Attribute-Based Encryption for Named Data Networking

Aleksandr Lenin
†

Peeter Laud
Cybernetica AS
peeter.laud@cyber.ee

ABSTRACT

We compare and discuss the applicability and trade-offs of different attribute-based encryption (ABE) schemes to the possible use-cases of content-centric networking requiring end-to-end encryption of data with fine-grained access control, where the nature of content producers and consumers may vary, as well as the required expressivity of policies. We also report on the choice and implementation of an ABE scheme, as well as the overheads associated with its use.

CCS CONCEPTS

• **Networks** → *Transport protocols*; • **Security and privacy** → **Security protocols**.

ACM Reference Format:
Aleksandr Lenin and Peeter Laud. 2021. Attribute-Based Encryption for Named Data Networking. In *8th ACM Conference on Information-Centric Networking (ICN '21), September 22–24, 2021, Paris, France.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3460417.3483371

1 INTRODUCTION

Named Data Networking (NDN) is one of five projects funded by the U.S. National Science Foundation under its Future Internet Architecture Program [19]. NDN changes the network layer in the network protocol stack, such that packets name content objects, rather than communication endpoints. This changes the semantics of the network from delivering packets to a given destination to fetching data identified by given names. NDN follows a data-centric security approach, in which the content producer signs all the data packets it generates. This ensures the integrity and authenticity of a data packet. It allows to decouple the consumer's trust from the network node that served the content, and replaces it with the trust towards the producer directly.

Signatures provide the data packets with properties similar to authentic channels in connection-centric networks — known source, and non-tampering. The properties of confidential channels — known sink, and non-observation — are captured by encryption primitives, whose use in content-centric networks requires the solutions to key distribution and management. To grant content access to authorized consumers, their group should be somehow separated out the group of recipients, with the publisher not required to keep a comprehensive list of them.

Encrypting content to a group of recipients, without having a comprehensive list of them, matches the functionality of *attribute-based encryption* (ABE) [6], more specifically *ciphertext-policy* ABE (CP-ABE) [3]. This primitive does away with crisp identities of the recipients, but assigns a set of *attributes* to each entity. ABE allows to encrypt a message for an *access policy* (essentially a boolean formula) over attributes, so that only users holding private keys with attributes matching the access structure can decrypt the message. In an ABE scheme, the users' keys are issued by some trusted party, usually called the key generation center (KGC).

2 CIPHERTEXT POLICY ATTRIBUTE BASED ENCRYPTION

The lifecycle of a CP-ABE instance starts with the KGC generating the *master secret key* (MSK) and the *master public key* (MPK). Using the MSK, the KGC can issue *private keys* corresponding to sets of attributes. Anyone, using the MPK, can encrypt messages for the access structure they've selected; the richness of supported access structures may vary among different CP-ABE schemes. The ciphertexts can be decrypted with private keys having the attributes that satisfy the access structure selected during encryption. Two different private keys cannot be combined to build a "stronger" private key.

Different CP-ABE schemes strike different trade-offs between its functionality and complexity parameters — expressiveness of access policies, length of keys and ciphertexts, complexity of computations. These should be compared for their usability in NDN and in various applications. This comparison and informed choice has been lacking in previous *Encryption-based access control* (EncBAC) [7, 11] proposals making use of ABE [14, 18, 21]. Even though a precise comparison, applicable in all scenarios, may be difficult, we have nevertheless ventured to compare different schemes, in order to select one for our application.

The comparison is simplified by most ABE schemes being instances of pairing-based cryptography [5, 10]. Hence the plaintexts in these schemes are elements of elliptic curve groups, suitable for encoding a key for symmetric encryption (e.g. AES). The sizes of keys and ciphertexts can be measured by the number of group elements it takes to encode them. The computational complexity of operations can be estimated as the number of expensive operations — exponentiations and pairings — it takes to perform them, even though the cost of performing multiple exponentiations does not always grow linearly with their number.

The most variety in CP-ABE schemes is in the policies they support. Some of them support all monotone formulas, others support a more restricted set. Some may also support non-monotone formulas, where a ciphertext can be decrypted only with keys that *do not* have a certain attribute. Some of them have a separate notion of revoking a private key. We have tried to characterize a number of proposed CP-ABE schemes and the policies they support, together

---

ICN '21, September 22–24, 2021, Paris, France

ACM ISBN 978-1-4503-8460-5/21/09…$15.00
https://doi.org/10.1145/3460417.3483371

Table 1: comparison of CP-ABE schemes (n: number of attributes; m: number of users; s: number of (non-negated) attributes in a key or policy; t: number of negated or revoked attributes in a key or policy; u and v: number of leaves and non-leaves in the access tree; E: exponentiation; P: pairing; size: number of group elements)

Reference	supported policies	private key size	ciphertext size	encryption cost	decryption cost
[3]	tree of threshold gates	$2s + 1$	$2u + 1$	$2(u + 1)E$	$2uP + vE$
[4]	Conjunction of attributes and their negations	2	3	$2E$	$3E + 2P$
[22]	Conjunction of attributes and their negations	$2n + 1$	2	$2E$	$1P$
[13]	Conjunction of attributes and their negations	$s + 2$	$t + 3$	0^*	$3P$
[17]	Formula in negation normal form	$4s + 2$	$3(s + t) + 2$	$(5(s + t) + 2)E + 1P$	$(3s + t(2(s + t) + 1))P$
[15]	boolean formula	$4s + 2$	$2s + 3t + 2$	$(2s + 3t + 2)E$	$2sP + (3v + 1)E$
[12]	AND-OR formula	$n + m + 1$	$3 + 2s$	$1P + 3(s + 1)E$	$1P$
[20]	(Private) conjunction of attributes	$m + 2$	4	$1P + 4E$	$4P$
[1]	monotone formula**	$3 + n$	$3 + s + t$	$(3 + s + t)E + 2P$	$4P + 2E$
[9]	boolean circuit	$2n + s + t$	3	$1P + (2 + s + t)E$	$2P + (s + t)E$

* Only multiplications of group elements are needed ** has the notion of *users* and their *revocation*

with the computation and storage/communication costs of using them; the comparison is presented in Table. 1. It is missing certain "one-time" costs — computation involved in the generation of MSK and MPK, and the private keys, as well as the sizes of MSK and MPK. As these costs are incurred significantly less often than the costs we consider, we believe it is fair to ignore them in our comparison.

Even though a scheme may "naturally" support only a restricted class of policies, e.g. a conjunction of attributes, it is still possible to encode more complex policies, although this may incur a significant computational cost. For example, support for disjunction in policies can be provided even if the scheme does not "naturally" have it — one may encrypt the same symmetric key several times, under different policies.

More significantly, instead of negating an attribute in the policy, we could introduce two different attributes — the "positive" and the "negative". In this way, monotone policies would be sufficient. However, each key would have a lot of attributes associated to it. Also, instead of having a separate notion of revoking a private key, we could introduce a separate attribute for each user of the system. Revoking that user would mean negating the corresponding attribute in future ciphertexts. But this will increase the total number of attributes in the system. The actual cost of such steps can be found from Table 1.

A system supporting the revocation of users would have a *revocation authority*), possibly equal to the KGC, who is trusted to make statements about the revocation status of users. This authority periodically publishes (under a predefined name) the list of revoked users, and signs it. A content producer obtains that list by posting an interest for this name, verifies the signature, and adds the negated attributes of revoked users to the policy of the ciphertext of any content it subsequently produces.

3 OUR APPLICATION

In our application, data is produced by a variety of devices, some of which are heavily resource-constrained. The decryption of data takes place in smartphones and more powerful devices. Hence we are looking for a scheme where the encryption cost and ciphertext size are small.

Our policies are mostly just conjunctions of attributes. However, we need to be able to revoke the decryption ability of individual devices that have been assigned private keys. The number of revoked devices is assumed to not grow large.

Hence we have chosen the Lubicz-Sirvent scheme [13] for our application. In addition to "normal" attributes, we will also introduce a separate attribute for each decryption device. The central authority distributes the decryption keys, as well as the lists of attributes of revoked devices. When encrypting, the data producers add the negations of revoked attributes to the policy.

We have implemented the scheme [13] and measured its overheads (compared to no confidentiality protection) in certain scenarios. In particular, we timed the overhead of encryption when downloading files of different sizes over NDN. Here the overhead consisted of downloading the ABE-encrypted symmetric key protecting the actual file, performing the ABE-decryption to obtain the symmetric key, and then decrypting the encrypted file. On a network consisting of a single node, with $1 - -10$ users, simultaneously downloading a $50 - -500$ MiB file decryptable for them all, we saw timing overheads in the range of $50 - -100\%$.

The scheme [13] was originally presented as using a *symmetric* pairing scheme. Due to the advances in cryptanalysis of pairing-based cryptography [2, 8], we have adapted their scheme to use asymmetric pairings. We have used the PARI library [16] for algebraic computations. The source code of the implementation will be made public.

Acknowledgment. This research has been funded by the Defense Advanced Research Projects Agency (DARPA) under contract HR0011-17-C-0111, as well as by European Regional Development Fund through the ICT Centre of Excellence "EXCITE". The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

[1] Asim, M., Ibraimi, L., and Petkovic, M. Ciphertext-policy attribute-based broadcast encryption scheme. In *Communications and Multimedia Security, 12th IFIP TC 6 / TC 11 International Conference, CMS 2011, Ghent, Belgium, October 19-21,2011. Proceedings* (2011), B. D. Decker, J. Lapon, V. Naessens, and A. Uhl, Eds., vol. 7025 of *Lecture Notes in Computer Science*, Springer, pp. 244–246.

[2] Barbulescu, R., Gaudry, P., Joux, A., and Thomé, E. A heuristic quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic.

In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings* (2014), P. Q. Nguyen and E. Oswald, Eds., vol. 8441 of *Lecture Notes in Computer Science*, Springer, pp. 1–16.

[3] Bethencourt, J., Sahai, A., and Waters, B. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy (SP '07)* (May 2007), pp. 321–334.

[4] Emura, K., Miyaji, A., Nomura, A., Omote, K., and Soshi, M. A ciphertext-policy attribute-based encryption scheme with constant ciphertext length. In *Information Security Practice and Experience: 5th International Conference, ISPEC 2009 Xi'an, China, April 13-15, 2009 Proceedings* (Berlin, Heidelberg, 2009), F. Bao, H. Li, and G. Wang, Eds., Springer Berlin Heidelberg, pp. 13–23.

[5] Galbraith, S. D., Paterson, K. G., and Smart, N. P. Pairings for cryptographers. *Discret. Appl. Math. 156*, 16 (2008), 3113–3121.

[6] Goyal, V., Pandey, O., Sahai, A., and Waters, B. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, October 30 - November 3, 2006* (2006), A. Juels, R. N. Wright, and S. D. C. di Vimercati, Eds., ACM, pp. 89–98.

[7] Hamdane, B., Msahli, M., Serhrouchni, A., and Fatmi, S. G. E. Data-based access control in named data networking. In *9th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing* (Oct 2013), pp. 531–536.

[8] Joux, A. A new index calculus algorithm with complexity $L(1/4 + o(1))$ in small characteristic. In *Selected Areas in Cryptography - SAC 2013 - 20th International Conference, Burnaby, BC, Canada, August 14-16, 2013, Revised Selected Papers* (2013), T. Lange, K. E. Lauter, and P. Lisonek, Eds., vol. 8282 of *Lecture Notes in Computer Science*, Springer, pp. 355–379.

[9] Junod, P., and Karlov, A. An efficient public-key attribute-based broadcast encryption scheme allowing arbitrary access policies. In *Proceedings of the 10th ACM Workshop on Digital Rights Management, Chicago, Illinois, USA, October 4, 2010* (2010), E. Al-Shaer, H. Jin, and M. Joye, Eds., ACM, pp. 13–24.

[10] Koblitz, N., and Menezes, A. Pairing-based cryptography at high security levels. In *Cryptography and Coding, 10th IMA International Conference, Cirencester, UK, December 19-21, 2005, Proceedings* (2005), N. P. Smart, Ed., vol. 3796 of *Lecture Notes in Computer Science*, Springer, pp. 13–36.

[11] Kurihara, J., Uzun, E., and Wood, C. A. An encryption-based access control framework for content-centric networking. In *2015 IFIP Networking Conference (IFIP Networking)* (May 2015), pp. 1–9.

[12] Li, Q., and Zhang, F. A fully secure attribute based broadcast encryption scheme. *I. J. Network Security 17*, 3 (2015), 255–263.

[13] Lubicz, D., and Sirvent, T. Attribute-based broadcast encryption scheme made efficient. In *Progress in Cryptology – AFRICACRYPT 2008: First International Conference on Cryptology in Africa, Casablanca, Morocco, June 11-14, 2008. Proceedings* (Berlin, Heidelberg, 2008), S. Vaudenay, Ed., Springer Berlin Heidelberg, pp. 325–342.

[14] Malik, A. M., Borgh, J., and Ohlman, B. Attribute-based encryption on a resource constrained sensor in an information-centric network. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking* (New York, NY, USA, 2016), ACM-ICN '16, Association for Computing Machinery, p. 217âĂŞ218.

[15] Roy, S., and Chuah, M. Secure data retrieval based on ciphertext policy attribute-based encryption (cp-abe) system for the dtns.

[16] The PARI Group. *PARI/GP version* `2.11.2`. Univ. Bordeaux, 2019. available from http://pari.math.u-bordeaux.fr/.

[17] Yamada, S., Attrapadung, N., Hanaoka, G., and Kunihiro, N. A framework and compact constructions for non-monotonic attribute-based encryption. In *Public-Key Cryptography – PKC 2014: 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings* (Berlin, Heidelberg, 2014), H. Krawczyk, Ed., Springer Berlin Heidelberg, pp. 275–292.

[18] Yu, Y., Afanasyev, A., and Zhang, L. Name-based access control. Tech. Rep. NDN-0034, NDN, 2016.

[19] Zhang, L., Afanasyev, A., Burke, J., Jacobson, V., Claffy, K., Crowley, P., Papadopoulos, C., Wang, L., and Zhang, B. *Named data networking*, 3 ed., vol. 44. Association for Computing Machinery, 2014, pp. 66–73.

[20] Zhang, L., and Yin, H. Recipient anonymous ciphertext-policy attribute-based broadcast encryption. *I. J. Network Security 20*, 1 (2018), 168–176.

[21] Zhang, Z., Yu, Y., Ramani, S. K., Afanasyev, A., and Zhang, L. NAC: Automating access control via Named Data. In *MILCOM 2018-2018 IEEE Military Communications Conference (MILCOM)* (2018), IEEE, pp. 626–633.

[22] Zhou, Z., and Huang, D. On efficient ciphertext-policy attribute based encryption and broadcast encryption: Extended abstract. In *Proceedings of the 17th ACM Conference on Computer and Communications Security* (New York, NY, USA, 2010), CCS '10, ACM, pp. 753–755.