



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Broadcast Encryption

Dave Lassalle
GSEC Practical 1.4c, Option 1
Submitted January 12, 2005

Contents

<u>Abstract</u>	2
<u>1 Introduction</u>	3
<u>2 How Broadcast Encryption Works</u>	3
<u>3 Space Requirements for Broadcast Encryption</u>	5
<u>4 Long-lasting Broadcast Encryption Schemes</u>	7
<u>5 Broadcast Encryption Versus Public-Key Cryptography</u>	8
<u>6 Protecting Digital Media in the Home</u>	10
<u>7 Conclusion</u>	11
<u>References</u>	12

Abstract

Broadcast encryption is a type of encryption scheme first proposed by Amos Fiat and Moni Naor in 1993. Their original goal was to prove that two devices, previously unknown to each other, can agree on a common key for secure communications over a one-way communication path. Broadcast encryption allows for devices that may not have even existed when a group of devices was first grouped together to join into this group and communicate securely. This paper describes broadcast encryption in general, how a few different broadcast encryption schemes work, differences between broadcast encryption and public-key cryptography, and some common uses of broadcast encryption, including cable TV systems and digital content protection.

© SANS Institute 2005, Author retains full rights.

1 Introduction

Traditionally, secure transmission of information has been achieved through the use of public-key cryptography. For this system to work, communicating devices must know about each other and agree on encryption keys before transmission. Broadcast encryption seeks to solve the problem of two devices, previously unknown to each other, agreeing upon a common key. This can allow for new devices, even if they did not exist when the encrypted data was made, to be added to a group of acceptable devices. Since the same data is being sent to all devices, instead of a separately encrypted message for each, broadcast encryption must also ensure that only those devices in the privileged group will be able to decode the message.

Broadcast encryption was proposed in the classic paper of the same name by Amos Fiat and Moni Naor in 1993 (Fiat and Naor 480). In my paper, I will provide a description of how broadcast encryption works, its space and computation requirements, and ways to make it resistant to pirate decoders over time. I will also provide a comparison of broadcast encryption to public-key cryptography and a description of how broadcast encryption can be used to protect digital content within the home.

2 How Broadcast Encryption Works

A true broadcast encryption scheme is one in which the same message is broadcast to all users, and those users in the privileged group recover the message while all others derive nonsense or nothing at all. The original broadcast encryption scheme designed by Fiat and Naor proposed the following scenario. There exists a key distribution center and a group of users. The center allocates predefined keys for all of the users. The center later wants to transmit to a privileged subset of users. For this to occur, this subset of users must recover a common key while not allowing any other users who are able to receive the transmission to recover this key. The privileged subset of users can be fixed, slowly changing, or rapidly changing.

Fiat and Naor's original scheme was a k -resilient broadcast encryption scheme, meaning it was "secure against a coalition of at most k non-privileged users" (Blundo and Cresti 288). This was also a zero-message scheme meaning the broadcast center did not have to broadcast a message for the users to be able to compute the key. It could be computed from information the user receives from the center, called the management key block, and from other users in the set. Since, their original design required a large amount of memory usage for users to store keys, Fiat and Naor proposed other schemes requiring fewer keys. However, these schemes were not unconditionally secure since they relied upon unproven complexity assumptions.

In this type of scheme, called a key predistribution scheme, there is a preprocessing phase in which the center, not knowing the privileged subset of users nor the common key, distributes a number of keys to the users. Then, in the broadcast-encryption phase, the center creates a set of messages containing the common key and encrypted using the keys it distributed in the previous phase (Blundo and Cresti 297). After broadcasting this set of messages, users in the privileged group can recover the common key. Even if they know all the broadcast messages of the other users, a coalition of non-privileged users cannot recover any information regarding the common key.

Lotspiech, Nusser, and Pestoni provide a simple explanation of how broadcast encryption works (57). Various broadcast encryption schemes are based on a key management block. This is a block of data located at the beginning of a broadcast or prerecorded onto some type of blank media, most often a smart card. From this key management block, each recipient can derive the management key. A device not in the privileged group of devices even with access to the encoded data will derive the wrong answer from the key management block. Restricted devices can attempt to process the key management block but they will not yield the correct key.

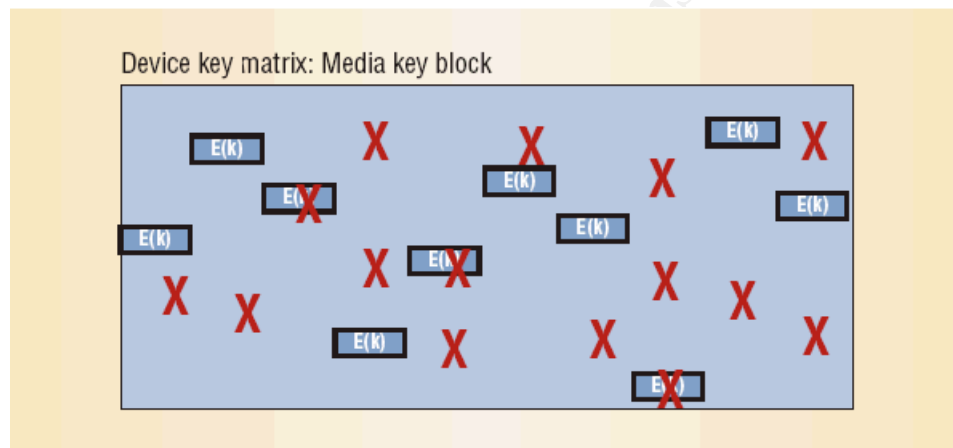
In their paper “New Constructions on Broadcast Encryption and Key Pre-Distribution Schemes,” Huang and Du compare typical broadcast encryption schemes and group key distribution schemes (1). The broadcast encryption scheme can be viewed as a “revocation scheme” (Huang and Du 1). Since the central authority predetermines the privileged subset of users and distributes the common key to this subset, it is easy to revoke access to the non-privileged users. However, it is non-trivial to allow new users to join the privileged group.

Group key distribution schemes, on the other hand, can provide join and revocation operations. The central authority, when the membership of the group changes, will choose a new session key, encrypt it, and broadcast it to the set of users. When a user is revoked, her secret information will not be used again. This makes it easy to allow the joining of new members; however, revoking multiple users becomes impractical due to all the extra updating. Huang and Du propose that the join and multiple revocation operations may be “intrinsically incompatible” with each other (2).

One type of broadcast encryption was developed by the 4C Entity formed by IBM, Intel, Matsushita, and Toshiba. This scheme, called content protection for recordable media (CPRM), uses a matrix of device keys that is 16 columns wide and over 2,000 rows tall. A device that knows the correct key position in the matrix can decrypt the value at that position to retrieve the management key. Each device actually has 16 different device keys associated with it, one from each column; but no two devices will have all 16 keys the same.

Figure 1: Device Key Matrix (Lotspiech, Nusser, Pestoni 58)

Figure 1, a scaled down version of the larger CPRM matrix, more clearly demonstrates how the device key matrix works. At first, a device can use any of its keys to decrypt the corresponding position in the matrix. The device begins with one key from each column, represented by $E(k)$ in the graphic. However, if one of the keys is compromised, it can be made unavailable. Locations in the matrix with a red X in the figure are keys that have been compromised. When a device is attempting to decode a message and encounters a key that has been crossed out like this, it can use one of its other keys.



In the extreme case that so many keys have been compromised that innocent devices are neglected access, the system is considered broken. Therefore, CPRM has a finite, although large, ability to withstand attacks from circumvention devices. Other, more complicated schemes have been devised that can provide unlimited revocation of pirate devices.

In 2000, Dalit Naor, Moni Naor, and Jeffrey Lotspiech designed another broadcast encryption scheme called a logical key hierarchy (41). In this design, devices are grouped into a tree structure. Leaf nodes in the same area of the tree can be denied access by the key distribution center. Their design reduced the size of the key management block to roughly the same size of a public key certificate revocation list, eliminating the remaining advantage of public key cryptography.

3 Space Requirements for Broadcast Encryption

There are two important size requirements in a broadcast encryption scheme:

the number of keys that must be stored in a device and the number of messages a broadcast center must send. Broadcast schemes are based on the idea of guaranteeing a maximum number of users who cannot gain any information about the keys or about the encrypted message. This number of users is usually represented by k , and the total number of users is represented by n .

The original scheme designed by Fiat and Naor required that every user in the broadcast group store $O(k \log k \log n)$ keys and that the broadcast center transmit $O(k^2 \log^2 k \log n)$ messages (Fiat and Naor 483). This scheme guaranteed that any coalition of k users could not acquire any information about the keys or the broadcast message and is called a k -resilient scheme. These complexities are too high for applications such as pay-TV systems because pirate decoders can easily obtain and analyze thousands of smart cards to break the system (Halevy and Shamir 48).

Based on results obtained by Blundo and Cresti, to obtain memory requirements smaller than those found by Fiat and Naor, you must resort to "unproven complexity assumptions" (Blundo and Cresti 288). These assumptions include stating that a "one-way function exists" or "extracting prime roots modulo a composite is hard" (Blundo and Cresti 288). They also showed that interaction among users, in which users are allowed to set up a common key, does not decrease the size of the information given to the users. The only way they found to reduce the amount of information users must store is to relax the security requirement. This is accomplished by allowing the users to compute a common key a finite number of times before information might not be completely secure.

One can think of these aspects in terms of the confidentiality, integrity, availability, or CIA, dimensions of computer security. On one hand is the secure information theoretic framework with a smaller scope in durability. This scheme is cryptographically secure and, thus, provides the best chance for confidentiality. However, over time it may degrade in usefulness or availability if a large number of keys are compromised. On the other hand, if a scheme is based on unproven complexity assumptions, although we generally regard these assumptions as being safe for secure transmission of information, this is still a relaxation in the security of the encryption and, therefore, the absoluteness in confidentiality. However, this relaxation allows for higher availability in that the cryptosystem will be less likely to be broken.

A simple solution to the space requirement issue is to give each user a symmetric key that only that user and the broadcast center know. The broadcast center can encrypt a session key using each user's symmetric key and broadcast these encrypted messages to each user. Finally, the broadcast center transmits the encrypted content to each user using his respective session key. This reduces the space requirements for the users to $O(1)$, but the transmission length is increased to $O(s)$ where s is the privileged set of users

(Halevy and Shamir 48). This scheme is only useful if the number of privileged users is small.

The logical tree hierarchy mentioned earlier improves on this by treating devices as the leaf nodes in a tree structure. Each vertex in the tree is assigned a key and every device knows the keys of its ancestor nodes. The memory requirements for the users in this case is $O(\log n)$, the number of nodes traced by following the path from the leaf node to the root node. Note that in this case, all devices know the key of the root node. A message can then be encrypted using one of these keys and all the leaf nodes in that node's subtree will be able to decrypt the message. This scheme works well if related users are grouped into the same subtrees rather than in random order.

Fiat and Naor proposed the subset difference (SD) scheme further improving on these requirements. A later improvement by Halevy and Shamir, called the layered subset difference (LSD) scheme, enables each user to store one kilobyte worth of keys on a smart card and the broadcast center can thereby revoke any number of users r out of about 256 million users by transmitting at most $4r$ messages and on average $2r$ messages (49).

The LSD method is based on creating the set of privileged users by performing inclusion and exclusion operations on subsets of users. Each subset has a key associated with it. Again, this can be most easily accomplished by grouping users in a balanced binary tree, with each vertex representing a key that all leaf nodes in that subtree know, and then including or excluding certain subtrees.

This nesting inclusion and exclusion of subsets allows the following scenario. Consider a football game being broadcast on a national level to a cable television company's subscribers. The television company allows all subscribers access to the broadcast, except for the local network where a blackout is in place. However, sports bars in the local viewing area with a special subscription are allowed to receive the broadcast, while any sports bar without the special subscription is still excluded. If the subscribers are grouped in a tree structure based on geography and subscription type, this operation could easily be performed using the LSD method. If the leaf nodes are not grouped in a logical way, essentially, the message will have to be encrypted using mostly leaf node keys and the number of messages broadcast will be on the order of the number of devices. This would be extremely impractical, so the grouping becomes very important (Halevy and Shamir 48).

4 Long-lasting Broadcast Encryption Schemes

As mentioned before, although it is not likely because of the large space of device keys, it is possible for all the keys to be compromised and for the encryption scheme to break. Garay, Staddon, and Wool proposed a way to extend the lifetime of a broadcast encryption scheme (333). They describe a

system in which keys for devices are stored on smartcards. When a pirate decoder has been found, the keys associated with its smartcard will be revoked. A user's keys can also be revoked if her subscription expires.

When all the keys in an innocent device have been revoked, its smartcard will have to be replaced with a new set of keys. Keys also need to be replaced if the contract for a given device has expired. Garay, Staddon, and Wool seek to minimize the number of smartcards that will need to be replaced in a given period of time they define as an epoch. At the end of an epoch, the service provider must compute which users need to have smart cards replaced to continue secure communications. Thus, the cost of such a scheme becomes directly related to the cost of periodically replacing a number of smart cards in each epoch.

They analyze three different broadcast encryption schemes: a bucket-based scheme proposed by Kumar, Rajagopalan, and Sahai (609); a deterministic scheme proposed by Gafni, Staddon, and Yin (372); and their own simple randomized scheme. Each scheme is similar in the number of cards that can be tolerated before recarding must take place. To explain their results, the authors provide three numerical experiments that vary the total number of keys, the number of keys per card, and different epoch lengths. This work provides for the case where permanent revocation of keys is desired, rather than short-term prevention from access to a particular message.

For situations in which pirate decoders provide themselves and other unprivileged users access to content, traitor tracing schemes can be employed. An effective traitor tracing scheme provides the following services (Fiat and Tassa 212):

- trace the source of piracy
- disconnect the traitor and dependent unauthorized users from further transmittal of information
- harm no legitimate users
- supply legal evidence of the pirate's identity.

Traitor-tracing schemes aim to make the construction of pirate decoders risky because once a compromised key is found, the smart card it came from can be revoked.

5 Broadcast Encryption Versus Public-Key Cryptography

Broadcast encryption and public-key cryptography differ in several ways. The main difference is in the allocation of keys. As mentioned before, in broadcast encryption, devices are given a set of keys and the master server contains a key

management block. Future devices can easily be added and be allowed into the privileged group. None of the devices needs to know about each other; all they know is that each belongs to the privileged group.

Public-key cryptography, on the other hand, is based on prior knowledge of participating devices. Senders use their own private key to encode messages, and recipients use a public key to decode messages. Since these keys must be known before messages can be exchanged, each device must know about every other device, meaning it must know and store the keys of other devices, it wants to communicate with.

A disadvantage of broadcast encryption, however, is that it cannot provide a nonrefutable signature. In public-key cryptography, forgery of a valid signature is an intractable problem without the actual signer's private key. Therefore, the true identity of an individual cannot be guaranteed with broadcast encryption. Broadcast encryption can only guarantee that one participant is in the same group as another; public-key cryptography can guarantee the participant's actual identity. Broadcast encryption replaces the digital signature with a message authentication code, or MAC, which is a weaker system. Any device that can process the key management block and generate the management key can verify the MAC.

An advantage of public-key cryptography systems is that they do not require a central authority. With broadcast encryption, a central authority, the broadcast center, must produce and assign key management blocks and assign device keys. In public-key systems like Pretty Good Privacy (PGP), users create their own certificates and exchange them throughout the system (Zimmerman 9). Some public-key systems do use a central authority for key distribution, but it is not a requirement. In other cases, public-key systems use a "web of trust" in which users verify the authenticity of the keys of other users' keys instead of a central authority performing this task (Feisthammel). For example, if user A trusts user B, and if user B has verified the key of user C, user A can be confident that the person using user C's key, with whom A is communicating, is actually user C. This is a simplified example. A real web of trust would involve more people verifying the authenticity of a user's key before another user would safely believe the identity of a previously unknown user.

Advantages of broadcast encryption include increased speed over public-key cryptography and the ability to adapt to attacks on the system. Since broadcast encryption performs simple symmetric operations and public-key cryptography uses exponentiation operations, the processor load on a broadcast encryption system can be up to 1,000 times less than the load to perform a public-key signature calculation (Lotspiech, Nusser, and Pestoni 58).

Also, as mentioned before, the ability to remove compromised keys from the system is a major advantage to provide longer life and durability to the system.

Without the ability to revoke compromised keys, the system degenerates into a shared secret system and is broken when one key is discovered by an unauthorized party. For example, the Content Scrambling System (CSS) encryption scheme for DVDs cannot be fixed without redesigning the system from scratch. If a broadcast encryption scheme had been used, the creators could have simply released new discs that disallowed pirate programs while not affecting legitimate users (Lotspiech, Nusser, and Petroni 57).

Copyright protection has become an important application for cryptography. Instead of a system in which only authorized users have access to the keys, copyright systems provide keys to all users because one cannot tell the difference. As a result, a box's sensitivity to "reverse engineering becomes an important issue" (Lotspiech, Nusser, and Petroni 58). Public-key systems perform a handshake at the link-level requiring keys to be placed in the link-level code where they might be easier to find by malicious users. On the other hand, since their systems are one-way, broadcast encryption schemes have the advantage that they can hide their keys much deeper in the software making the keys more difficult for malicious users to discover.

Broadcast encryption is not applicable in all applications; however, it is especially useful in content protection. It can be useful in pay-TV systems, distributing copyrighted information of CD and DVD disks, and multicasting music and video on the internet. For example, its low overhead, revocation ability, and resistance to reverse engineering are of utmost importance in consumer electronics. The disadvantages of broadcast encryption do not present any problems to its use in content protection.

Since it is not necessary to know exactly what particular device another device is connecting to, just that it can be granted access, nonrefutable signatures are not necessary. Computing the MAC will allow a device to know that it is communicating with a compliant device, and that is sufficient. Finally, having a central authority is actually a desirable quality in content protection systems which focus on licensing content to specific users. Relating back to the CIA view of computer security, the MAC provides integrity to the devices, that a given device may know it is communicating with another approved device. The central authority provides the confidentiality and availability aspects by encrypting transmissions to privileged users and only making these communications available to authorized users.

6 Protecting Digital Media in the Home

When broadcast encryption was first designed, the creators primarily saw it as a means to provide conditional access, allowing only privileged users or devices access to a message. An ideal use for this was granting access to premium cable TV channels to paying customers. This original use for broadcast encryption has turned out to be less important than another application: media

content protection.

Content protection is an important topic since home users now have access to all types of media in digital form. Since the millionth digital copy is just as good as the original, protecting digital content and the rights of its creators is of increasing concern. Traw discusses the two primary methods for content protection: licensing and technology (42). Licensing is a more effective approach for some users because, given enough time, the technically savvy will probably find ways around technological protections. However, providing secure and effective technology-based protections will keep the masses of casual users from infringing on others' intellectual property.

This can be viewed as a type of defense-in-depth strategy for major media publishers. Defense-in-depth has long been known as a beneficial quality in the realm of security, providing for several methods to secure data in case any one of them is found to be deficient. These two content protection styles, licensing and technology, are both effective at stopping users of different skill levels from accessing content they are not entitled to. Ripley, et. al. show that, as the skill level of the user increases from "casual copier" to "professional pirate," licensing becomes increasingly important as the more effective approach (50). On the other hand, technology becomes less important because, while an advanced cryptosystem will stop the casual copier, the professional pirate will usually find ways to defeat a system.

Broadcast encryption is most appropriate for content protection in the home in devices such as prerecorded DVD media and CD media. Traw discusses the use of CPRM, described earlier, as an effective way to provide technological protection of intellectual property (44). CPRM is a good method because it is effectively hard to break the scheme. Also, if some keys are compromised, the system is not rendered useless, like the CSS encryption method for DVDs; instead, the unauthorized devices can be blocked from future communications when they are discovered.

7 Conclusion

Broadcast encryption is not applicable for all situations. However, there are some cases in which it could prove to be better than public-key cryptography. One such example is the case of DVD encryption. The Content Scrambling System in DVDs has already been broken, and it is easy to find programs that can decrypt encoded DVDs. Since the key scheme for CSS is a shared-secret scheme, it cannot be changed dynamically like a broadcast encryption scheme, so the security breach cannot be fixed. This is an example of why broadcast encryption is an important area of research and should be used when appropriate.

Professionals implementing security systems should know about this

technology and implement it when it is appropriate. Broadcast encryption schemes provide many benefits over other technologies especially when used in the realm of content protection.

© SANS Institute 2005, Author retains full rights.

References

- Blundo, Carlo, and Antonella Cresti. "Space Requirements for Broadcast Encryption." *Advances in Cryptology* (Eurocrypt 1994). Lecture Notes in Computer Science 950. Springer-Verlag, 1995: 287-298.
- Feisthammel, Patrick. "PGP: Explanation of the web of trust of PGP." 7 Oct. 2004. <<http://www.rubin.ch/pgp/weboftrust.en.html>>.
- Fiat, Amos, and Moni Naor. "Broadcast Encryption." *Advances in Cryptology* (Crypto 1993). Lecture Notes in Computer Science 773. Springer-Verlag, 1994: 480-491.
<<http://citeseer.ist.psu.edu/rd/30505408%2C421418%2C1%2C0.25%2CDownload/http://citeseer.ist.psu.edu/cache/papers/cs/20701/http:zSzzSzwww.wisdom.weizmann.ac.ilzSz%7EnaorzSzPAPERSzSzSzbroad.pdf/fiat94broadcast.pdf>>.
- Gafni, E., J. Staddon, and Y. Yin. "Efficient Methods for Integrating Broadcast Encryption and Traceability." *Advances in Cryptology* (Crypto 1999). Lecture Notes in Computer Science 1666. Springer-Verlag, 1999: 372-387.
- Garay, Juan A., Jessica Staddon, and Avishai Wool. "Long-lived Broadcast Encryption." *Advances in Cryptology* (Crypto 2000). Lecture Notes in Computer Science 1880. Springer-Verlag, 2000: 333-352.
<<http://citeseer.ist.psu.edu/cache/papers/cs/14825/http:zSzzSzwww.bell-labs.comzSzuserzSzgarayzSzbcast-enc.pdf/long-lived-broadcast-encryption.pdf>>.
- Halevy, Dani and Adi Shamir. "The LSD Broadcast Encryption Scheme." *Advances in Cryptology* (Crypto 2002). Lecture Notes in Computer Science 2442. Springer-Verlag, 2002: 47-60.
- Huang, Scott C.-h., and Ding-zhu Du. "New Constructions On Broadcast Encryption Key Pre-Distribution Schemes." *Technical Report 04-027*. 24 Jun. 2004: <https://www.cs.umn.edu/tech_reports_upload/tr2004/04-027.pdf>.
- Kumar R., S. Rajagopalan, and A. Sahai. "Coding Constructions for Blacklisting Problems without Computational Assumptions." *Advances in Cryptology* (Crypto 1999). Lecture Notes in Computer Science 1666. Springer-Verlag, 1999: 609-623.
- Lotspiech, Jeffrey, Steffan Nusser, and Florian Pestoni. "Broadcast Encryption's Bright Future." *IEEE Computer* 35.8 (August 2002): 57-63.

Naor, Dalit, Moni Naor, Jeffrey Lotspiech. "Revocation and Tracing Routines for Stateless Receivers." *Advances in Cryptology (Crypto 2001)*. Lecture Notes in Computer Science 2139. Springer-Verlag, 2001: 41-62.

Ripley, Michael, et. al. "Content Protection in the Digital Home." *Intel Technology Journal*. 15 November 2004: 48-56. 9 January 2005.
<http://www.intel.com/technology/itj/2002/volume06issue04/art05_protection/p01_abstract.htm>.

Traw, Brendan S. "Protecting Digital Content Within the Home." *IEEE Computer* 34.10 (October 2001): 42-47.

Zimmermann, Paul R. *The Official PGP User's Guide*. Cambridge: MIT Press, 1995.

© SANS Institute 2005, Author retains full rights.