International Journal of Science Engineering and Advance Technology

# Secure Broker-Less Publish/Subscribe Systems Using Identity-Based Encryption

Vundrajavarapu Ajay Kumar[1], Sri K.Satyanarayana Raju[2]
[1] M.Tech (IT), S.R.K.R Engineering College, A.P., India.
[2]Asst Professor, Dept. of Information Technology, S.R.K.R Engineering College, A.P., India.

***Abstract — publish–subscribe is a messaging pattern where senders of messages, called publishers, do not program the messages to be sent directly to specific receivers, called subscribers. provisioning of basic security mechanisms such as authentication and confidentiality is highly challenging in a content based publish/subscribe system. Authentication of publishers and subscribers is difficult to achieve due to the loose coupling of publishers and subscribers. Likewise, confidentiality of events and subscriptions conflicts with content-based routing. This paper presents a novel approach to provide confidentiality and authentication in a broker-less content-based publish/subscribe system. The authentication of publishers and subscribers as well as confidentiality of events is ensured, by adapting the pairing-based cryptography mechanisms, to the needs of a publish/subscribe system. Furthermore, an algorithm to cluster subscribers according to their subscriptions preserves a weak notion of subscription confidentiality. In addition to our previous work [20], this paper contributes 1) use of searchable encryption to enable efficient routing of encrypted events, 2) multicredential routing a new event dissemination strategy to strengthen the weak subscription confidentiality, and 3) thorough analysis of different attacks on subscription confidentiality. The overall approach provides fine-grained key management and the cost for encryption, decryption, and routing is in the order of subscribed attributes. Moreover, the evaluations show that providing security is affordable w.r.t. 1) throughput of the proposed cryptographic primitives, and 2) delays incurred during the construction of the publish/subscribe overlay and the event dissemination.***

***Keywords —Content-based,publish/subscribe,peer-to-peer, broker-less, security, identity-based encryption***

## I. Introduction

Published messages are characterized into classes, without knowledge of what, if any, subscribers there may be. Similarly, subscribers express interest in one or more classes, and only receive messages that are of interest, without knowledge of what, if any, publishers there are.Pub/sub is a sibling of the message queue paradigm, and is typically one part of a larger message-oriented middleware system. Most messaging systems support both the pub/sub and message queue models in their API, e.g. Java Message Service (JMS).This pattern provides greater network scalability and a more dynamic network topology, with a resulting decreased flexibility to modify the Publisher and its structure of the data published.
publishers inject information into the pub/sub system, and subscribers specify the events of interest by means of subscriptions. Published events are routed to their relevant subscribers, without the publishers knowing the relevant set of subscribers, or vice versa. This decoupling is traditionally ensured by intermediate routing over a broker network [10]. In more recent systems, publishers and subscribers organize themselves in a broker-less routing infrastructure, forming an event forwarding overlay [14]. Content-based pub/sub is the variant that provides the most expressive subscription model, where subscriptions define restrictions on the message content. Its expressiveness and asynchronous nature is particularly useful for large-scale distributed applications such as news distribution, stock exchange, environmental monitoring, traffic control, and public sensing. Not surprisingly, pub/sub needs to provide supportive mechanisms to fulfill the basic security demands of these applications such as access control and confidentiality. Access control in the context of pub/sub system means that only authenticated publishers are allowed to disseminate events in the network and only those events are delivered to authorized subscribers. Moreover, the content of events should not be exposed to the routing infrastructure and a subscriber should receive all relevant events without revealing its subscription to the system. Solving these security issues in a content-based pub/sub system imposes new challenges. For instance, end-to-end authentication using a public key infrastructure (PKI) conflicts with the loose coupling between publishers and subscribers, a key requirement for building scalable pub/sub systems. For PKI, publishers must maintain the public keys of all interested subscribers to encrypt events. Subscribers must know the public keys of all relevant publishers to verify the authenticity of the received events. Furthermore, traditional mechanisms to provide confidentiality by encrypting the whole event message conflict with the content-based routing paradigm. Hence, new mechanisms are needed to route encrypted events to subscribers without knowing their subscriptions and to allow subscribers and publishers authenticate each other without knowing each other. In the past, most research

has focused only on providing expressive and scalable pub/sub systems, but little attention has been paid for the need of security. Existing approaches toward secure pub/sub systems mostly rely on the presence of a traditional broker network [20], [2], [9], [12], [7], [18], [16]. These either address security under restricted expressiveness, for example, by using only keyword matching for routing events [12], [20] or rely on a network of (semi-)trusted brokers [19], [17], [12]. Furthermore, existing approaches use coarse-grain epoch based key management and cannot provide fine-grain access control in a scalable manner [12], [20]. Nevertheless, security in broker-less pub/sub systems, where the subscribers are clustered according to their subscriptions, has not been discussed yet in the literature. Building on our results of [13], this paper presents a new approach to provide authentication and confidentiality in a broker-less pub/sub system. Our approach allow subscribers to maintain credentials according to their subscriptions. Private keys assigned to the subscribers are labeled with the credentials. A publisher associates each encrypted event with a set of credentials. We adapted identity-based encryption (IBE) mechanisms [4], [8] 1) to ensure that a particular subscriber can decrypt an event only if there is a match between the credentials associated with the event and the key; and 2) to allow subscribers to verify the authenticity of received events. Furthermore, we address the issue of subscription confidentiality in the presence of semantic clustering of subscribers. A weaker notion of subscription confidentiality is defined and a secure overlay maintenance protocol is designed to preserve the weak subscription confidentiality. In addition to [3], we also present 1) extensions of the cryptographic methods to provide efficient routing of encrypted events by using the idea of searchable encryption, 2) "Multicredential routing" a new event dissemination strategy which strengthens the weak subscription confidentiality, and 3) a thorough analysis of different attacks on subscription confidentiality.

## II .PROBLEM STATEMENT

The provisioning of basic security mechanisms such as authentication and confidentiality is highly challenging in a contentbased publish/subscribe system. Authentication of publishers and subscribers is difficult to achieve due to the loose coupling of publishers and subscribers. Likewise, confidentiality of events and subscriptions conflicts with content-based routing. This paper presents a novel approach to provide confidentiality and authentication in a broker-less content-based publish/subscribe system. The authentication of publishers and subscribers as well as confidentiality of events is ensured, by adapting the pairing-based cryptography mechanisms, to the needs of a publish/subscribe system. Furthermore, an algorithm to cluster subscribers according to their subscriptions preserves a weak notion of subscription confidentiality.

## III .RELATED WORK

**Content-Based Publish/Subscribe:**

For the routing of events from publishers to the relevant subscribers, we use the content-based data model. The event space, denoted by , is composed of a global ordered set of d distinct attributes ($A_i$): f($A_1$, $A_2$, ... $A_d$). Each attribute Ai is characterized by a unique name, its data type, and its domain. The data type can be any ordered type such as integer, floating point, and character strings. The domain describes the range of possible attribute values. The operator O typically includes equality and range operations for numeric attributes and prefix/suffix operations for strings. An event consists of attributes and associated values. An event is matched against a subscription f if the values of attributes in the event satisfy the corresponding constraints imposed by the subscription. We consider pub/sub in a setting where there exits no dedicated broker infrastructure. Publishers and subscribers contribute as peers to the maintenance of a self-organizing overlay structure. To authenticate publishers, we use the concept of advertisements in which a publisher announces beforehand the set of events.

**Attacker Model**

Our attacker model is similar to the commonly used honestbut-curious model [12], [11]. There are two entities in the system: publishers and subscribers. Both the entities are computationally bounded and do not trust each other. Moreover, all the peers (publishers or subscribers) participating in the pub/sub overlay network are honest and do not deviate from the designed protocol. Likewise, authorized publishers only disseminate valid events in the system. However, malicious publishers may masquerade the authorized publishers and spam the overlay network with fake and duplicate events. We do not intend to solve the digital copyright problem; therefore, authorized subscribers do not reveal the content of successfully decrypted events to other subscribers. Subscribers are, however, curious to discover the subscriptions of other subscribers and published events to which they are not authorized to subscribe. Similarly, curious publishers may be interested to read events published in the system. Furthermore, passive attackers outside the pub/sub overlay network can eavesdrop the communication and try to discover content of events and subscriptions. Finally, we assume the presence of secure channels for the distribution of keys from the key server to the publishers and subscribers. A secure channel can be easily realized by using transport layer mechanisms such as Transport Layer Security (TLS) or Secure Socket Layer (SSL).

There are three major goals for the proposed secure pub/ sub system, namely to support authentication, confidentiality, and scalability. Authentication. To avoid noneligible publications, only authorized publishers should be able to publish events in the system. Similarly, subscribers should only receive those messages to which they are authorized to subscribe. Confidentiality. In a broker-less environment, two aspects of confidentiality are of interest: 1) the events are only visible to authorized subscribers and are protected from illegal modifications, and 2) the subscriptions of subscribers are confidential and unforgeable. Scalability. The secure pub/sub system should scale with the number of subscribers in the system. Three aspects are important to preserve scalability: 1) the number of keys to be managed and the cost of subscription should be independent of the number of subscribers in the system, 2) the key server and subscribers should maintain small and constant numbers of keys per subscription, and 3) the overhead because of rekeying should be minimized without compromising the fine-grained access control.

Identity-Based Encryption

While a traditional PKI infrastructure requires to maintain for each publisher or subscriber a private/public key pair which has to be known between communicating entities to encrypt and decrypt messages, identity-based encryption [6] provides a promising alternative to reduce the amount of keys to be managed. In identity-based encryption, any valid string which uniquely identifies a user can be the public key of the user. A key server maintains a single pair of public

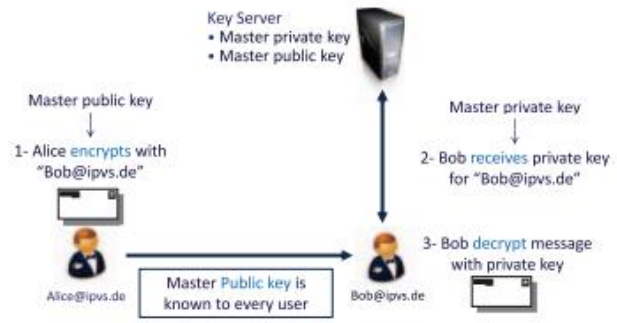


Fig. 1. Identity-based encryption.

and private master keys. The master public key can be used by the sender to encrypt and send the messages to a user with any identity, for example, an e-mail address. To successfully decrypt the message, a receiver needs to obtain a private key for its identity from the key server. Fig. 1 shows the basic idea of using identity-based encryption. We want to stress here that although identity-based encryption at the first glance appears like a highly centralized solution, its properties are ideal for highly distributed applications. A sender needs to know only a single master public key to communicate with any identity. Similarly, a receiver only obtains private keys for its own identities. Furthermore, an instance of central key server can be easily replicated within the network. Finally, a key server maintains only a single pair of master keys and, therefore, can be realized as a smart card, provided to each participant of the system. Although identity-based encryption has been proposed some time ago, only recently pairing-based cryptography (PBC) has laid the foundation of practical implementation of identity-based encryption. Pairing-based cryptography establishes a mapping between two cryptographic groups by means of bilinear maps. This allows the reduction of one problem in one group to a different usually easier problem in another group. We utilize bilinear maps for establishing the basic security mechanisms in the pub/sub system and, therefore, introduce here the main properties.

In our approach, publishers and subscribers interact with a key server. They provide credentials to the key server and in turn receive keys which fit the expressed capabilities in the credentials. Subsequently, those keys can be used to encrypt, decrypt, and sign relevant messages in the contentbased pub/sub system, i.e., the credential becomes authorized by the key server. A credential consists of two parts: 1) a binary string which describes the capability of a peer in publishing and receiving events, and 2) a proof of its identity. The latter is used for authentication against the key server and verification whether the capabilities match the identity of the peer. While this can happen in a variety of ways, for example, relying on challenge response, hardware support, and so on, we pay attention mainly at expressing the capabilities of a credential, i.e., how subscribers and publishers can create a credential. This process needs to account for the many possibilities to partition the set of events expressed by an advertisement or subscription and exploits overlaps in subscriptions and publications. Subsequently, we use the term credential only for referring to the capability string of a credential. The keys assigned to publishers and subscribers, and the ciphertexts, are labeled with credentials. In particular, the identity-based encryption ensures that a particular key can decrypt a particular ciphertext only if there is a match between the credentials of the ciphertext and the key. Publishers and subscribers maintain separate private keys for each authorized credential. The public keys are generated by a string concatenation of a credential, an epoch for key revocation, a symbol distinguishing publishers from subscribers, and some additional parameters. The public keys can be easily generated by any peer without contacting the key server or other peers in the system. Similarly, encryption of events and their verification using public keys do not require any interaction. Due to the loose coupling between publishers and subscribers, a publisher does not know the set of relevant subscribers in the system. Therefore, a published event is encrypted with the public key of all possible credentials, which authorizes a subscriber to successfully decrypt the event. The ciphertexts of the encrypted event

are then signed with the private key of the publisher, as shown in Fig. 2. The overlay network is maintained according to the containment relationship between the subscriptions. Subscribers with coarser subscriptions are placed near the root and forward events to the subscribers with less coarser subscriptions. To maintain such a topology, each subscriber should know the subscription of its parent and child peers. When a new subscriber arrives, it sends the connection request (CR) along with its subscription to a random peer in the overlay network. The connection request is forwarded by possibly many peers in the overlay network before it reaches the right peer to connect. Each forwarding peer matches the subscription in the request with the subscription of its parent and child peers to decide the forwarding direction. Maintaining a relationship between subscriptions clearly contradicts subscription confidentiality. Therefore, we show the approach to ensure a weaker notion of subscription confidentiality.

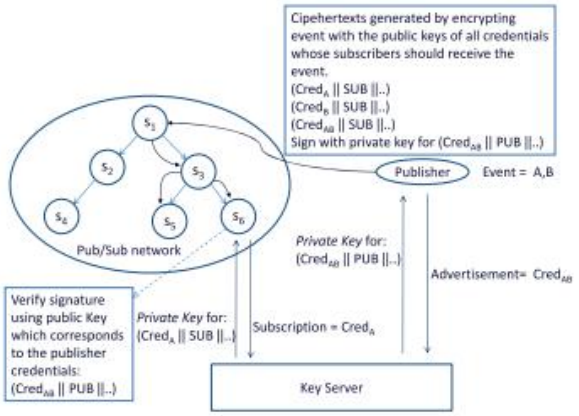


Fig. 2. Approach overview: Publisher has credentials to publish events with two attributes A and B. Subscriber $s_6$ has credentials to receive events with attribute A.

Numeric Attributes

The event space, composed of d distinct numeric attributes, can be geometrically modeled as a d-dimensional space such that each attribute represents a dimension in the space. With the spatial indexing approach, the event space is hierarchically decomposed into regular subspaces, which serve as enclosing approximation for the subscriptions, advertisements, and events. The decomposition procedure divides the domain of one dimension after the other and recursively starts over in the created subspaces. Fig. 3 visualizes the advancing decomposition with the aid of a binary tree. Subspaces are identified by a bit string of "0" and "1"s. A subspace represented by dz1 is covered by the subspace represented by dz2, if dz2 is a prefix of dz1. Subscription or advertisement of a peer can be composed of several subspaces. A credential is assigned for each of the mapped subspace. For instance, in Fig. 3, f2 is mapped to two subspaces and therefore possesses two credentials f000; 010g. An event can be approximated by the smallest

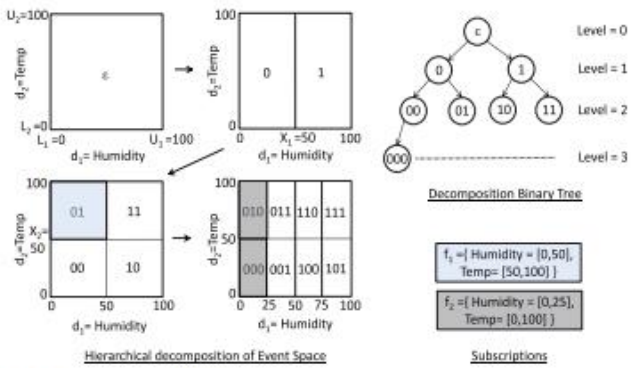


Fig. 3. Numeric attribute.

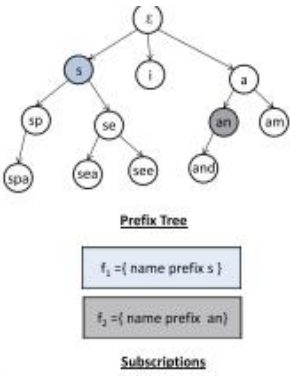


Fig. 4. Prefix matching.

(finest granularity) subspace that encloses the point represented by it. To deliver the encrypted event, a ciphertext must be generated for each subspace that encloses the event so that the peer whose subscription mapped to any of these subspaces should be able to successfully decrypt the event. For example, an event 0010 is enclosed by the five subspaces 0010, 001, 00, 0, and . For an event space with a large set of numeric attributes, the number of mapped subspaces and, therefore, credentials for a subscription can be very large. This affects the scalability of the system. We address this problem by decomposing the domain of each attribute into subspaces separately. The spatial indexing procedure is the same as above; however, in this case, a separate decomposition tree is built for each attribute. Each peer receives credentials separately for each attribute in its subscription. The number of credentials maintained for each subscription or advertisement is bounded.

## IV. Conclusion

The approach is highly scalable in terms of number of subscribers and publishers in the system and the number of keys maintained by them. In particular, we have developed mechanisms to assign credentials to publishers

and subscribers according to their subscriptions and advertisements.In this paper, we have presented a new approach to provide authentication and confidentiality in a broker-less contentbased pub/sub system. Private keys assigned to publishers and subscribers, and the ciphertexts are labeled with credentials. We adapted techniques from identitybased encryption 1) to ensure that a particular subscriber can decrypt an event only if there is a match between the credentials associated with the event and its private keys and 2) to allow subscribers to verify the authenticity of received events. Furthermore, we developed a secure overlay maintenance protocol and proposed two event dissemination strategies to preserve the weak subscription confidentiality in the presence of semantic clustering of subscribers. The evaluations demonstrate the viability of the proposed security mechanisms and analyze attacks on subscription confidentiality.

## REFERENCES

[1] E. Anceaume, M. Gradinariu, A.K. Datta, G. Simon, and A. Virgillito, "A Semantic Overlay for Self- Peer-to-Peer Publish/ Subscribe," Proc. 26th IEEE Int'l Conf. Distributed Computing Systems (ICDCS), 2006.

[2] J. Bacon, D.M. Eyers, J. Singh, and P.R. Pietzuch, "Access Control in Publish/Subscribe Systems," Proc. Second ACM Int'l Conf. Distributed Event-Based Systems (DEBS), 2008.

[3] W.C. Barker and E.B. Barker, "SP 800-67 Rev. 1. Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher," technical report, Nat'l Inst. of Standards & Technology, 2012.

[4] J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext-Policy Attribute-Based Encryption," Proc. IEEE Symp. Security and Privacy, 2007.

[5] D. Boneh, G.D. Crescenzo, R. Ostrovsky, and G. Persiano, "Public Key Encryption with Keyword Search," Proc. Int'l Conf. Theory and Applications of Cryptographic Techniques on Advances in Cryptology (EUROCRYPT), 2004.

[6] D. Boneh and M.K. Franklin, "Identity-Based Encryption from the Weil Pairing," Proc. Int'l Cryptology Conf. Advances in Cryptology, 2001.

[7] S. Choi, G. Ghinita, and E. Bertino, "A Privacy-Enhancing Content-Based Publish/Subscribe System Using Scalar Product Preserving Transformations," Proc. 21st Int'l Conf. Database and Expert Systems Applications: Part I, 2010.

[8] V. Goyal, O. Pandey, A. Sahai, and B. Waters, "Attribute-Based Encryption for Fine-Grained Access Control of Encrypted Data," Proc. ACM 13th Conf. Computer and Comm. Security (CCS), 2006.

[9] M. Ion, G. Russello, and B. Crispo, "Supporting Publication and Subscription Confidentiality in Pub/Sub Networks," Proc. Sixth Int'l ICST Conf. Security and Privacy in Comm. Networks (SecureComm), 2010.

[10] H.-A. Jacobsen, A.K.Y. Cheung, G. Li, B. Maniymaran, V. Muthusamy, and R.S. Kazemzadeh, "The PADRES Publish/ Subscribe System," Principles and Applications of Distributed Event-Based Systems. IGI Global, 2010.

[11] M. Jelasity, A. Montresor, G.P. Jesi, and S. Voulgaris, "PeerSim: A Peer-to-Peer Simulator," http://peersim.sourceforge.net/, 2013.

[12] H. Khurana, "Scalable Security and Accounting Services for Content-Based Publish/Subscribe Systems," Proc. ACM Symp. Applied Computing, 2005.

[13] A. Lewko, A. Sahai, and B. Waters, "Revocation Systems with Very Small Private Keys," Proc. IEEE Symp. Security and Privacy, 2010.

[14] B. Lynn, "The Pairing-Based Cryptography (PBC) Library," http://crypto.stanford.edu/pbc/, 2010.

[15] F.P. Miller, A.F. Vandome, and J. McBrewster, Advanced Encryption Standard. Alpha Press, 2009.

[16] M. Nabeel, N. Shang, and E. Bertino, "Efficient Privacy Preserving Content Based Publish Subscribe Systems," Proc. 17th ACM Symp. Access Control Models and Technologies, 2012.

[17] L. Opyrchal and A. Prakash, "Secure Distribution of Events in Content-Based Publish Subscribe Systems," Proc. 10th Conf. USENIX Security Symp., 2001.

[18] L.I.W. Pesonen, D.M. Eyers, and J. Bacon, "Encryption-Enforced Access Control in Dynamic Multi-Domain Publish/Subscribe Networks," Proc. ACM Int'l Conf. Distributed Event-Based Systems (DEBS), 2007.

[19] P. Pietzuch, "Hermes: A Scalable Event-Based Middleware," PhD dissertation, Univ. of Cambridge, Feb. 2004.

[20] C. Raiciu and D.S. Rosenblum, "Enabling Confidentiality in Content-Based Publish/Subscribe Infrastructures," Proc. IEEE Second CreatNet Int'l Conf. Security and Privacy in Comm. Networks (SecureComm), 2006.