

# Privacy Preserving Context Aware Publish Subscribe Systems

Mohamed Nabeel<sup>1</sup>, Stefan Appel<sup>2</sup>, Elisa Bertino<sup>1</sup>, and Alejandro Buchmann<sup>2</sup>

<sup>1</sup> Purdue University, West Lafayette, IN, USA,  
{nabeel, bertino}@cs.purdue.edu

<sup>2</sup> TU Darmstadt, Darmstadt, Germany,  
{appel, buchmann}@dvs.tu-darmstadt.de

**Abstract.** Modern pub/sub systems perform message routing based on the message *content* and allow subscribers to receive messages related to their subscriptions and the current *context*. Both content and context encode sensitive information which should be protected from third-party brokers that make routing decisions. In this work, we address this issue by proposing an approach that assures the confidentiality of the messages being published and subscriptions being issued while allowing the brokers to make routing decisions without decrypting individual messages and subscriptions, and without learning the context. Further, subscribers with a frequently changing context, such as location, are able to issue and update subscriptions without revealing the subscriptions in plaintext to the broker and without the need to contact a trusted third party for each subscription change resulting from a change in the context. Our approach is based on a modified version of the Paillier additive homomorphic cryptosystem and a novel group key management scheme. The former construct is used to perform privacy preserving matching, and the latter construct is used to enforce fine-grained encryption-based access control on the messages being published. We optimize our approach in order to efficiently handle frequently changing contexts. We have implemented our approach in a prototype using an industry strength JMS broker middleware. The experimental results show that our approach is highly practical.

## 1 Introduction

The publish/subscribe (pub/sub) paradigm is a well known approach for disseminating information between multiple interested parties in a decoupled and asynchronous manner [7]. Message producers submit messages to a broker network which routes them to interested subscribers. Subscribers express their interest by issuing subscriptions. Content-based pub/sub systems allow subscribers to express their interest based on the message content. This content can be an arbitrary payload, e.g., a set of attribute-value (att/val) pairs, XML documents, or combinations of different types. The supported message content depends on the pub/sub middleware which performs message routing and matching. In many systems it is common to use att/val pairs to describe content and to express subscription filters as logical expressions on these attributes (e.g., `MessageType == StockTickMessage ∧ StockPrice > 38`).

Context-sensitive message dissemination extends pub/sub content dissemination by taking into account the subscriber context [6]. Subscribers express interest in messages based on their current context, e.g., their current location. A major challenge for context-sensitive message dissemination is that the context of subscribers, and thus the subscriptions, change frequently over time, e.g., as the location changes. Context-dependent information dissemination is however a crucial requirement in many application scenarios. One example is a traffic information system (TIS) where information about the traffic situation is provided. Due to the characteristics of this scenario a pub/sub middleware supporting context-sensitive message dissemination is an appropriate infrastructure. Information about the current traffic situation is published and participants express their interest in information with subscriptions. As subscribers move and thus their contexts constantly change, subscriptions need to change accordingly, e.g., subscribers would typically be interested in traffic information along the route they are traveling.

A major shortcoming of existing context-based pub/sub approaches is that they do not assure privacy. In such systems, the broker receives subscriptions in plaintext and is thus aware of the context of subscribers. In case of the TIS example, this implies that brokers are aware of the exact position of subscribers. In order to assure privacy, we propose an approach to construct a privacy-preserving context-based pub/sub system. We extend and improve our preliminary work [15] in order to propose a new security model and construct our privacy preserving context-based pub/sub system.

In our previous model [15], each subscriber is required to submit a new subscription via a secure channel to a trusted third-party (TTP) that encrypts the subscription in a special way. This special encryption operation is called *blinding* and the encrypted value *blinded value*. Such blinded values are semantically secure (IND-CPA secure) where two blindings of the same value result in two different blinded values. The subscriber then registers with such blinded subscriptions at an untrusted broker. Such an approach allows the use of honest-but-curious brokers<sup>3</sup> to perform matching and routing on encrypted notifications using the blinded subscriptions. Although such an approach is privacy preserving, it is not suitable for context-based pub/sub systems. The reason is that as the context of subscribers changes frequently, subscriptions have to be updated often and involving a TTP to blind every subscription is not any longer feasible. Further, the previous model does not support fine-grained access control of notifications. Thus, a new security model and mechanisms are required whereby subscribers are allowed to create their own blinded subscriptions without compromising the security of the overall system and enforce fine-grained access control of notifications. In this work, we achieve our first objective by allowing authorized subscribers to create blinded subscriptions after obtaining some public security parameters at the time of registration. After the initial interactions, subscribers are not required to contact a TTP unless the public security parameters are updated. We achieve the second objective by introducing a fine-grained encryption-based access control mechanism. An advantage of such an approach compared to approaches based on shared secrets is that no secret information

---

<sup>3</sup> Brokers are obliged to follow the protocol, but they are curious to learn as much as possible during the execution of the protocol.

is given to subscribers to generate blinded subscriptions and therefore our approach avoids the problem of leakage of shared secrets by malicious subscribers.

Each notification in our approach is encrypted twice. The first encryption, referred to as blinding operation, is performed to blind each attribute value in the notification that is used by brokers to perform matching operations. The notification blinding is similar to the subscription blinding operation mentioned earlier except that the two operations use different blinding parameters so that certain parameters cancel off when a blinded notification and a subscription are homomorphically added by multiplying them. It should be noted that brokers cannot decrypt individual blinded values and they only learn a randomized difference between subscription and notification values when they perform matching operations. The second encryption, referred to as broadcast encryption, is performed to encrypt the payload of notifications based on fine-grained access control policies (ACPs). According to current initiatives on identity management [17], fine-grained ACPs are specified using the attributes of subscribers, referred to as *identity attributes*. Our broadcast encryption is based on a recently proposed group key management (GKM) scheme, referred to as attribute based GKM (AB-GKM) [20, 14, 16]. In the AB-GKM scheme, unlike conventional GKM schemes [2, 9], subscribers are allowed to dynamically derive the data decryption keys based on the attribute credentials they possess and some public information provided by publishers.

We also provide support for multiple publishers that produce messages with overlapping attribute sets; a blinded subscription may match notifications from several publishers. Especially for context-based pub/sub, it is crucial to support multiple publishers located in the same context, e.g., the same geographical region. We thus introduce context managers as TTPs in our approach. Context managers provide publishers as well as subscribers with information required to publish encrypted/blinded notifications and to issue blinded subscriptions. Once publishers and subscribers obtain the required security parameters, the context manager is responsible for controlling the level of protection. It decides when to renew and redistribute security parameters to publishers and subscriber in order to reduce the risk of adversaries learning the content of notifications and subscriptions.

We implement our scheme based on the Java Message Service (JMS), the de-facto industry standard for messaging. We chose Apache ActiveMQ as JMS broker and extended it to support subscription evaluation on encrypted data. This allows us to perform a realistic evaluation of our approach since ActiveMQ is used in many real-world production pub/sub systems.

Our paper is structured as follows. Section 2 introduces the cryptographic constructs used in our approach. Section 3 introduces context-based pub/sub systems and presents an overview of our solution. Sections 4 and 5 present the technical details of the modified Paillier cryptosystem which is used to blind subscriptions and notifications, and our overall scheme. Section 6 shows experimental results for various algorithms and the overall system implemented in Apache ActiveMQ. Section 7 discusses related work and Section 8 concludes the paper.

## 2 Background

### 2.1 Paillier homomorphic cryptosystem

The *Paillier homomorphic cryptosystem* is a public key cryptosystem by Paillier [18] based on the “Composite Residuosity assumption (CRA)”. The Paillier cryptosystem is homomorphic in that, by using a public key, the encryption of the sum  $m_1 + m_2$  of two messages  $m_1$  and  $m_2$  can be computed from the encryption of  $m_1$  and  $m_2$ . Our approach and protocols are inspired by how the Paillier cryptosystem works. Hence, we provide some internal details of the cryptosystem below so that readers can follow the rest of the paper.

#### **Key generation:**

Set  $n = pq$ , where  $p$  and  $q$  are two large prime numbers. Set  $\lambda = \text{lcm}(p-1, q-1)$ , i.e., the least common multiple of  $p-1$  and  $q-1$ . Randomly select a base  $g \in \mathbb{Z}/(n^2)^\times$  such that the order of  $g_p$  is a multiple of  $n$ . Such a  $g_p$  can be efficiently found by randomly choosing  $g_p \in \mathbb{Z}/(n^2)^\times$ , then verifying that  $\text{gcd}(L(g_p^\lambda \pmod{n^2}), n) = 1$ , where  $L(u) = (u-1)/n$ . for  $u \in S_n = \{u < n^2 \mid u = 1 \pmod{n}\}$ . In this case, set  $\mu = (L(g_p^\lambda \pmod{n^2}))^{-1} \pmod{n}$ . The public encryption key is a pair  $(n, g_p)$ . The private decryption key is  $(\lambda, \mu)$ , or equivalently  $(p, q, \mu)$ .

#### **Encryption $E(m, r)$ :**

Given plaintext  $m \in \{0, 1, \dots, n-1\}$ , select a random  $r \in \{1, 2, \dots, n-1\}$ , and encrypt  $m$  as  $E(m, r) = g_p^m \cdot r^n \pmod{n^2}$ . When the value of  $r$  is not important to the context, we sometimes simply write a short-hand  $E(m)$  instead of  $E(m, r)$  for the Paillier ciphertext of  $m$ .

#### **Decryption $D(c)$ :**

Given ciphertext  $c \in \mathbb{Z}/(n^2)^\times$ , decrypt  $c$  as  $D(c) = L(c^\lambda \pmod{n^2}) \cdot \mu \pmod{n}$ .

In the construction of our pub/sub system, the Paillier homomorphic cryptosystem is used in a way that public and private keys are judiciously distributed among publishers, subscribers, and brokers such that the confidentiality and privacy are assured based on homomorphic encryption. A detailed description of the construction is presented in Section 4.

### 2.2 Attribute Based Group Key Management

Broadcast Group Key Management (BGKM) schemes [24, 20] are a special type of GKM scheme whereby the rekey operation is performed with a single broadcast without requiring private communication channels. Unlike conventional GKM schemes, BGKM schemes do not give subscribers private keys. Instead subscribers are given a secret which is combined with public information to obtain the actual private keys. Such schemes have the advantage of requiring a private communication only once for the initial secret sharing. The subsequent rekeying operations are performed using one broadcast message. Further, in such schemes achieving forward and backward security requires only to change the public information and does not affect the secrets given to existing subscribers. However, BGKM schemes do not support group membership policies over a set of attributes. A recently proposed attribute based GKM (AB-GKM)

scheme [13] provides all the benefits of BGKM schemes and also supports attribute based access control policies (ACPs).

Subscribers are required to show their identity attributes to the group controller to obtain secrets using the AB-GKM scheme. In order to hide the identity attributes from the group controller while allowing only valid subscribers to obtain secrets, we utilize oblivious commitment based envelope (OCBE) protocols [8]. We omit the details of the OCBE protocols due to the page limit.

The idea behind the AB-GKM scheme is as follows. A separate BGKM instance for each attribute condition, which is a predicate over an attribute, is constructed. The ACP, a Boolean expression over attribute conditions, is embedded in an access structure  $\mathcal{T}$ .  $\mathcal{T}$  is a tree with the internal nodes representing threshold gates and the leaves representing BGKM instances for the attributes.  $\mathcal{T}$  can represent any monotonic policy. The goal of the access tree is to allow deriving the group key for only the subscribers whose attributes satisfy the access structure  $\mathcal{T}$ .

The AB-GKM scheme consists of five algorithms: **Setup**, **SecGen**, **KeyGen**, **KeyDer** and **ReKey**. **Setup** initializes the system. **SecGen** generates a unique secret for each attribute condition. For a given ACP, **KeyGen** creates a symmetric key, public information and an access structure. **KeyDer** derives the symmetric key given one or more secrets and public information. **ReKey** regenerates the symmetric key and public information.

### 3 Overview

Our approach requires a modification of the matching algorithm inside the message broker to support the evaluation of blinded subscriptions against blinded notifications without decrypting them. Our system also supports fine-grained encryption based access control over notifications. Publishers encrypt notifications so that only authorized subscribers can derive the key and decrypt the notifications. In order to assure privacy, publishers and subscribers must perform an initialization to obtain secrets and public parameters that they later need for encryption and blinding operations. In this section, we give an overview of the modified Paillier cryptosystem and our context-based pub/sub model and present our system architecture. We describe the initialization phase as well as the regular runtime behavior. Finally, we present the trust model assumed in our approach.

#### 3.1 Modified Paillier Cryptosystem

In our work, we adapt the Paillier cryptosystem so that brokers can perform matching operations without decrypting individual subscriptions and notifications. A high level overview of the modifications we perform to the Paillier cryptosystem and the rationale behind our modifications are provided below. We first shift the computation towards encryption so that decryption is computationally more efficient than the Paillier decryption. We also allow brokers to perform certain operations without knowing the private key. Such shifting of computation improves the performance of the overall

pub/sub system since publishers and subscribers, which perform encryption, are typically distributed to many nodes while brokers have to handle notifications from many publishers and subscribers. Thus, by making decryption efficient, we eliminate a bottleneck in the system and improve the overall efficiency. We also blind the encrypted values and make  $\mu$ , a parameter of the Paillier cryptosystem, public so that individual values cannot be decrypted, but a blinded subscription and a blinded notification can be multiplied together to obtain the difference. In order to make the correct matching decision by calculating the difference, we limit the domain size ( $l$  bits) of the subscription and notification values. We assume that  $l$  is much smaller compared to the plaintext space of the Paillier cryptosystem,  $n$ . For example, the subscription and notification for the attribute age, which may take values from 0 to 200, can be represented using a domain size of 8 bits. Since the domain size is much smaller than the plaintext space, brokers can make the matching decision by calculating the difference as follows: If the difference between a notification and a subscription value is in the first half of the plaintext space, the difference is positive and the notification value is greater than or equal to the subscription value. Otherwise, the difference is negative and the notification value is less than the subscription value.

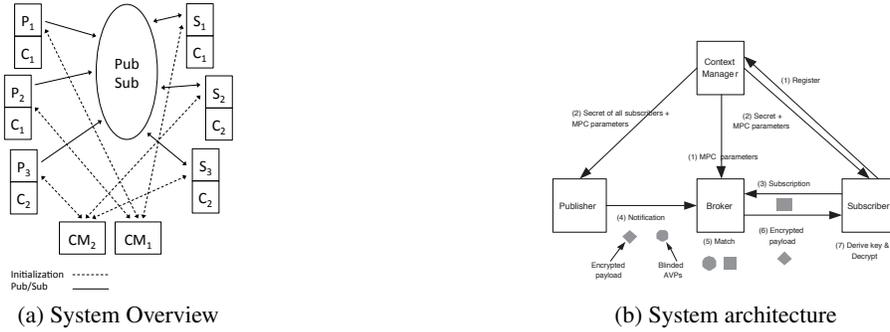
The above modifications allow brokers to make matching decisions without learning the actual values. However the modified matching protocol still reveals the actual difference between the notification value and subscription values which leaks information about these values. In order to address this issue, we introduce controlled random values to the subscription and notification blinding operations so that the difference is randomized. The brokers can still make correct matching decisions by comparing which half the computed difference falls in the plaintext space, without however learning the actual difference.

### 3.2 System Architecture

We assume that logical expressions based on att/val pairs are used by subscribers to express their interest in notifications. We distinguish between two sets of attributes, namely context set and static set, and an additional payload which by itself can be a set of att/val pairs. The context set contains attributes representing the subscriber context, e.g., location. The static set represents general attributes describing the message, e.g., message type.

Our approach consists of four entities (see Figure 1a and Figure 1b): publisher, subscriber, broker, and context manager. In order to assure the privacy of subscriptions, subscribers must hide the content of their subscriptions from brokers. Further, since the attributes in the context set frequently change, subscribers must be able to update their subscriptions without contacting a TTP, referred to as the context manager. In our approach, publishers and subscribers communicate with context managers only either during the initialization phase or when security parameters change in order to obtain secrets and public parameters.

A context manager is responsible for a certain context in the system. The context is described by a set of attributes, such as location with attributes *latitude* and *longitude*. To assure the privacy of the context, subscribers must be able to issue and update



**Fig. 1.** System details: Relations between publishers (P), subscribers (S), contexts (C), and context managers (CM)

subscriptions with blinded context set attribute values and publishers must be able to publish notifications and blind context set attribute values.

When publishers join the system and want to publish messages for a certain context, they first contact the responsible context manager. After successful authentication the context manager provides some security parameters and a set of secrets corresponding to valid subscribers to the publishers which allow them to blind and encrypt notifications. The blinding operation is performed using the modified Paillier cryptosystem described in the next section and the encryption is performed using the symmetric key generated using the AB-GKM's key generation algorithm.

A similar bootstrapping process is necessary for subscribers. They contact the context manager and receive some other security parameters which allow them to issue blinded subscriptions for a certain context and secrets for identity attributes they have. Further, our scheme allows the subscribers to update their subscriptions within a context without contacting the context manager.

Brokers only receive blinded and encrypted notifications, and blinded subscriptions. Upon receiving a notification, they execute the matching operation on the blinded notification and the blinded subscriptions to make forwarding decisions. If a notification matches a subscription, the broker strips the blinded portion of the notification and sends only the encrypted notification to matching subscribers. It should be noted that brokers do not learn the individual notification and subscription values during matching operations.

Once subscribers receive encrypted notifications, using the secrets obtained from the context manager during the initialization phase, they can derive the decryption key using the AB-GKM's key derivation algorithm. The AB-GKM scheme makes sure that only valid subscribers can derive the key and hence decrypt notifications. Since subscribers are not given decryption keys during the registration and they must dynamically derive the keys, our approach can efficiently handle subscriber revocations and additions as well as access control policy changes without affecting the existing subscribers.

### 3.3 Trust Model

We consider threats and assumptions from the point of view of publishers and subscribers with respect to third-party brokers. We assume that brokers are honest but curious; they perform pub/sub operations correctly, but are curious to know the notifications and subscriptions. In other words, brokers are not trusted for the confidentiality of the notifications and subscriptions. The context manager is fully trusted. Publishers are trusted to keep the secrets obtained from the context manager confidential and to perform notification blinding and encryption as specified. Subscribers are not trusted in our system. They can decrypt encrypted notifications only if they have valid credentials. Brokers may collude with one another as well as with malicious subscribers.

## 4 Modified Paillier Cryptosystem

In this section we provide the details of our modified Paillier cryptosystem.

### Making $\mu$ Public:

Recall that in the original Paillier cryptosystem, the tuple  $(\lambda, \mu)$  is the private key. However,  $\mu$  does not need to be private since it is hard to decrypt an encrypted message by only knowing  $\mu$ . In order to decrypt, one needs to know both  $\lambda$  and  $\mu$ . It can be shown that if a probabilistic polynomial time (PPT) adversary can obtain  $\lambda$  from  $\mu$ , it can solve the discrete logarithm problem (DLP). Since DLP is a known hard problem, it is hard to obtain  $\lambda$  from  $\mu$ . Hence, we can make  $\mu$  public while achieving the same security guarantees as the unmodified Paillier cryptosystem. We take advantage of this fact in order to shift the computation towards encryption and make decryption light weight.

### Shifting the Computation:

With the above modification, the new public and private keys are  $(n, g_p, \mu)$  and  $\lambda$  respectively. First, we modify the Paillier cryptosystem so that anyone can decrypt using the new public key, but only those holding the private key can encrypt. This is similar to how digital signatures work. The following equations show the modifications to the encryption and decryption algorithms: **Encryption**  $E'(m, r, \lambda) = E(m, r)^\lambda = g_p^{m\lambda} \cdot r^{n\lambda} \pmod{n^2} = c$ . and **Decryption**  $D(c) = L(c \pmod{n^2}) \cdot \mu \pmod{n}$ .

It should be noted that one can perform all the homomorphic operations on the modified Paillier cryptosystem similar to the unmodified Paillier cryptosystem as the above modification only shift the computation from decryption to encryption.

### Computing Differences (but not Individual Values):

With the shift of computation described above, anyone can find the difference by simply decrypting each value. However, such an approach does not assure the privacy of individual values. Therefore, we introduce an additional parameter to the encryption operation in order to allow one to compute the difference while at the same time not allowing the decryption of individual values.

Assume that there are two values  $x_1$  and  $x_2$ . We perform the following modification to the encryption operation so that a decryptor can learn the difference  $(x_1 - x_2)$  without learning either  $x_1$  or  $x_2$ . We call the modified encryption as blinding operation. The modified encryption  $E''(x_1, x_2)$  outputs  $x'_1$  and  $x'_2$  where  $x'_1 = g^t \cdot E'(x_1, r_1) \pmod{n^2}$  and  $x'_2 = g^{-t} \cdot E'(-x_2, r_2) \pmod{n^2}$ .

Notice that even though the decryptor knows  $\mu$ , it can decrypt neither  $x'_1$  nor  $x'_2$  as they are modular multiplied with  $g^t$  and  $g^{-t}$  respectively. Due to the additive homomorphic property, the following holds:  $x'_1 \cdot x'_2 = E'(x_1 - x_2, r_3)$ .

Since the multiplication of  $x'_1$  and  $x'_2$  cancels the blinding parameters, anyone can compute the difference as follows using the public key of the modified Paillier cryptosystem:  $D(x'_1 \cdot x'_2) = x_1 - x_2$ .

**Allowing Comparison:**

Recall that in Section 3.1 we introduced the notion of domain size which is much smaller than the plaintext space of the Paillier cryptosystem. According to our assumptions,  $0 \leq x_1, x_2 \leq 2^l$  where  $l$  is the domain size and  $2^l \ll n$ . Let the difference of  $x_1$  and  $x_2$  be  $d$ . Due to the restriction on the domain size,  $d$  is either between 0 and  $2^l$  or  $n - 2^l$  and  $n$ . We use this fact to compare the numbers; if  $d \leq 2^l$ , then  $x_1 \geq x_2$  and if  $d > n - 2^l$ , then  $x_1 < x_2$ .

During the above comparison process, the party performing the comparison learns the difference  $d$  which leaks certain information about the actual values. Hence, the comparison is not privacy preserving. We introduce a technique to randomize the difference so that it is difficult to learn the difference yet the party can learn the exact comparison result.

Notice that in the above calculation, we only utilize a small range of the plaintext space to make the comparison decision. We utilize the unused space in the plaintext space in the above calculation to randomize the difference while still allowing one to make the correct matching decision. The key idea is to expand the difference from  $0 - 2^l$  to  $0 - n/2$  and  $(n - 2^l) - n$  to  $n/2 - n$  by introducing controlled random values to the encryption operation. We introduce two random values  $r_p$  and  $r_q$  during the encryption operation shown below:  $x''_1 = g^t \cdot E'(x_1, r_1)^{r_p} E'(r_q) \pmod{n^2}$  and  $x''_2 = g^{-t} \cdot E'(-x_2, r_2)^{r_p} \pmod{n^2}$ . The decryption results in the following output:  $D(x''_1 \cdot x''_2) = r_p(x_1 - x_2) + r_q = d'$ .

$r_p$  and  $r_q$  are randomly selected so that  $d' \leq n/2$  if  $x_1 \geq x_2$  and  $d' > n/2$  if  $x_1 < x_2$ . Each time a party performs the comparison it gets a different  $d'$  due to the random values and thus the difference preserves the privacy of the individual values under comparison.

## 5 Privacy-Preserving Brokering Scheme

In this section, we describe in detail our approach to construct a privacy preserving context-aware publish subscribe system using the modified Paillier cryptosystem presented in Section 4 and the AB-GKM scheme.

As introduced in Section 3, there are four entities in our system: context manager, publisher, subscriber, and broker. The context manager acts as a TTP and generates the parameters for the modified Paillier cryptosystem and manages secrets obtained by the SecGen algorithm of the AB-GKM scheme to subscribers based on the identity attributes they possess. The context manager maintains a set of contexts  $\mathcal{C}$  and a set of secrets issued to subscribers. Each context  $C_i \in \mathcal{C}$  is a tuple of the following form:  $C_i = \langle \lambda_i, \mu_i, t_i, r_i \rangle$ , where  $\lambda_i$  and  $\mu_i$  are Paillier parameters for  $E'$  and  $D'$  algorithms.  $t_i$  and  $r_i$  are random values.

Brokers match notifications with subscriptions within the same context only.  $\mu_i$  values are public.  $\lambda_i$  and  $t_i$  values are private to the context manager.

**Subscriber Registration:**

Each subscriber registers with the context manager. Let the context of a random subscriber be  $C_i$ . During the registration, the subscriber receives the following values from the context manager:  $E'(-r_i)$ ,  $E'(-1)$ , and  $g^{-t_i} \cdot E'(-r_i)$ .

These parameters are used by the subscriber to blind subscriptions. Since  $\mu_i$  is public, the subscriber may decrypt  $E'(-r_i)$  using  $D'$  and obtain  $r_i$ . However, the subscriber can recover neither  $g^{-t_i}$  nor  $t_i$  from  $g^{-t_i} \cdot E'(-r_i)$ .

Using the SecGen algorithm of the AB-GKM scheme, each subscriber  $i$  also obtains secrets  $s_{ij}$  for each identity attribute  $j$  they possess from the context manager. These secrets are later used to derive the decryption key using the key derivation (KeyDer) algorithm of the AB-GKM scheme and decrypt notifications.

It should be noted that the identity attributes are not revealed to the context manager in plaintext as the SecGen algorithm internally utilizes the OCBE protocols. Thus the privacy of the identity attributes are preserved from the context manager.

**Publisher Registration:**

Each publisher also registers with the context manager. Let the context of a random publisher be  $C_i$ . During the registration, the publisher receives the following values from the context manager:  $E'(r_i)$ ,  $E'(1)$ , and  $g^{t_i} \cdot E'(r_i)$ .

Similar to subscriber registration, these parameters are used by the publisher to blind notifications. Since  $\mu_i$  is public, the publisher may decrypt  $E'(r_i)$  using  $D'$  and obtain  $r_i$ . Notice that the context manager may provide  $E'(1)$  and  $r_i$ , and allow the publisher to compute  $E'(r_i)$  homomorphically instead of providing the value directly. Also, notice that the publisher can recover neither  $g^{t_i}$  nor  $t_i$  from  $g^{t_i} \cdot E'(r_i)$ .

In addition to the above modified Paillier cryptosystem parameters, each publisher also obtains the set of secrets issued to subscribers using the SecGen algorithm. These secrets are used to selectively encrypt notifications based on the identity attributes that subscribers possess. The publisher first uses the key generation (KeyGen) algorithm of the AB-GKM scheme to generate the encryption key based on these secrets and then encrypts the notifications using the generated key. Notice that these secrets do not reveal the actual identity attributes of subscribers. Thus the identity attributes of subscribers are preserved from publishers as well.

**Notifications:**

Assume that a publisher wants to publish a notification for the attributes  $a_1$  and  $a_2$  with values  $v_1$  and  $v_2$  respectively. The publisher first blinds  $v_1$  and  $v_2$  to create  $v'_1$  and  $v'_2$  respectively using the modified Paillier cryptosystem presented in Section 4. We show the blinding operation for a general value  $v$  as follows:  $v' = g^{t_i} \cdot E'(r_i) \cdot E'(r_i(v - 1)) \cdot E'(r_v) = g^{t_i} \cdot E'(r_i v + r_v)$ , where  $r_v$  is a controlled random value selected by the publisher.  $E'(r_i(v - 1))$  is homomorphically computed using  $E'(r_i)$ . Notice that this value can be computed efficiently using fast multiplication.

Based on the ACP and the secrets issued to subscribers, the publisher generates the encryption key  $k$  using the KeyGen algorithm of the AB-GKM scheme. It then encrypts the payload of the notification ( $a_1 = v_1$ ,  $a_2 = v_2$ ) using the key  $k$ . We denote the encrypted payload as  $E_k(\text{payload})$ . The publisher sends the blinded and encrypted

notification  $((a_1 = v'_1, a_2 = v'_2), E_k(\text{payload}))$  to brokers. Notice that brokers cannot decrypt any of the blinded values as well as the encrypted payload.

An advantage of having two sets of encrypted values for each notification is that our approach allows to perform privacy preserving matching and enforce fine-grained ACPs independently.

**Subscriptions:**

Assume that a subscriber wants to subscribe for the attribute  $a_1$  with the value  $x$ . The subscriber blinds  $x$  and creates  $x'$  as follows:  $x' = g^{-t_i} \cdot E'(-r_i) \cdot E'(r_i(1-x)) = g^{-t_i} \cdot E'(-r_i x) \cdot E'(r_i(1-x))$  is homomorphically computed using  $E'(-r_i)$ .

The tuple  $(a, x', \alpha)$ , where  $\alpha = \{<, >\}$ , is sent to brokers. Similar to the blinded notifications, notice that brokers cannot decrypt  $x'$ .

The tuple  $(a, x', \alpha)$  represents a single attribute condition and we call such a subscription an *atomic subscription*. A subscription, in general, can be a Boolean expression over a set of atomic subscriptions and is called a *composite subscription*. Notice that the atomic subscription intentionally leaves the equality comparison operator. The motivation behind such a scheme is to further hide subscriptions and notifications from brokers. In our scheme, equality subscriptions are performed using range queries so that brokers cannot distinguish between equality subscriptions and range queries. In order to submit an equality subscription for attribute  $a$  with the value  $x$ , the subscriber submits the query  $(a, x_1, >) \wedge (a, x_2, <)$ , where  $x_1$  and  $x_2$  are the blinded values of  $x - 1$  and  $x + 1$ . Since the blinded values are semantically secure, the conjunctive query does not reveal any information about the range and therefore brokers cannot distinguish an equality subscription from a general range subscription.

**Broker Matching:**

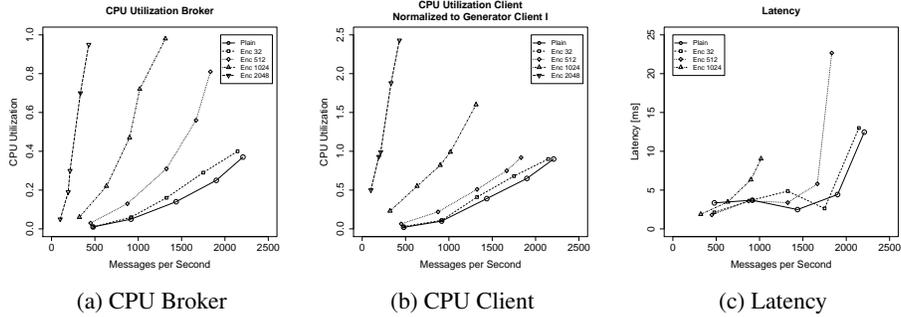
For each context  $C_i$ , brokers receive  $\mu_i$ . Assume that for the context  $C_i$ , a broker has received the blinded notification and subscription values  $v'_1$  and  $x'$  respectively for the attribute  $a_1$ . As mentioned above, we emphasize that the broker can decrypt neither  $v'_1$  nor  $x'$ . As described in Section 4, the broker computes the randomized difference  $d'$  as follows:  $d' = D'(v' \cdot x') = r_i(v - x) + r_v$

It decides  $v_1 > x$  if  $d' \leq n/2$  and  $d' \neq 0$ ,  $v_1 < x$  otherwise. The above matching algorithm is described for an atomic subscription. Usually notifications contain more than one attribute and the broker has to match such notifications with either atomic or composite subscriptions. The matching for a composite subscription is performed by evaluating each atomic subscription in the subscription and evaluating the Boolean expression. After successful matching, the broker forwards only the encrypted payload  $E_k(\text{payload})$  to matching subscribers. Subscribers having valid credentials can derive the key  $k$  using the KeyDer algorithm of the AB-GKM scheme and access the payload of the notification.

## 6 Implementation and Evaluation

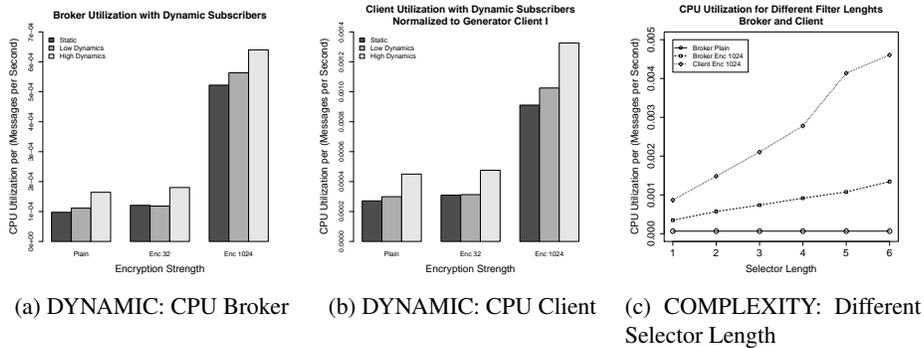
We extended Apache ActiveMQ, a production strength and widely used messaging middleware, with our proposed mechanisms for privacy preserving pub/sub. We used a distributed setup with one 8 core and one 16 core Intel Xeon machine as message generating clients connected to an 8 core Intel Xeon 3.5 GHz machine running Linux 2.6

that acts as our extended ActiveMQ broker. An extended discussion of the evaluation is presented in our technical report [11].



**Fig. 2.** Results CONSTANT Scenario: CPU Utilization and Latency for Different Message Rates and Blinding Strength

Figures 2a, 2b, and 2c show CPU utilization and latency for the CONSTANT scenario where a constant number of entities are used.



**Fig. 3.** Results DYNAMIC and COMPLEXITY Scenarios: Frequent Re-Subscriptions and Different Filter Lengths

Figures 3a and 3b show the CPU utilization for the DYNAMIC scenario where subscribers leave and join the system at a certain rate to simulate context changes and user churn. The results show that the additional overhead of joining subscribers is not the factor that dominates CPU utilization at the broker. Further, the increase occurs independent of blinding which shows that the overhead of joining is inherent to the

broker. On the client side CPU utilization is increased by about one third in the High Dynamics configuration compared to the Static configuration.

The results for the COMPLEXITY scenario, where subscribers use different message selector lengths, are shown in Figure 3c. For unencrypted subscriptions an increase in CPU utilization is not observable. For blinded subscriptions the CPU utilization of the broker increases slightly with increasing complexity. The utilization on the client side increases faster since for each message all attributes have to be blinded, but the broker does not necessarily evaluate the whole message selector.

## 7 Related Work

In this section, we compare our approach with existing work on secure content based pub/sub systems, and search over encrypted data.

**Secure content based pub/sub systems:** Most of prior work on data confidentiality in the context of content based pub/sub systems is based on the assumption that brokers are trusted with respect to the privacy of the subscriptions [1, 22, 12]. However, when such an assumption does not hold, both publication confidentiality and subscription privacy are at risk. Further, such approaches limit brokers' ability to make routing decisions based on the content of the messages and thus their applicability is very limited. Approaches have also been proposed to assure confidentiality/privacy in the presence of untrusted third-party brokers. These approaches however suffer from several limitations [19, 23, 10, 5]: inaccurate content delivery, because of the limited ability of brokers to make routing decisions based on content; weak security protocols; lack of privacy guarantees. For example, some of these approaches are prone to false positives, that is, sending irrelevant content to subscribers.

**Search over encrypted data:** Search on encrypted data is a privacy-preserving technique used in the *outsourced storage model* where a user's data are stored on a third-party server and encrypted using the user's public key. The user can use a query in the form of an encrypted token to retrieve relevant data from the server, whereas the server does not learn any more information about the query other than whether the returned data matches the search criteria. There have been efforts to support simple equality queries [21, 3] and more recently complex ones involving conjunctions and disjunctions of range queries [4]. These approaches cannot be applied directly to the pub/sub model.

## 8 Conclusions

We proposed an approach to construct a privacy preserving context-based pub/sub system. Our approach assures the confidentiality of notifications and subscriptions from third-party brokers while allowing the brokers to perform matching operations. Further, publishers are able to enforce fine grained ACPs over encrypted notifications. Our solution is based on a modified Paillier cryptosystem and a recent group key management scheme. Unlike the existing approaches, in our approach, publishers and subscribers are able to generate notifications and subscriptions without contacting a TTP. We implemented our approach in ActiveMQ and the experimental results show that our approach

is practical and efficient. As part of future work, we plan to investigate performance improvement techniques and specifically the trade-off between subscriber privacy and the message routing efficiency.

## References

1. E. Bertino, B. Carminati, E. Ferrari, B. Thuraisingham, and A. Gupta. Selective and authentic third-party distribution of XML documents. *IEEE TKDE*, 16(10):1263–1278, Oct. 2004.
2. E. Bertino and E. Ferrari. Secure and selective dissemination of XML documents. *ACM TISS*, 5(3):290–331, 2002.
3. D. Boneh, G. Crescenzo, R. Ostrovsky, and G. Persiano. Public-key encryption with keyword search. In *EUROCRYPT*, 2004.
4. D. Boneh and B. Waters. Conjunctive, subset, and range queries on encrypted data. *CRYPTO*, 2007.
5. S. Choi, G. Ghinita, and E. Bertino. A privacy-enhancing content-based publish/subscribe system using scalar product preserving transformations. In *DEXA*, 2010.
6. G. Cugola, A. Margara, and M. Migliavacca. Context-aware publish-subscribe: Model, implementation, and evaluation. In *ISCC*, 2009.
7. P. Eugster, P.A. Felber, R. Guerraoui, and A. Kermarrec. The many faces of publish/subscribe. *ACM Computing Surveys*, 35(2):114–131, 2003.
8. J. Li and N. Li. OACerts: Oblivious attribute certificates. *IEEE TDSC*, 3(4):340–352, 2006.
9. G. Miklau and D. Suci. Controlling access to published data using cryptography. In *VLDB*, 2003.
10. K. Minami, A. J. Lee, M. Winslett, and N. Borisov. Secure aggregation in a publish-subscribe system. In *WPES*, 2008.
11. M. Nabeel, S. Appel, E. Bertino, and A. Buchmann. Privacy preserving context aware publish/subscribe systems. Technical Report 2013-1, Purdue University, CERIAS, 2013.
12. M. Nabeel and E. Bertino. Secure delta-publishing of XML content. In *ICDE*, 2008.
13. M. Nabeel and E. Bertino. Towards attribute based group key management. In *CCS*, 2011.
14. M. Nabeel, E. Bertino, M. Kantarcioglu, and B. M. Thuraisingham. Towards privacy preserving access control in the cloud. In *CollaborateCom*, 2011.
15. M. Nabeel, N. Shang, and E. Bertino. Efficient privacy preserving content based publish/subscribe systems. In *SACMAT*, 2012.
16. M. Nabeel, N. Shang, and E. Bertino. Privacy preserving policy based content sharing in public clouds. *IEEE TKDE*, 2012.
17. OpenID. <http://openid.net/> [Last accessed: July 18, 2012].
18. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, 1999.
19. C. Raiciu and D. S. Rosenblum. Enabling confidentiality in content-based publish/subscribe infrastructures. In *Securecomm*, 2006.
20. N. Shang, M. Nabeel, F. Paci, and E. Bertino. A privacy-preserving approach to policy-based content dissemination. In *ICDE*, 2010.
21. D. X. Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *SP*, 2000.
22. M. Srivatsa and L. Liu. Securing publish-subscribe overlay services with eventguard. In *CCS*, 2005.
23. M. Srivatsa and L. Liu. Secure event dissemination in publish-subscribe networks. In *ICDCS*, 2007.
24. X. Zou, Y. Dai, and E. Bertino. A practical and flexible key management mechanism for trusted collaborative computing. *INFOCOM*, 2008.