

Programming Project #1

Due: Friday, February 9th, 2007.

1 Overview

For the first programming assignment you will be implementing a secure content distribution system similar to AACCS, which is used for HD-DVDs. The DVD manufacturers have content that they would like to distribute widely yet retain control over. Only those DVD players that have an appropriate set of keys should be able to recover the content. The manufacturer should be able to revoke the privileges of a player whose keys were compromised. This revocation policy must support disabling of individual players.

For the purposes of this project, there are a maximum of 2^{32} players. The players can be viewed as the leaves of a binary tree of height 33. Each node in the tree contains an AES key, all of which are kept secret from consumers. A specific player will contain (in a file) the 33 keys associated with the path from the root to leaf nodes. The manufacturer should be able to create content that cannot be decrypted by any revoked player.

Conceptually, the system works in three phases:

- In the *key generation phase*, a master AES key (K_{AACCS}) is used to generate node keys. So when a new DVD player is “created”, a file is generated with the 33 player keys. This key file is protected with a password-based key.
- In the *manufacturing phase*, the DVD manufacturer encrypts content using a set of keys that is derived from the key tree and the revocation list.
- In the *playback phase*, the DVD player attempts to decrypt the content using its player keys. If the player has not been revoked, the original content should be recovered.

The required security features are :

- Secure storage (password-protected) of player keys.
- Encryption of all content with AES in the counter mode (CTR).
- Integrity check for the protected content using Message Authentication Codes (MACs).

We will examine each of these features in detail in the following section.

2 Security Features

2.1 Secure generation and storage of player keys

Before a player can be used, its set of 33 keys must be created and securely stored in a file. The key file will be generated by a command-line utility, that takes a 32-bit player number and a password. This utility will then generate the 33 keys representing a path from the root to the leaf, corresponding to the given player number. The keys will then be stored in an encrypted file. The file encryption key will be generated using the provided password.

Note that the full tree of all keys is never stored in one place, for security and storage reasons. So, you will have to derive all player keys, as needed, from K_{AACs} . We will provide you a representation of the unique IDs for each tree node, but you will need to come up with a way to derive an AES key from a node ID. We will also provide a method that you can use to convert a password into an AES key.

2.2 Content Encryption

Content must be encrypted by the DVD manufacturer using AES in CTR mode, with a new, random *title key* for each title encrypted. In order to ensure that revoked players cannot decrypt content, a revocation list must be available to the manufacturer. Then, the *title key* K_t will be encrypted under an appropriate subset (*cover*) of the keys in the tree, and stored along with the encrypted content. This content header should also contain a *content title* string, which is descriptive metadata about the content. The content header plus the encrypted content must have an associated MAC.

Note that the key used for encrypting content, K_{enc} , and the key used for MAC'ing K_{mac} must be derived from the *title key*:

- $K_{enc} = HMAC(K_t, "enc")$
- $K_{mac} = HMAC(K_t, "mac")$

Along with each encrypted *title key*, the content header should contain an encoded node ID, for efficient decryption. Otherwise, the player would have to try every combination of (player key, encrypted *title key*). Your decryption solution should run in time proportional to $k + l$ as opposed to $k * l$, where k is the number of keys in the player and l is the number of encrypted copies of the *title key* in the content header

2.2.1 Counter Mode (CTR)

Counter mode encryption generates a pseudorandom sequence by encrypting successive values of a counter. Formally, encryption of a message (m_0, m_1, \dots, m_n) is $(IV, E(k, IV) \oplus m_0, E(k, IV + 1) \oplus m_1, \dots, E(k, IV + n) \oplus m_n)$. Similarly as in other modes of encryption, new IV should be chosen randomly each time.

2.3 Integrity Check using MACs

The key used to generate MACs for the content should be derived from the *title key*. This will enable detection of an attacker tampering with the protected content.

2.4 Content Decryption

Before attempting to decrypt content, the DVD player must access its player keys from an encrypted file. Then it will try to determine the *title key* for the encrypted content. There are three possible outcomes:

- If none of the player keys can be used to recover the *title key*, then this player has been revoked, and an appropriate error message should be output.
- If the MAC is invalid, then the protected content has been tampered with, and an appropriate error message should be output.
- Otherwise, the original unencrypted content should be output.

3 Components

3.1 PlayerKeys

PlayerKeys implements functionality related to the generation of the player keys. The main purpose is to provide a command line utility for generating an encrypted file containing keys for a given player.

In generating player keys, the PlayerKeys utility will take the *master AACs password* (used to derive K_{AACs}), a player number n and a *player password* (used to encrypt the player key file). The output filename is *player_n.key*. The same *player password* will be used by a DVDPlayer to decrypt its key file. Note that all nodes in the key tree are derived from the master key K_{AACs} . Rather than persistently storing K_{AACs} you will derive it from the *master password* provided on the command line. The same *master password* must be used for all operations on a particular key tree.

- PlayerKeys will generate an AES key, K_{enc} , and a MAC key, K_{mac} , from the *player password* shared with player n .
- Then it will encrypt and MAC the player keys for DVD player n :

$$E[K_{enc}, \text{player_keys}] || MAC[K_{mac}, E[K_{enc}, \text{player_keys}]]$$

- Note that the list of player keys should include some identifying information for each key. This will make determining the appropriate player key to use (in the decryption step) more efficient.
- Then PlayerKeys will write all of this to the given output file and exit.

3.2 KeyTree

The KeyTree class implements all functionality querying of the tree of all node keys. The main functions provided for you are:

- An API (*getPathNodes*) for retrieving a list of nodeIDs starting from the root and ending at a leaf (i.e. a player).
- An API (*getCoverSet*) that takes as input a list of revoked player numbers and returns a list of keys (a cover) needed to encrypt a *title key* under. This will be used by the DVDManufacturer.
- An API (*createAESKeyMaterial*) for converting a password into a byte array that can be used as an AES key.

3.3 DVDManufacturer

The DVDManufacturer will be run as a command line utility that takes a filename (of some content) and a *content title* (a metadata string). As with the PlayerKeys utility, you will derive K_{AACS} from a master password provided on the command line, and use it to derive any needed node keys. The revocation list file is expected to be at *revoke.lst*, and the output file name is the input filename with a **.enc** extension.

Before encryption is started, a new random *title key* K_t is generated, and encryption/MAC keys K_{enc} and K_{mac} are derived from it (as described above). First, the content must be encrypted with K_{enc} using AES in CTR mode. Next, the KeyTree is queried for a list of nodes under whose keys K_t will be encrypted. From this list you will derive a set S of AES keys. Then, the content header is generated and the header and content are MAC'd using K_{mac} .

Finally, an output file will be generated:

$$[\text{content_header}] \parallel E[K_{enc}, \text{content}] \parallel \text{MAC}[K_{mac}, [\text{content_header}] \parallel E[K_{enc}, \text{content}]]$$

Where the *content.header* consists of K_t encrypted under each key in S , as well as some identifying information for each encrypted key and a *content title*. This identifying information should allow a player to determine which of its keys to use to try and decrypt K_t .

3.4 DVDPlayer

The DVDPlayer will be run as a command line utility that takes the player key file password, and the filename of some encrypted content. The player keyfile is expected to be *player.n.key* and the output file (if generated) will be the encrypted filename without a **.enc** extension. Before decryption is started, the player key file is read in, decrypted, and the MAC checked. Then, the protected content is read in and the player attempts to recover the *title key* K_t from the content header. Given K_t , the MAC will be verified and the content decrypted.

4 Implementation

You will be using the JCE (Java Cryptographic Extensions) while programming for this assignment. You should spend some time getting familiar with the provided framework.

4.1 Getting the code

Download the *pp1.tar.gz* file linked on site to a directory in your account. Untar and unzip using the following command:

```
tar xvzf pp1.tar.gz
```

This should create the source tree for the project under the *pp1/* directory.

4.2 Description of the code

Here is a brief description of the files we provide. The files you need to change are in **bold**

Makefile	Makefile for the project
revoke.lst	Revocation list file
aacs/PlayerKeys.java	Used to generate player keyfile
aacs/KeyTree.java	key tree utilities
aacs/DVDManufacturer.java	The encryption program
aacs/DVDPlayer.java	The decryption program

4.3 Running the code

To build the project, enter the *pp1* directory and type *make*. To run the system, follow these steps:

1. Run the PlayerKeys utility to generate a set of keys for a player *n*. You will also specify a password for deriving the AACS Master Key and a keyfile password. The output filename will be *player_n.key*.

```
elaine21:~/pp1> java aacs/PlayerKeys <AACSPwd> n <keyfile pwd>
```

2. Encrypt some content in a file. You will specify a metadata string (e.g. a title) and a content file name. You will also specify a password for deriving the AACS Master Key. Note: the revocation list file is in *revoke.lst* and the output file will be *content file name.enc* .

```
elaine21:~/pp1> java aacs/DVDManufacturer <AACSPwd> <content title>
<content file name>
```

3. Decrypt the encrypted content file. You will specify a player number, the keyfile password, and an encrypted content file. The player key file is expected to be *player.n.key* and the output file name will be the encrypted file name without the **.enc** extension.

```
elaine21:~/pp1> java aacs/DVDPlayer n <keyfile pwd>
<encrypted content file name>
```

Note: Your solution will be tested on the elaine machines. So, please test your code on one of the elaines before submitting.

4.4 Crypto Libraries and Documentation

Javas security and cryptography classes are divided into two main packages: `java.security.*` and `javax.crypto.*`. They have been integrated into Java 2 Platform Standard Edition v 1.5. Classes for cryptographic hashing and digital signatures (not required for project 1) can be found in `security`, whereas ciphers and MACs are located in the JCE.

The following are some links to useful documentation :

- Java API
<http://java.sun.com/j2se/1.5.0/docs/api>
- JCE Reference Guide
<http://java.sun.com/j2se/1.5.0/docs/guide/security/jce/JCERefGuide.html>
- Java Tutorial
<http://java.sun.com/docs/books/tutorial/>
- Chapter 6 from Java Cryptography by Jonathon Knudsen
<http://www.oreilly.com/catalog/javacrypt/chapter/ch06.html>

Some classes/interfaces you may want to take a look at:

- `javax.crypto.KeyGenerator`
- `javax.crypto.SecretKey`
- `javax.crypto.IvParameterSpec`
- `javax.crypto.Mac`
- `javax.crypto.Cipher`
- `javax.crypto.CipherInputStream`

- javax.crypto.CipherOutputStream
- javax.crypto.SecretKeyFactory

- java.security.MessageDigest
- java.security.SecureRandom

- java.math.BigInteger

5 Miscellaneous

5.1 Questions

- We strongly encourage you to use the class newsgroup (su.class.cs255) as your first line of defense for the programming projects. TAs will be monitoring the newsgroup daily and, who knows, maybe someone else has already answered your question.
- You can also email the staff at cs255ta@cs.stanford.edu

5.2 Deliverables

In addition to your well-decomposed, well-commented solution to the assignment, you should submit a README containing the names, leland usernames and SUIDs of the people in your group, a description of the design choices you made in implementing each of the required security features, and a short answer to this question:

5.2.1 Breaking AACCS

Consider the following two types of hackers:

1. A hacker who recovers her DVD Player keys and anonymously publishes them on the web.
2. A hacker who purchases a DVD, recovers a title key, and anonymously publishes it on the web.

How would the AACCS licensing authorities deal with each type of hacker?

When you are ready to submit, make sure you are in your *pp1* directory and type:
`/usr/class/cs255/bin/submit.`