

Attribute-Based Encryption with Break-Glass

Achim D. Brucker¹, Helmut Petritsch¹, and Stefan G. Weber²

¹ SAP Research, Vincenz-Priessnitz-Str. 1, 76131 Karlsruhe, Germany
{achim.brucker,helmut.petritsch}@sap.com

² Telecooperation Group, Technische Universität Darmstadt, Hochschulstrasse 10,
64289 Darmstadt, Germany
sweber@tk.informatik.tu-darmstadt.de

Abstract. *Attribute-based Encryption* (ABE) allows for implementing fine-grained decentralized access control based on properties or attributes a user has. Thus, there is no need for writing detailed, user-based policies in advance. This makes ABE in particular interesting for implementing security mechanisms in dynamic environments such as ubiquitous computing, disaster management, or health-care. For supporting the latter two application areas, common ABE approaches lack one important feature: break-glass, i. e., the controlled overriding of access control restrictions.

In this paper we present an integration of break-glass into an approach for end-to-end secure information sharing using ABE techniques.

Keywords: Access control, break-glass, attribute-based encryption, disaster management, trusted computing platform.

1 Introduction

The distribution and availability of digital information in every-day's life and work brings along new opportunities for providing situation-dependent support to the individual. This ubiquity of information also creates new challenges in protecting both the provided information and the privacy of its users. Access control mechanism should allow the fine-grained specification of access control policies and protect the users privacy. Additionally, in several application domains, mechanisms for the controlled override of access restrictions need to be supported.

Recently introduced, Attribute-based Encryption (ABE) techniques [2, 13, 14] have drawn attention for realizing decentralized access control in large and dynamic networks and ubiquitous computing environments [8, 15, 16]. By allowing to implement an efficient targeted broadcast encryption, ABE provides means for designing flexible and scalable access control systems, without the need of administrating large access control policies. Moreover, this approach allows for a decentralized enforcement of access control policies by cryptographically binding the policies to data objects.

This makes ABE especially attractive in dynamic environments such as grid-computing, disaster management, or the health-care area. Albeit, for the latter

two areas, ABE creates particular challenges for providing one important feature: *break-glass*, i. e., the traceable, ad-hoc override of access control policies. Traditionally, break-glass access control systems are implemented in systems using user-based access control policies and (centralized) policy decisions points.

In this paper, we present an integration of fine-grained break-glass concepts into a system for end-to-end secure information sharing based on ABE. In more detail, our contributions are an ABE security infrastructure for ubiquitous data sharing applications supporting: 1. an ABE scheme with multi-leveled break-glass access control and 2. secure logging that allows for analyzing the actions of users done during break-glass access rights in particular.

The rest of the paper is structured as follows: after introducing the scenario and requirements in Sect. 2, we present preliminaries of our work in Sect. 3. In Sect. 4, we detail concepts building the foundation for the implementing architecture presented in Sect. 5. Sect. 6 and Sect. 7 conclude our work.

2 Scenario and Requirements

Our integration of break-glass into ABE is motivated by the rise of distributed disaster management information systems (DMISs), i. e., systems supporting crises management teams. For the operational headquarters, we assume that each of them is supported by a DMIS providing support for maintaining the current situation, planing, and simulation. Moreover, DMISs provide means for the efficient and secure communication within an operational headquarter and to the outside. The latter includes the communication between several operational headquarters (e. g., see [3]) and the communication with the forces in the field using mobile devices. For the field forces, we assume that they are equipped with mobile communication devices, i. e., smart-phones that provide a digital communication channel between the field forces and the operational headquarters. In the following, we focus on broadcast-based communication between the operational headquarters and the field forces. In particular, we do not assume a communication infrastructure providing confidential point-to-point channels. Instead, confidentiality of broadcasted messages is guaranteed by means of ABE, implementing an application level end-to-end encryption.

Let us assume that in an emergency situation, a firefighter decides to override the access control that prevents him from reading the communication between the police officers and the police headquarter. Of course, such overrides need to be made on case by case basis. Thus, the firefighter needs to be able to decide if a certain message addressed to the police is valuable for him or not. Moreover, overrides needs to be logged for post-hoc audits. To allow the efficient and traceable break-glass access, we require that the mobile devices can store a sensible amount of messages (e. g., all messages sent in the last 30 minutes), messages are classified in a format readable even if the message is, due to access restrictions, not readable, and overrides of access control restrictions and all action taken on information obtained by overriding access control restrictions are logged immutably.

3 Technical Background

3.1 Attribute-Based Encryption

Attribute-based Encryption (ABE) [2, 14] is a public key cryptography primitive generalizing Identity-based Encryption (IBE) [7]: while in IBE, a single receiver of a confidential message is described by a single string associated with his identity, in Attribute-based Encryption (ABE) a group of receivers is described by a combination of several descriptive attributes, which is also called an attribute policy. Therefore, each user is associated with a set of person-related attributes such as job description and status, and receives related private keys analogues as well. Especially, in the cryptographic operations, attributes directly map to a users' credential, i. e., the attributes present components of the private key.

In current ABE implementations, the policy itself is cryptographically bound to the data object, but can still be read by any receiver [13]. Moreover, using ABE, a sender does not know the public keys of the recipients. Thus, messages can be encrypted before any receiver must own the required decryption keys. Albeit, a trusted third party, called attribute authority, must be able to produce any possible private key.

In [15], we show that ABE allows for building a security architecture for secure data sharing, also suitable for emergency management applications, i. e., it can be used to implement an end-to-end encryption allowing operational headquarters to target dynamic groups of first responders. Albeit, the so far proposed concept lacks support for a controlled override of access restrictions in emergency cases.

3.2 Break-Glass

Introduced in [1], *break-glass* refers to quick means for extending a person's access rights in exceptional cases. Of course, the usage of exceptional access rights needs to be documented for later audits and reviews. Usually, break-glass solutions are based on authenticating the user and, therefore, are not directly applicable to ABE-based access control system.

Based on our break-glass approach presented in [4], we assume an access control policy p based on an access control model \mathcal{A} . A policy maps access control relevant information, e. g., subjects, resources, and actions on resources to an access control decision. A policy p *refines* a policy p' (written $p \sqsubseteq p'$) if and only if p is at least as restrictive as p' . We write p^\top for the policy that allows all actions and p^\perp for the policy that allows no action. The relation $_ \sqsubseteq _$ defines a partial order on a set of policies. Consequently, $(P_{\mathcal{A}}, \sqsubseteq, p^\perp, p^\top)$ is a lattice, where $P_{\mathcal{A}}$ be the set of all policies of the access control model \mathcal{A} . During normal operations, the *regular policy* p^{reg} is put into place; we call the set $L_{\mathcal{A}} = \{p \mid p \in P_{\mathcal{A}} \wedge p^{\text{reg}} \sqsubseteq p \wedge p \neq p^{\text{reg}}\}$ of policies that are refined by the regular policy, *emergency policies* of the policy p^{reg} and require $(P_{\mathcal{A}} \setminus p^\perp, \sqsubseteq, p^{\text{reg}}, p^\top)$ to be a lattice. At runtime, policies can be *active* or *inactive*, whereas only active emergency policies, denoted as $L_{\mathcal{A}}^{\text{act}} \subseteq L_{\mathcal{A}}$, contribute to the access control decision. An access evaluating to “deny” regarding the regular policy, but granted by an active emergency policy $p \in L_{\mathcal{A}}$ is called *override access*.

4 Integrating Break-Glass Access Control and ABE

For encoding a lattice of (emergency) policies in ABE, we use a hierarchy of *emergency attributes* ($\mathbf{a} \in \mathfrak{A}$). While emergency attributes are technically similar to regular attributes, their usage implies further obligations: first, every use of an emergency attribute requires user confirmation and, second, all access based on the use of emergency attributes needs to be logged for later audit. To each emergency policy p^j , we assign an emergency attribute \mathbf{a}^i representing the length of the shortest path (with respect to \sqsubseteq) from the regular policy p^{reg} to p^j . The length of the shortest override path provides a measurement of the severity of an override, i. e., an access requiring a “high” emergency attribute is more critical than one requiring a “low” emergency attributes. Of course, if several override attributes allow for accessing a specific resource, the lowest one (i. e., the one with the shortest path) should be chosen.

We assume that the different emergency organization agreed on using a common subset (i. e., a sub-lattice) of emergency attributes. Still, this does not result in a sharing of emergency attributes, e. g., let \mathfrak{A}_F (\mathfrak{A}_P) be the set of emergency attributes of the fire brigade (police) then $\mathfrak{A}_F \cap \mathfrak{A}_P = \emptyset$ holds. We denote the emergency attribute \mathbf{a}^i representing polices that have a minimal distance of i to the regular policy with $\mathbf{a}_F^1 \in \mathfrak{A}_F$ ($\mathbf{a}_P^1 \in \mathfrak{A}_P$) for the fire brigade (police).

Let us consider a situation in which firefighters should be able to access the status messages of the police by overriding the regular access control. In more detail, status updates should be accessible to

1. police officers-in-charge (i. e., with regular access): $p_P^{\text{reg}} = (\text{police}_P \wedge \text{officer-in-charge}_P)$
2. every police member under a low emergency attribute: $p_P^1 = (\text{police}_P)$,
3. every officer-in-charge of the fire brigade under a low emergency attribute: $p_F^1 = (\text{firebrigade}_F \wedge \text{officer-in-charge}_F)$, and
4. every member of the fire brigade under a high emergency attribute: $p_F^2 = (\text{firebrigade}_F)$.




Fig. 1. A simple emergency attribute hierarchy

Fig. 1 illustrates the resulting hierarchy of emergency attributes.

To support this scenario we derive the following ABE break-glass policy:

$$\underbrace{(p_P^{\text{reg}}) \overline{\vee} (\mathbf{a}_P^1 \wedge p_P^1)}_{\text{police}} \overline{\vee} \underbrace{(\mathbf{a}_F^1 \wedge p_F^1) \overline{\vee} (\mathbf{a}_F^2 \wedge p_F^2)}_{\text{fire brigade}}$$

Notably, the first part of the policy can only be decrypted by the police and the second part can only be decrypted by the fire brigade (to avoid a global key escrow, we require separate attribute authorities for each organization). Here, $_ \overline{\vee} _$ denotes a disjunction with lazy (also called short-circuit) evaluation semantics, i. e., if the left hand side already grants access, the right hand side is

ignored. This semantics allows that the regular policy has a higher precedence than any other policy and that policies under a lower emergency attribute have a higher precedence than policies under a higher emergency attribute.

The information exchange over organizational boundaries (i. e., between separate attribute authorities) requires an understanding of the semantics of the attributes between organization. Of course, only the set of attributes that should be used in shared policies needs to be exchanged. Within the DMIS foreign attributes can be grouped into virtual attributes (similar to e-mail aliases); this allows for maintaining an easy to use interface while, internally, mapping different organizational structures.

For supporting the activation and deactivation of emergency attributes by a central authority, there are mainly two different approaches: 1. the status of all emergency attributes is broadcasted to and enforced by clients, i. e., messages are encrypted with all emergency attributes. 2. only emergency attributes active during encryption are used. The first approach requires a high trustworthiness of the clients, e. g., by manipulating status updates a device could be forced to continue to use a high attribute. The second approach cannot prohibit the usage of an inactive emergency attribute, if this emergency attribute has been active during encryption. Thus, we propose a combination of these two approaches: i. e., decryption is only possible if the required emergency attribute is active during both encryption and decryption. First, this gives the information provider maximal control over his data. Second, for access over organizational boundaries, the information receiving organization can restrict the use of foreign information.

Messages sent to clients are self-contained, i. e., the user has to decide on a per message basis, if a message is relevant or not. For this decision the user requires some information about the information contained in the message. Such a classification may be sensitive itself (e. g., it may allow to infer secrets from the frequency of some type of message), so it has to be protected too. Thus, a message which allows override access consists of two parts: the *inner part* containing the message, encrypted with the regular policy and the emergency policies and the *outer part* containing some classification information about the message, whereas this classification can be protected by another policy. The outer policy must be less restrictive than the inner policy, i. e., users permitted to decrypt the inner part must be permitted to decrypt the outer part.

5 Security Infrastructure

Our approach for supporting ABE with break-glass on mobile devices requires a trusted software stack, which provides secure storage and logging functionalities. Such a trusted software stack can be provided by using a standard hardware component called *Trusted Platform Module* (TPM) (see Fig. 2). In particular, we rely on the *root of trust for storage* (RTS) and the *root of trust for reporting* (RTR), which both can be provided by a TPM.

In a registration phase, every member of the field forces receives a personal device containing the private keys representing the members attributes. After activation of such a personal device, e.g., by entering a PIN, the device allows for decrypting messages under the regular policy and the active emergency levels. Received messages are stored in their encrypted form, which allows to store them on a non-secured area.

The trusted base of the mobile device provides a secure storage compartment inside the TPM. However, the TPM can also protect larger amounts of stored data, i.e., the private keys for the attributes of the user, the status of current active emergency attributes, and the logs of break-glass accesses. Hereby, only the keys used for storage encryption remain in the TPM, while encrypted *key blobs* are stored outside the TPM. The TPM allows to provide a secure software stack, i.e., an isolated computing environment which runs on the mobile device. The module verifies that only certified applications access to certain, e.g., security critical, areas of the mobile device. As such, it guards the execution of critical steps, e.g., all steps relating to keys that shall only be available under additional obligations. The secure software stack allows to implement secure channels.

In case of a break-glass access, the following functional steps are executed on the device: 1. *Display of Policies*: The classification of the message is displayed to the user. 2. *Selection of Message*: Based on the classification, the user selects a message. Hereby, he confirms the break-glass action on the chosen single message. 3. *Granting Access*: After evaluating the current emergency level, a copy of the required emergency attribute is released to the secure software stack. 4. *Retrieval of regular Keys*: Also, the regularly available attributes are transferred to the secure computing environment. 5. *Decryption and Deletion*: Regular attributes and emergency attributes are used for decrypting the selected message. Afterwards, the copies of the attributes are deleted. 6. *Logging and Display of Content*: The break-glass access is logged, i.e., an entry containing policy, time and emergency level is signed with the private key of the TPM and added to the log. Finally, the content of the decrypted message is displayed to the user. This mechanism guarantees that in case of a break-glass access, first, the user has to confirm the override, second, only active emergency levels are used for a single decryption and, third, the emergency access is logged to a secure memory storage. As mobile devices can be physically manipulated, log entries stored on the device are synchronized with the central storage during online phases or when the operation is finished and the device is returned to the operational center.

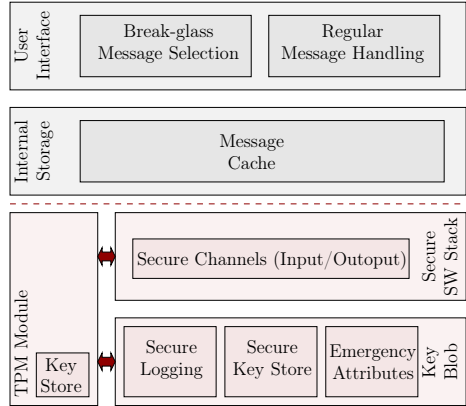


Fig. 2. Overview of the Security Infrastructure

6 Related Work

Our work extends previous work on ABE [2, 13, 14] (in particular CP-ABE [2]) and break-glass access control [1, 4, 5] (in particular [4]). In contrast to [6], our break-glass solution allows for the secure and traceable overruling of access control of mobile devices without the need of an online validation.

Research on disaster management information systems spans from high-level organizational views up to questions of how field forces can be supported by “mobile response devices” [10], whereas questions from collaboration of organizations, visualization and decision-making are examined [12]. It is often scenario-oriented, e. g., Johnson [9] draws conclusions from a flooding in UK in 2007 and points out, among others, a need for a better collaboration of organizations.

Only a few works in the disaster management domain explicitly consider security issues; in [15], we propose an ABE-based security architecture for sharing sensitive data in emergency management applications. Tailored to a different application context, Huang and Verma [8] present an ABE-based framework for establishing secure communications and enforcing access control policies within vehicular ad hoc networks. Yu et al. [16] use ABE for enforcing fine-grained access control policies in wireless sensor networks. Although these works are similar to ours, they lack mechanisms for break-glass. Levin et al. [11] present an architecture for managing emergency information for first responders. Their assumptions about the environment are similar to ours, but their approach requires additional special hardware and does not support multiple emergency levels.

7 Conclusion

Securely providing the right information at the right time is a highly challenging, but important task. Based on practical experiences we introduced a realistic scenario in the emergency management domain. Here, an individual field force may need immediate access to certain information to save lives and property.

Our experiences show that modern cryptographic techniques are important building blocks for realizing security and access control mechanisms in large scale decentralized systems and ubiquitous computing environments. However, to design a secure system, it is important to carefully integrate cryptographic mechanisms with further security concepts. In this paper, we have shown how ABE, which is among others appropriate for flexible, fast changing environments, can be combined with break-glass solutions which provide support for systems, where policies may be overridden in case of an emergency. Such systems are especially relevant for first response teams as fire brigades or for more static applications scenarios like health information on a chip-card, where sensitive information has to be protected but be accessible in emergency situations to save a patients life.

Acknowledgments. This work has been supported by the German “Federal Ministry of Education and Research” in the context of the project “SoKNOS” and by CASED (www.cased.de).

References

- [1] Break-glass: An approach to granting emergency access to healthcare systems. White paper, Joint NEMA/COCIR/JIRA Security and Privacy Committee (2004)
- [2] Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: IEEE Symposium on Security and Privacy, pp. 321–334. IEEE Computer Society Press, Los Alamitos (2007)
- [3] Brucker, A.D., Hutter, D.: Information flow in disaster management systems. In: International Conference on Availability, Reliability and Security (ARES), IEEE Computer Society Press, Los Alamitos (2010)
- [4] Brucker, A.D., Petritsch, H.: Extending access control models with break-glass. In: Carminati, B., Joshi, J. (eds.) ACM symposium on access control models and technologies (SACMAT), pp. 197–206. ACM Press, New York (2009)
- [5] Ferreira, A., Cruz-Correia, R., Antunes, L., Farinha, P., Oliveira-Palhares, E., Chadwick, D., Costa-Pereira, A.: How to break access control in a controlled manner. In: IEEE International Symposium on Computer-Based Medical Systems (CBMS), pp. 847–854 (2006)
- [6] Gardner, R.W., Garera, S., Pagano, M.W., Green, M., Rubin, A.D.: Securing medical records on smart phones. In: ACM workshop on Security and privacy in medical and home-care systems (SPIMACS), pp. 31–40. ACM Press, New York (2009)
- [7] Gentry, C.: IBE (Identity-Based Encryption). In: Handbook of Information Security, vol. 2, pp. 575–592. John Wiley & Sons, Chichester (2006)
- [8] Huang, D., Verma, M.: ASPE: attribute-based secure policy enforcement in vehicular ad hoc networks. *Ad Hoc Networks* 7(8), 1526–1535 (2009)
- [9] Johnson, C.W.: Complexity, structured chaos and the importance of information management for mobile computing in the UK floods of 2007. In: [12], pp. 1–11 (2007)
- [10] Lachner, J., Hellwagner, H.: Information and communication systems for mobile emergency response. In: Kaschek, R., Kop, C., Steinberger, C., Fliedl, G. (eds.) Information Systems and e-Business Technologies (UNISCON). LNBIP, vol. 5, pp. 213–224. Springer, Heidelberg (1974)
- [11] Levin, T.E., Dwoskin, J.S., Bhaskara, G., Nguyen, T.D., Clark, P.C., Lee, R.B., Irvine, C.E., Benzel, T.: Securing the dissemination of emergency response data with an integrated hardware-software architecture. In: Chen, L., Mitchell, C.J., Martin, A. (eds.) Trust 2009. LNCS, vol. 5471, pp. 133–152. Springer, Heidelberg (2009)
- [12] Löffler, J., Klamm, M. (eds.): Mobile Information Technology for Emergency Response (MobileResponse). LNCS, vol. 5424. Springer, Heidelberg (2009)
- [13] Pirretti, M., Traynor, P., McDaniel, P., Waters, B.: Secure attribute-based systems. In: ACM conference on Computer and communications security (CCS), pp. 99–112. ACM Press, New York (2006)
- [14] Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
- [15] Weber, S.G.: Securing first response coordination with dynamic attribute-based encryption. In: World Congress on Privacy, Security, Trust and the Management of e-Business (CONGRESS), pp. 58–69. IEEE Computer Society Press, Los Alamitos (2009)
- [16] Yu, S., Ren, K., Lou, W.: FDAC: Toward fine-grained distributed data access control in wireless sensor networks. In: IEEE Conference on Computer Communications (INFOCOM). IEEE Computer Society Press, Los Alamitos (2009)