem is very important, little can be found on this subject in the research literature, with the exception of [6], [7] and [8].

Designing an effective security solution requires, among other things, understanding the points of vulnerability and the attack models. Important issues that must be addressed include: (1) choice of encryption function(s) to use, (2) how

to perform key management, and (3) at what granularity to encrypt data. The main challenge is to introduce security functionality without incurring too much of overhead in terms of both performance and storage. The problem is further exacerbated since stored data may comprise both sensitive as well as non-sensitive components and access to the latter should not be degraded simply because the former must be protected.

In this paper, we argue that adding privacy as an afterthought results in sub-optimal performance. Efficient privacy measures require fundamental changes to the underlying storage subsystem implementation. We propose such a storage model and develop appropriate key management techniques which minimize the possibility of key and data compromise. More concretely, our main contribution is a novel DBMS storage model that facilitates efficient implementation of encryption techniques. Our approach is involves grouping sensitive data, in order to minimize the number of necessary encryption operations, thus minimizing cryptographic overhead.

Model: We envision a client-server scenario. The client has a combination of sensitive and non-sensitive data stored in a database at the server, with the sensitive data stored in encrypted form. Whether or not the two parties are co-located does not make a difference in terms of security. The server's added responsibility is to protect the client's sensitive data, i.e., to ensure its confidentiality and prevent unauthorized access. (Note that maintaining availability and integrity of stored data is an entirely different requirement.) This is accomplished through the combination of encryption, authentication and access control.

Trust in Server: The level of trust in the database server can range from fully trusted to fully untrusted, with several intermediate points. In a fully untrusted model, the server is not trusted with the client's cleartext data which it stores. (It may still be trusted with data integrity and availability.) Whereas, in a fully trusted model, the server essentially acts as a remote (outsourced) database storage for its clients.

Our focus is on environments where server is partially trusted. We consider one extreme of fully trusted server neither general nor particularly challenging. The other extreme case of fully untrusted server corresponds to the so-called "Database-as-a-Service" (DAS) model [9]. In a DAS model, a client does not even trust the server with cleartext queries; hence, it involves the server performing encrypted queries over encrypted data. The DAS model is interesting in its own right and presents a number of challenges. However, it also significantly complicates query processing at both client and server sides.

1.1 Potential Vulnerabilities

Our model has two major points of vulnerability with respect to client's data:

- **Client-Server Communication:** Assuming that client and server are not co-located, it is vital to secure their communication since client queries can involve sensitive inputs and server's replies carry confidential information.

– **Stored Data:** Typically, DBMS-s protect stored data through access control mechanisms. However, as mentioned above, this is insufficient, since server's secondary storage might not be constantly trusted and, at the very least, sensitive data should be stored in encrypted form.

All client-server communication can be secured through standard means, e.g., an SSL connection, which is the current *de facto* standard for securing Internet communication. Therefore, communication security poses no real challenge and we ignore it in the remainder of this paper. With regard to the stored data security, although access control has proven to be very useful in today's databases, its goals should not be confused with those of data confidentiality. Our model assumes potentially circumvented access control measures, e.g., bulk copying of server's secondary storage. Somewhat surprisingly, there is a dearth of prior work on the subject of incorporating cryptographic techniques into databases, especially, with the emphasis on efficiency. For this reason, our goal is to come up with a database storage model that allows for efficient implementation of encryption techniques and, at the same time, protects against certain attacks described in the next section.

1.2 Security and Attack Models

In our security model, the server's memory is trusted, which means that an adversary can not gain access to data currently in memory, e.g., by performing a memory dump. Thus, we focus on protecting secondary storage which, in this model, can be compromised. In particular, we need to ensure that an adversary who can access (physically or otherwise) server's secondary storage is unable to learn anything about the actual sensitive data.

Although it seems that, mechanically, data confidentiality is fairly easy to obtain in this model, it turns out not be a trivial task. This is chiefly because since incorporating encryption into existing databases (which are based on today's storage models) is difficult without significant degradation in the overall system performance.

Organization: The rest of the paper is organized as follows: section 2 overviews related work and discusses, in more detail, the problem we are trying to solve. Section 3 deals with certain aspects of database encryption, currently offered solutions and their limitations. In section 4, we outline our DBMS storage model which avoids the relatively high cost of the current solutions by re-organizing the page layout to enable efficient cryptographic operations. We also address encryption of indexes and other database related operations affected by our model. Section 5 consists of experiments with an implementation of our storage model using the TPC-H data set. We conclude the paper in section 6.

2 Background

Incorporating encryption into databases seems to be quite a recent development among the industry database providers [2] [5], and not a lot of research has been

devoted to this subject in terms of efficient implementation models. A survey inspecting the types of techniques used and offered by modern database providers can be found [10].

Some recent research has focused on providing database as a service (DAS) in an untrusted server model [7] [9]. Some of this work has dealt with analyzing how data can be stored securely at the server in such a way to allow a client to execute SQL queries directly on the encrypted tuples. As far as trusted server models, one approach that has been investigated involves the use of tamper resistant hardware (smart card technology) to perform the encryption operations at the server side [8].

2.1 The Problems

The implementations of encryption into today's DBMS's have often been incomplete as several factors have been neglected. Below we point out a few of these.

Lack of efficiency: Consistently, the added security measures introduce a large computational overhead to the running time of the general database operations. The lack of efficiency in terms of incorporating the encryption into existing databases is mainly due the underlying storage model used by the DBMS's. It seems difficult to find an efficient encryption scheme for current database products without modifying the way in which records are stored in blocks on disk. The large overhead encountered by the addition of encryption to databases has been demonstrated [10], where a comparison is performed among queries performed on several pairs of identical data sets, of which one contains encrypted information while the other does not.

Inflexibility: Depending on the level of encryption granularity, it might not be feasible to separate sensitive from non-sensitive fields when encrypting. As an example, if row level encryption is used and only one out of several attributes of the relation needs to be kept confidential, a considerable amount of computational overhead would exist due to the un-necessary encryption performed. We discuss different levels of encryption granularities in section 3.2, and will just mention here that the more granular level chosen, the more flexibility available in terms of selecting the attribute values to encrypt.

Meta data files: Many vendors seem content with being able to claim that they offer security along with their database products. Some of these provide an incomplete solution by only allowing for encryption of the actual records while disregarding meta-data and log files which can be used to reveal the sensitive fields.

Indexes are not encrypted: Several vendors do not permit for encryption of indexes, while others allow the users to build them based on the encrypted values, thereby losing a few of the most obvious characteristics of an index, namely range searches, as the encryption algorithms used will not be order preserving.

By not encrypting an index that has been constructed on a sensitive attribute (e.g., social security number), the encryption of records becomes meaningless. Encryption of the index is discussed in section 4.6.

3 Database Encryption

As is well-known, there are two main classes of encryption algorithms: *conventional* and *public-key*. Although both can provide data confidentiality, their goals and performance differ widely. Conventional, (also known as symmetric-key) encryption algorithms require the encryptor and decryptor to share the same key. Such algorithms can achieve high bulk encryption speeds, as high as 100s of Mbits/sec. However, they suffer from the problem of *secure key distribution*, i.e., the need to securely deliver the same key to all necessary entities.

Public-key cryptography solves the problem of key distribution by allowing an entity to create its own public/private key-pair. Anyone with the knowledge of an entity's public key can encrypt data for this entity, while only someone in possession of the corresponding private key can decrypt the respective data. While elegant and useful, public key cryptography typically suffers from slow encryption speeds (up to 3 orders of magnitude slower than conventional algorithms) as well as secure public key distribution and revocation issues.

To take advantage of their respective benefits and, at the same time, to avoid drawbacks, it it usual to bootstrap secure communication by having the parties use a public-key algorithm (e.g., RSA [11]) to agree upon a secret key, which is then used to secure all subsequent transmission via some efficient conventional encryption algorithm, such as AES [12].

Due to its clearly superior performance, we use symmetric-key algorithms for encrypting data stored at the server. It is also worth noting that our particular model does not warrant using public key encryption at all.

3.1 Encryption Modes and Their Side-Effects

A typical conventional encryption algorithm offers several modes of operation. They can be broadly classified as *block* or *stream* cipher modes.

Stream ciphers involve creating a key-stream based on a fixed key (and, optionally counter, previous ciphertext, or previous plaintext) and combining it with the plaintext in some way (e.g, by xor-ing them) to obtain ciphertext. Decryption involves reversing the process: combining the key-stream with the ciphertext to obtain the original plaintext. Along with the initial encryption key, one needs to keep stated information regarding the parameters used to create the key-stream, so that it can be re-created at a later time.

Block ciphers take as input a sequence of fixed-size plaintext blocks (e.g., AES uses 128-bit blocks) and output the corresponding ciphertext block sequence. It is usually necessary to pad the plaintext before encryption in order to have it align with the desired block size. This can cause certain overhead in terms of storage space, resulting in the some data *expansion*. A chained block cipher

(CBC) mode is a blend of block and stream modes; in it, a sequence of input plaintext blocks is encrypted such that each ciphertext block is dependent on all preceding ciphertext blocks and, conversely, influences all subsequent ciphertext blocks.

We use a block cipher in the CBC mode. Reasons for choosing block (over stream) ciphers include the added complexity of implementing stream ciphers, specifically, avoiding re-use of key-streams. This complexity stems from the dynamic nature of the stored data: the contents of data pages may be updated frequently, requiring the use of a new key-stream. In order to remedy this problem, a certain amount of state would be needed to help create appropriate distinct key-streams whenever stored data is modified.

3.2 Encryption Granularity

Encryption can be performed at various levels of granularity. In general, finer encryption granularity affords more flexibility in allowing the server to choose what data to encrypt. This is important since the stored data may include non-sensitive fields which ideally should not be encrypted (if for no other reason than to reduce overhead). The obvious encryption granularity choices are:

- **Attribute value:** smallest achievable granularity; each attribute value of a tuple is encrypted separately.
- **Record/row:** each row in a table is encrypted separately. Consequently, if we only need to retrieve certain tuples and we know where they are stored, we do not have to decrypt the entire table.
- **Attribute/column:** a more selective approach whereby only certain sensitive attributes (e.g., credit card numbers) are encrypted.
- **Page/block:** this approach is geared for automating the encryption process. Whenever a page/block of sensitive data is stored on disk, the entire block is encrypted. One such block might contain one or more tuples, depending on the number of tuples fitting into one page (a typical page is 16 Kbytes).

As mentioned above, we need to avoid encrypting non-sensitive data. If a record contains only a few sensitive fields, it would be wasteful to use row- or page-level encryption. However, if an entire table was to be encrypted, it would be advantageous to work at the page level. This is because encrypting fewer large pieces of data is always considerably more efficient than encrypting several smaller pieces. Indeed, this is supported by our experimental results in section 3.6.

3.3 Key Management

Key management is clearly an important aspect of the PPC model. We propose a key management scheme based on a two-level hierarchy consisting of a single master key and multiple sub-keys. Sub-keys are associated with individual tables or pages and are used to encrypt the data therein. The generation of all keys is the responsibility of the database server. The sub-keys are encrypted under the

master key. Certain precautions need to be taken in the event that the master key is believed to be compromised; in particular, re-keying strategies must be considered.

3.4 Re-keying and Re-encryption

There are two types of re-keying: periodic and emergency. The former is useful since it is generally considered good practice to periodically change data encryption keys, especially, for data stored over a long term. Folklore has it, that the benefit of periodic re-keying is to prevent potential key compromise. However, this is not the case in our setting, since an adversary can always copy the encrypted database from untrusted secondary storage and compromise keys at some (much) later point, via, e.g., a brute-force attack.

Emergency re-keying is done whenever key compromise is suspected or expected. For example, if a trusted employee (e.g., a DBA) who has access to encryption keys is about to be fired or re-assigned, the risk of this employee mis-using the keys must be considered. Consequently, to prevent potential compromise, all affected keys should be changed before (or at the time of) employee termination or re-assignment.

3.5 Key Storage

Clearly, where and how the master key is stored influences the overall security of the system. The master key needs to be in possession of the DBA, stored on a smart card or some other hardware device or token. Presumably, this device is plugged in to the database server during normal operation. However, in this case, it is possible for a DBA to abscond with the master key or somehow leak it. This should trigger emergency re-keying, whereby a new master key is created and all keys previously encrypted under the old master key are updated accordingly.

3.6 Encryption Costs

Advances in general processor and DSP design continuously yield faster encryption speeds. However, even though encryption data rates can be very high, there remains a constant start-up cost associated with each individual encryption operation. This cost dominates overall processing time when encrypting small amounts of data, such as individual records or attribute values.

Experiments: Recall our earlier claim that encrypting the same amount of data using few encryption operations with large data units is more efficient than many operations with small data units. Although this claim is quite intuitive, we still chose to run an experiment to support it. The experiment consisted of encrypting 10 Mbytes using both large and small unit sizes: blocks of 100-, 120-, and 16K-bytes. The two smaller data units represent average sizes for records in the TPC-H data set [13], while the last unit of 16-Kbytes was chosen as it is

Table 1. Many small vs few large data blocks encrypted. All times in msec include initialization and encryption cost.

Encryption Alg	**100 Bytes * 100,000**	**120 Bytes * 83,3333**	**16 KBytes * 625**
AES	365	334	194
DES	372	354	229
Blowfish	5280	4409	170

the default page size used in MySQL's InnoDB table type. We used MySQL to implement our proposed storage model, and describe the details in section 4

We then performed the following operations: 100,000 encryptions of the 100-byte unit, 83,333 encryptions of the 120-byte unit, and 625 encryptions of the 16-Kbyte unit. As the hardware platform we used a Linux box with a 2.8 Ghz Pentium IV and 1-Gbyte of memory. We used the OpenSSL library [14] to write a script to run the encryption operations. The three algorithms used were: DES [15], Blowfish [16], and AES [12], of which the first two operate on 8-byte data blocks, while the latter uses 16-byte blocks. The measurements for encrypting 10-Mbytes of data, including the initialization cost associated with each invocation of the encryption algorithms, are are shown in Table 1.

We point out that a constant initialization cost is associated with each algorithm, and that this cost becomes significant when invoking the encryption algorithm numerous times. Blowfish is the fastest of the three algorithms in terms of encryption speed, but it also incurs the largest initialization cost, and this is clearly illustrated during the encryption of the smaller data units. All algorithms display reduced encryption costs when using the 16-Kbyte blocks.

The main conclusion we draw from these results is that encrypting the same amount of data using fewer large blocks is clearly more efficient than using several smaller blocks. The difference in the cost between encrypting the smaller data units with that of the 16-Kbyte block is due to the start-up cost associated with the initialization of the encryption algorithms.

It is clearly advantageous to minimize the total number of encryption operations, while ensuring that data to be encrypted matches up with the encryption algorithm's block size, in order to ensure minimum padding. One obvious way is to cluster together the sensitive data which needs to be encrypted. This is exactly what we achieve in the storage model proposed in section 4.2.

4 Partition Plaintext and Ciphertext (PPC)

The majority of today's database systems use the N-ary Storage Model (NSM) [17] which we now describe.

4.1 N-ary Storage Model (NSM)

NSM stores records from a database continuously starting at the beginning of each page. An offset table is used at the end of the page to locate the beginning of

each record. NSM is optimized for transferring data to and from secondary storage and offers excellent performance when the query workload is highly selective and involves most record attributes. It is also popular since it is well-suited for online transaction processing; more so than another prominent storage model, the Decomposed Storage Model (DSM) [1].

NSM and Encryption: Even though NSM has been very successful as a storage model for RDBMS's, it is rather ill-suited for incorporating encryption. This is especially the case when a record has both sensitive and non-sensitive attributes. We will demonstrate, via an example scenario, exactly how the computation and storage overheads are severely increased when using encryption within the NSM model. We will work with a relation that has four attribute values: *EmpNo, Name, Department, and Salary.* Of these, only *Name* and *Salary* are sensitive and should, therefore, be encrypted. Figure 1 shows the NSM record structure.

Because only two attributes are sensitive, we would encrypt at the attribute level so as to avoid unnecessary encryption of non-sensitive data (see section 3.2). Consequently, we need one encryption operation for each attribute-value. [3]

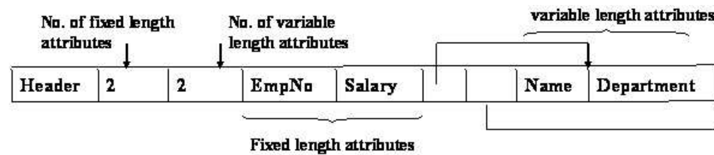

Fig. 1. NSM structure for our sample relation.

As described in section 3.1, using a symmetric-key algorithm in block cipher mode requires padding the input to match the block size. The padding can result in significant overhead when encrypting multiple values each needing a certain amount of padding. For example, since AES [12] uses 16-byte input blocks, if we needed to AES-encrypt a 2-byte attribute value, 14 bytes of padding would be required for each encryption.

To reduce costs outlined above, we must avoid small non-continuous sensitive plaintext values. Instead, we need to cluster them in some way, thereby reducing the number of encryption operations. Another potential benefit would be the reduced amount of padding involved: per cluster, as opposed to per attribute value.

Optimized NSM: Since using encryption in NSM is quite costly, we now suggest an optimized version of NSM. It involves storing all encrypted attribute values of one record sequentially (and, similarly, for all plaintext values). With

[3] If we instead encrypted at record or page level, non-sensitive attributes *EmpId* and *Department* would be also encrypted, thus requiring additional encryption operations. Even worse, for selection queries that only involve non-sensitive attributes, the cost of decrypting the data would still apply.

this optimization, one record ends up consisting of two parts: the ciphertext attributes followed by the plaintext (non-sensitive) attributes. The optimized version reduces padding overhead and eliminates multiple encryptions operations within one record. However, each record is still stored individually, meaning that, for each record, one encryption operation is needed. Moreover, each record is individually padded.

4.2 Partition Plaintext Ciphertext Model (PPC)

Our approach is to cluster encrypted data together while retaining the NSM's benefits. Partition Attribute Across (PAX) was recently proposed as an alternative model to NSM; it involves partitioning a page into mini-pages to improve upon cache performance [18]. Each mini-page represents one attribute in the relation. It contains the value for this attribute of each record stored in the page. Our model, referred to as Partition Plaintext and Ciphertext (PPC), employs an idea similar to that of PAX, in that pages are split into two mini-pages, based on plaintext and ciphertext attributes, respectively. Each record is likewise split into two sub-records.

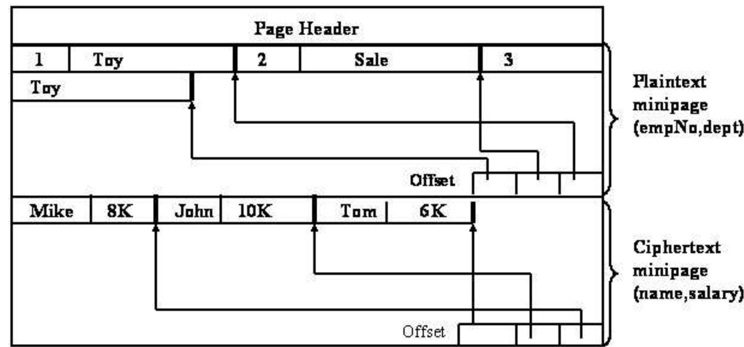

Fig. 2. Sample PPC page

PPC Overview: The main motivation behind PPC is to reduce encryption costs, including computation and storage costs, while keeping the NSM storage schema. We thus take advantage of NSM while enabling efficient encryption. Implementing PPC on existing DBMS's that use NSM requires only modifications to the page layout. The PPC model stores the same amount of records on each page as does NSM.

Within a page, PPC vertically partitions a record into two sub-records, one of which contains the plaintext, while the other – ciphertext, attributes. Both sub-records are organized in the same manner as NSM records. PPC stores all plaintext sub-records in the first part of the page, which we call a plaintext mini-page. The second part of the page consists of the ciphertext mini-page.

Each mini-page has the same structure as a regular NSM page and records within two mini-pages are stored in the same relative order. At the end of each mini-page is an offset table pointing to the end of each sub-record. Thus, a PPC page can be viewed as two NSM mini-pages. Specifically, if a page does not have any ciphertext, PPC is identical to NSM. Current database systems using NSM would only need to change the way they access pages in order to incorporate the PPC model.

Figure 2 shows an example of a PPC page. In it, three records are stored within one page. The plaintext mini-page contains non-sensitive attribute values EmpNo and Department and the ciphertext mini-page stores encrypted Name and Salary attributes. The advantage of encryption at the mini-page level can be seen by observing that only one encryption operation is needed *per page*, and, of course, only one decryption is required when the page is brought into memory and any sensitive attributes are accessed.

The PPC page header contains two mini-page pointers, in addition to the typical page header fields which include the starting address for both the plaintext and ciphertext mini-page.

Support from the Buffer Manager: When a PPC page is brought into a buffer slot, depending upon the nature of the access, the page may first need to be decrypted. The buffer manager, besides supporting a *write bit* to indicate whether a page has been modified, also needs to support an *encryption bit* to determine whether the ciphertext mini-page has already been decrypted. Initially, when the page is brought into a buffer slot, its write bit is off, and its encryption bit is on. Each read or update request to the buffer manager indicates whether a sensitive field needs to be accessed.

The buffer manager processes read requests in the obvious manner: if the encryption bit is on, it requests the mini-page to be decrypted and then resets the encryption bit. If the bit is off, the page is already in memory in plaintext form. Insertion, deletion and updating of records is also obvious: the write bit is set, while the encryption bit is only modified if any ciphertext has to be decrypted.

Whenever a page in the buffer is chosen by the page replacement policy to be sent to secondary storage, the buffer manager checks if the page's encryption bit is off and the write bit is on. If so, the cipher mini-page is first re-encrypted before being stored.

4.3 Analysis and comparisons of the PPC Model

We now compare NSM with the proposed PPC model. As will be seen below, PPC outperforms NSM irrespective of the encryption level of granularity in NSM. The two main advantages of PPC are: 1) considerably fewer encryption operations due to clustering of sensitive data, and 2) overhead for queries involving only non-sensitive attributes.

NSM with attribute-level encryption: The comparison is quite straightforward. NSM with attribute-level encryption requires as many encryption operations per record as there are sensitive attribute values. Records are typically small enough such that a large number can fit into one page, resulting in a large number of encryption operations per page. As already stated, only one decryption is per page is needed in the PPC model.

NSM with record-level encryption: One encryption is required for each record in the page. PPC requires only one encryption per page.

NSM with page level encryption: one encryption per page is required in both models. The only difference is for queries involving both sensitive and non-sensitive attributes. NSM performs an encryption operation regardless of the types of attributes, whereas, PPC only encrypts as necessary.

Optimized NSM: We introduced the optimized NSM model in section 4.1. It is similar to NSM with record-level encryption (each record requires one encryption). It differs from NSM in that extra overhead is incurred for non-sensitive queries. Again, PPC requires only one encryption operation per page as opposed to one per record for the optimized NSM.

Note that, for each comparison, the mode of access is irrelevant, i.e., whether records within the page are accessed sequentially or randomly, as is the case with the DSS and OLTP workloads. For every implementation, other than NSM with page-level encryption, one still needs to access and perform an encryption operation on the record within the page, regardless of how the record is accessed.

From the above comparisons, we observe that PPC has the same costs as the regular non-encrypted NSM when handling non-sensitive queries. On the other hand, when sensitive attributes are involved, it costs a single encryption operation. We thus conclude that *PPC is superior to all NSM variants* in terms of encryption-related computation overhead.

Although we have not performed a detailed comparison of storage overheads (due to padding), we claim that PPC requires the same (or less) amount of space than any of the NSM variants. In the most extreme case, encrypting at the attribute level requires each sensitive attribute to be padded. A block cipher operating on 128-bit blocks (e.g., AES) would, on the average, add 64 bits of padding to each encrypted unit. Only encrypting at the page level would minimize padding overhead, since only one unit is encrypted. This is the case for both NSM with page-level encryption and PPC.

4.4 Database operations in PPC

Basic database operations (insertion, deletion, update and scan) in PPC are implemented similar to their counterparts in NSM, since each record in the former is stored as two sub-records conforming to the NSM structure in the two NSM mini-pages. During insertion, a record is split into two sub-records,

each inserted into its corresponding mini-page (sensitive or non-sensitive). As described in Section 4.2, the buffer manager determines when the ciphertext mini-page needs to be en/de-crypted. Implementation of deletion and update operations is straight-forward.

When running a query, two scan operators are invoked, one each for the plaintext and ciphertext mini-page. Each scan operator sequentially reads a sub-record in the corresponding mini-page. If the predicate associated with the scan operator refers to an encrypted attribute, the scan operator indicates in its request to the buffer manager that it will be accessing an encrypted attribute. Scans could be implemented either using sequential access to the file containing the table, or using an index. Indexing on encrypted attributes is discussed in section 4.6 below.

4.5 Schema change

PPC stores the schema of each relation in the catalog file. Upon adding or deleting an attribute, PPC creates a new schema and assigns to it a unique version ID. All schema versions are stored in the catalog file. The header in the beginning of each plaintext and ciphertext sub-record contains the schema version that it conforms to. The advantage of having schemas in both plaintext and ciphertext sub-records is that, when retrieving a sub-record, only one lookup is necessary to determine the schema that the record conforms to.

Adding or deleting an attribute is handled similar to NSM, with the two following exceptions: (1) a previously non-sensitive attribute is changed to sensitive (i.e., it needs to be encrypted), and (2) a sensitive attribute is changed to non-sensitive, i.e., it needs to be stored as plaintext.

To handle the former, a new schema is first created and assigned a new version ID. Then, all records containing this attribute are updated according to the new schema. This operation can be executed in a lazy fashion (pages are read and translated asynchronously), or synchronously, in which case the entire table is locked and other transactions prevented from accessing the table until the reorganization is completed.

Changing an attribute's status from sensitive to non-sensitive can be deferred since it is of immediate importance whether the attribute is currently encrypted. For each accessed page, the schema comparison operation will indicate whether a change is necessary. At that time, the attribute will be decrypted and moved from the ciphertext to the plaintext mini-page.

4.6 Encrypted index

Index data structures are crucial components of any database system, mainly in terms of providing efficient range and selection queries. We need to determine the potential impact of encryption on the performance of the index. Note that, an index built upon a sensitive attribute must be encrypted, since it contains attribute values present in the actual relation.

There are basically two approaches for building an index based on sensitive attribute values. In the first, the index is based upon *ciphertext*, while, in the second, the intermediate index is based upon *plaintext* and the final index is obtained by encrypting the intermediate index. Based upon characteristics which make encryption in the PPC model efficient, we choose to encrypt at the page level, thereby encrypting each node independently. Whenever a specific index is needed in the processing of a query, the necessary parts of the data structure are brought into memory and decrypted. This approach provides full index functionality while keeping the index itself secure.

There are certain tradeoffs associated with either of the two approaches. Since encryption does not preserve order, the first approach is infeasible when the index is used to process range queries. Note that exact-match selection queries are still possible if encryption is deterministic.[4] When searching for an attribute value, the value is simply encrypted under the same encryption key as used in the index before the search. The second approach, in contrast, can support range queries efficiently, albeit, with additional overhead of decrypting index pages.

Given the tradeoff, the first strategy appears preferable for hash-indices or B-tree indices over record identifiers and foreign keys where data access is expected to be over equality constraints and not a range. The second strategy is preferable when creating B-trees where access may include range queries over data. Since the second strategy incurs additional encryption overhead, we discuss its performance in further detail.

When utilizing a B-tree index, we can assume that the plaintext representation of the root node already resides in memory. Using a tree of depth two, we only need to perform two I/O operations for a selection: one for the node at levels 1 and 2, and, correspondingly, two decryption operations. As described in section 3.6, we measure encryption overhead based upon the number of encryption operations performed. Since there is one decryption for each accessed node, the total overhead amounts to the number of accessed nodes multiplied by the start-up cost of the underlying encryption algorithm.

5 Experiments

We created an implementation of the PPC model based on MySQL version 4.1.0-alpha. This version provides all the necessary components of a database management system. Specifically, we modified the InnoDB storage model, by altering its page and record structure such as to create the plaintext and ciphertext mini-pages. We utilized the OpenSSL cryptographic library to write our encryption code, and used the Blowfish encryption algorithm. The experiments were conducted on a Windows XP platform with a 1.8 Ghz Pentium 4 processor and 384 MB RAM.

[4] If encryption is non-deterministic (randomized), which is common practice for preventing ciphertext correlation and dictionary attacks, the same cleartext encrypted twice yields two different ciphertexts.

5.1 PPC implementation details

Each page in InnoDB is by default 16KB. Depending on the number of encrypted attributes in the tables to be stored, we would split the existing pages into two parts, such as to accommodate the respective plaintext and ciphertext mini-pages. The partitioning of one record into its plaintext and ciphertext sub-record only takes place when the record is written to its designated page. The InnoDB record manager was modified such as to partition the records and store them in their corresponding mini-pages.

If a record is to be read from disk during a query execution, it is first accessed by InnoDB which then converts it to the MySQL record format. We modified this conversion functionality in the record manager, such as to determine whether the ciphertext part of the record is needed in the current query. If this is the case, the respective ciphertext mini-page is first decrypted before the ciphertext and plaintext sub-records are combined to a regular MySQL record. The ciphertext mini-page is only be re-encrypted if and when it is written back to disk.

5.2 Overview of experiments

For each of our experiments, we compared running times between *NSM with no encryption*, *NSM with page level encryption*, and *PPC*. We hereby refer to NSM with no encryption as *NSM* and NSM with page level encryption as *NSM-page*. Our goal was to verify that PPC would outperform NSM-page and, in the case where no sensitive values are involved in a query, achieve similar performance as NSM. Lastly, when all attributes involved are encrypted, we expected PPC and NSM-page to perform roughly equally.

As previously noted, PPC handles mixed queries very well, that is, queries involving both sensitive and non-sensitive attributes. To realize this feature, we created a schema whereby at least one attribute from each table was encrypted, and more than one in the larger tables (2 attributes in *lineitem* and 3 in *orders*). Specifically, we encrypted the following attributes: *l_partkey, l_shipmode, p_name, s_acctbal, ps_supplycost, c_name, o_orderdate, o_custkey, o_totalprice, n_name, and r_name.*

The three experiments we ran were all based on the TPC-H data set. First, we measured the loading time during bulk insertion of a 100-, 200-, and 500-MB database. Our second experiment consisted of running TPC-H query number 1, which is solely based on the *lineitem* table, while varying the number of encrypted attributes involved in the query such as to analyze PPC's performance. Finally, we compared the query response time of a sub-set of TPC-H queries, attempting to identify and highlight the properties of the PPC scheme.

5.3 Bulk Insertion

We compared the bulk insertion loading time required for a 100-, 200-, and 500-MB TPC-H database for each of the three models. The schema described in section 5.2 was used to define which attributes to encrypt in PPC. On average,

NSM-page and PPC incur a 24% and 15% overhead, respectively, in loading time when compared to NSM. As expected, PPC outperforms NSM-page because less data is encrypted, since not all attributes are considered sensitive.

It should be noted that a more detailed PPC implementation would perform some page reorganization if the relations inserted contained variable length attributes, such as to make best use of the space in the mini-pages. As stated in [18], the PAX model suffers a 2-10% performance penalty due to page reorganization, depending on the desired degree of space utilization in each mini-page. However, the PAX model creates as many mini-pages as there are records, while we consistently only create two.

5.4 Varying number of encrypted attributes

In this experiment, we wanted to run a query such as to isolate a single table and analyze the performance of PPC as we increased the number of encrypted attributes. Query 1, which only involves table *lineitem*, was chosen for this purpose, and executed over a 200 MB database. It contains 16 attributes, of which 7 are involved in the query. We then computed the running time for NSM and NSM-page, both of which remain constant as the number of encrypted attributes varies. Eight different instances of PPC were run, the first having no attributes encrypted and the last having 14 (we did not encrypt the attributes used as primary keys). Figure 3 illustrates our results.

A notable observation is that the overhead added when encrypting additional attributes in PPC is rather minimal. As expected, PPC with no (or only a few) attributes encrypted performs almost identically as NSM. Also, as more attributes are encrypted, PPC's performance begins to resemble that of NSM-page. The two last instances of PPC have relatively longer query execution times than the previous ones. This is due to the encryption of the *lineitem* variables *l_shipinstruct and l_comment*, which are considerably larger than any other in the table.

5.5 TPC-H Queries

In section 4.3 we described the advantages of the PPC model, which can be summarized as the reduced number of encryption operations required in addition to achieving similar performance as NSM for queries not involving sensitive attributes. PPC and NSM-page both require only one encryption operation per page, but NSM-page executes the operation irregardless of whether there are encrypted attributes in the query or not. In the following experiment we attempt to exploit the advantages that PPC offers, by comparing its performance with NSM-page when executing a chosen subset of the TPC-H queries on a 200 MB database. As in the bulk insertion experiment, we utilized our pre-defined schema (see section 5.2) to determine which attributes to encrypt when using PPC.

In figure 4 we refer to individual queries to highlight some of the interesting observations from the experiment. Queries 1 and 6 are range queries and, according to our PPC schema, have no sensitive attributes involved. The running

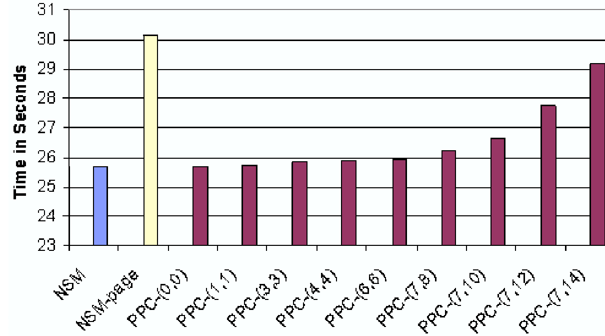

Fig. 3. Varying the number of encrypted attributes in TPC-H Query 1. PPC-(x, y) indicates that y attributes in lineitem were encrypted with x of them appearing in the query

time of NSM and PPC are, as expected, almost identical. However, since the *tables* involved do contain encrypted attributes, NSM-page will suffer from over-encryption and consequently have to decrypt the records involved in the query, thereby adding to its running time. Due to the simplicity of these two queries, the NSM-page encryption overhead is relatively small.

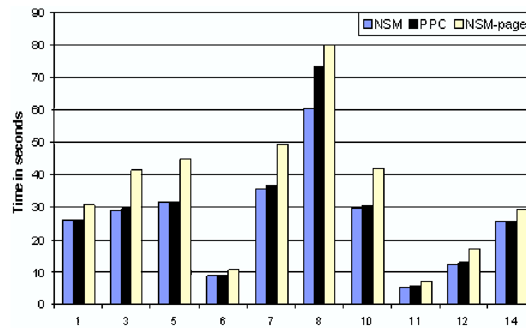

Fig. 4. Comparison of running times from selected TPC-H queries over a 200 MB database for NSM, PPC, and NSM-page.

Query 8 involves encrypted attributes from each of the three largest tables (*lineitem, partsupp, orders*), causing both PPC and NSM-page to incur relatively significant encryption costs. Again, PPC outperforms NSM-page because it has to decrypt less data. In contrast, query 10 involves encrypted attributes from four tables, of which only one is large (table *orders*). In this case, PPC performs well in comparison to NSM-page, because the latter needs to decrypt *lineitem* in addition to the other involved tables.

Overall, based on the 10 queries depicted in figure 4, PPC and NSM-page incur 6% and 33% overhead, respectively, in query response time. We feel that this experiment displayed the superior performance of PPC when executing queries involving both sensitive and non-sensitive attributes.

6 Conclusion

In this paper, we proposed a new DBMS storage model that facilitates an efficient implementation of encryption techniques. Our approach is based on grouping sensitive data, so as to minimize the number of necessary encryption operations, thus, greatly reducing encryption overhead.

References

1. Copeland, G. P., Khoshafian, S. F.: A Decomposition Storage Model. ACM SIGMOD International Conference on Management of Data. (1985) 268-269
2. Oracle Corporation: Database Encryption in Oracle9i. url=otn.oracle.com/deploy /security/oracle9i. (2001)
3. Department of Health and Human Services (U.S.). Gramm-Leach-Bliley (GLB) Act. FIPS PUB 81. url=www.ftc.gov/bcp/conline/pubs/buspubs/glbshort.htm. (1999)
4. Federal Trade Commission (U.S.): Health Insurance Portability and Accountability Act (HIPAA). url=www.hhs.gov/ocr/hipaa/privacy.html. (1996)
5. IBM Data Encryption for IMS and DB2 Databases, Version 1.1. url=http://www-306.ibm.com/software/data/db2imstools/html/ibmdataencryp.html. (2003)
6. He, J., Wang, M.: Cryptography and Relational Database Management Systems. Proceedings of the 5th International Database Engineering and Applications Symposium. (2001) 273-284
7. Hacigümüş, H., Iyer, B., Li, C., Mehrotra, S.: Executing SQL over Encrypted Data in the Database Service Provider Model. ACM SIGMOD Conference on Management of Data (2002)
8. Bouganim, L., Pucheral, P.: Chip-Secured Data Access: Confidential Data on Untrusted Servers. VLDB Conference, Hong Kong, China. (2002)
9. Hacigümüş, H., Iyer, B., Mehrotra, S.: Providing Database as a Service. ICDE. (2002)
10. Karlsson, J. S.: Using Encryption for Secure Data Storage in Mobile Database Systems. Friedrich-Schiller-Universitat Jena. (2002)
11. Rivest, R. L., Shamir, A., Adleman, L. M.: A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. Communications of the ACM. vol. 21, (1978)
12. NIST. Advanced Encryption Standard. FIPS PUB 197. (2001)
13. TPC Transaction Processing Performance Council. url=http://www.tpc.org
14. OpenSSL Project. url=http://www.openssl.org
15. NIST. Data Encryption Standard (DES). FIPS 46-3. (1993)
16. Schneier, B.: Description of a New Variable-Length Key, 64-Bit Block Cipher (Blowfish). Fast software Encryption, Cambridge Security Workshop Proceedings. (1993) 191-204
17. Ramakrishnan, R., Gehrke, J.: Database Management Systems, 2nd Edition, WCB/McGraw-Hill. (2000)
18. Ailamaki, A., DeWitt, D. J., Hill, M. D., Skounakis, M.: Weaving Relations for Cache Performance. The VLDB Journal. (2001) 169-180