Using AES in GCM Mode for Tape Encryption

Second Draft, December 2, 2005

Revision history:

- Dec-2-2005, Modified by Glen Jaquette and Shai Halevi
- Aug-3-2005, Initial draft by Shai Halevi

Preface

This document is an early draft, intended to be used as a basis for inclusion of the GCM-AES transform in the future p1619.1 standard. It incorporates some of the verbal agreements reached in the P1619.1 teleconference on August 11, 2005. All aspects of the current draft are open for discussion.

1 Introduction

The Galois/Counter Mode of Operation (GCM) uses an underlying block cipher to encrypt messages of arbitrary length in an authenticated manner. GCM is specified in a document titled "The Galois/Counter Mode of Operation (GCM)" by David A. McGrew and John Viega. The current version of that document is from May 31, 2005, and it is available off the modes-of-operation web-page as the Computer Security Resource Center of the US National Institute for Standards and Technology (NIST/CSRC), at the URL

http://csrc.nist.gov/CryptoToolkit/modes/proposedmodes/gcm/gcm-revised-spec.pdf

The present draft uses the specification in the (normative Section 2 of the) above-mentioned document as an integral part of the proposed standard. Below we refer to that document as *the GCM document*.

The GCM mode of operation defines two transformations: namely encryption and decryption. These transformations depend on two parameters: one is the choice of the underlying block cipher, and the other is the length of the authentication tag. (The authentication tag is part of the output of GCM encryption and part of the input for GCM decryption). For the purpose of tape encryption, we restrict ourselves to the case of using AES as the underlying block cipher, and using 128-bit authentication tag. With these choices, the interfaces of the two transformations as specified in the GCM document are as follows:

1.1 GCM Encryption and Decryption

Encryption: The GCM encryption routine expects four inputs:

- A secret key K, to be used with the underlying block cipher. AES is defined to support key lengths of 128-, 192- or 256-bits long.
- An initialization vector IV that (in principle) can be of any length between 1 and 2^{64} bits.

- A plaintext P that can be of any length between 0 and (2^{39}–256) bits.
- Additional authenticated data AAD that can be of any length between 0 and 2^{64} bits.

This procedure has two outputs:

- A ciphertext C that has the same length as the input plaintext P.
- An authentication tag T, which in our case is of length exactly 128 bits.

Below we denote the GCM encryption routine (using AES) by

```
(C,T) := GCM-AES-enc(K; IV, P, AAD)
```

The security of GCM relies on the secret key being secret, and on the IV being used as a nonce. That is, GCM only offers security as long as the same value for the IV is never used for encryption of more than one plaintext under the same key.

Decryption: The GCM decryption routine has five inputs: the key K, initialization vector IV, ciphertext C, additional authenticated data AAD, and tag T, all as above. Its output is either the plaintext P as above, or the special signal *fail*. Below we denote the GCM decryption routine (using AES) by

```
P/fail := GCM-AES-dec(K; IV, C, AAD, T)
```

2 Using GCM-AES for Tape Encryption

Specifying the use of GCA-AES for tape encryption involves two aspects. One aspect is further restricting the allowed inputs to the encryption routine, and the other is specifying of what gets written to tape. These two aspects are treated separately in the next two subsections.

2.1 The inputs to GCM-AES-enc

In addition to the restrictions from the GCM document, the tape encryption standard further restricts the allowed input to the GCM encryption routine as follows:

- The key K shall be 256-bits long (i.e. use of 128- and 192-bit keys is not allowed)
- The initialization vector IV shall be exactly 96-bit long.
- The Plaintext P shall corresponds to one tape-record. That is, an entity that is always written or read in its entirety. (Put in other words: an application cannot read from tape or write to tape part of a record, only a full record.)
- The AAD input shall consist of exactly one 128-bit block.

It is noted that the third bullet above in particular requires that different tape records must be encrypted with different IVs if encrypted with the same key.

2.2 What is Stored on Tape?

When encrypting a tape using GCM-AES, the tape-drive will write to tape *for each record* the following information:

- AAD – metadata associated with the encryption of the following tape-record. This field SHALL contain enough information to let a tape-drive reading this record identify the key that was used to encrypt it (e.g., a key-identifier). It MAY also include other information.
- IV - the IV that was used
- C - the ciphertext
- T - the authentication tag
- Other – things that are needed for proper formatting of the tape (e.g. markers delimiting the start and end of the cipher text associated with a tape-record)

An example of a tape record with these fields is depicted in Figure 1:

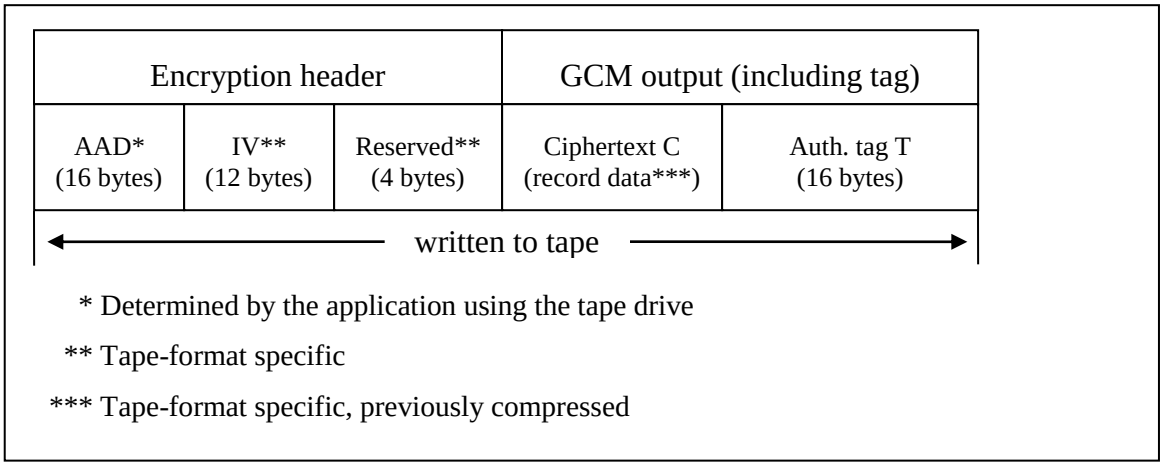

Figure 1. Format of an encrypted tape record

In addition, the implementation may write to tape also an algorithm identifier (e.g. "GCM-AES for P1619.1") but this need not be written with every record, and may be stored elsewhere on the tape cartridge.

A comment about the AAD field. The AAD field was initially meant as a key-identifier, specifying the key that was used to encrypt the current tape-record. This could be, for example, an index into a table of wrapped keys that are stored elsewhere on the tape cartridge, or an identifier that is used by an auxiliary key-management application. (Note, however, that the AAD field has only 16 bytes so it cannot be used directly to store a wrapped key.) The reason that this field is authenticated is that someone may want to store additional information there, on top of just pointing to the key.