Modernizing the OpenPGP Message Format

draft-ford-openpgp-format-00

Bryan Ford
Swiss Federal Institute of Technology (EPFL)

IETF 94 – November 3, 2015

Possible Goals for Discussion

- Modernize cryptographic suite
  - Deprecate SHA-1, shift to authenticated encryption
- Metadata protection for encrypted files
  - Leave no byte unencrypted
  - Padding to minimize leakage via length
- Partial-file integrity protection [DKG, CFRG list]
  - Streaming-mode incremental integrity checking
  - Integrity-protected random access
- Others???

Cryptographic Suite

Modernizing cipher suite, especially MACs

- Ditch SHA-1, Modification Detection packet
- Support authenticated encryption (AEAD)

Which scheme(s)? Some options:

- AES-GCM: well-established, safe if not shiny
- Keccak/SHA-3 sponge: newly standardized
- ChaCha20-Poly1305: fast, popular “alt-crypt”
- Future: CAESAR competition winner, finalist(s)

What about passphrase? Adopt scrypt/???

Cryptographic Suite

Some technical format issues:

- Repurpose existing packets (tags 18+19)
  or define new AEAD-protected packet (tag 20?)
- Merge MAC check into encrypted data packet?
  Safe to assume MACs are always fixed-length?
- AEAD nonce: always explicitly transmitted?
  Implicitly defined (e.g., by counting within file)?
- “Additional Data” (AD): any use in OpenPGP?

Metadata Protection

Should encrypted files leak all this metadata?

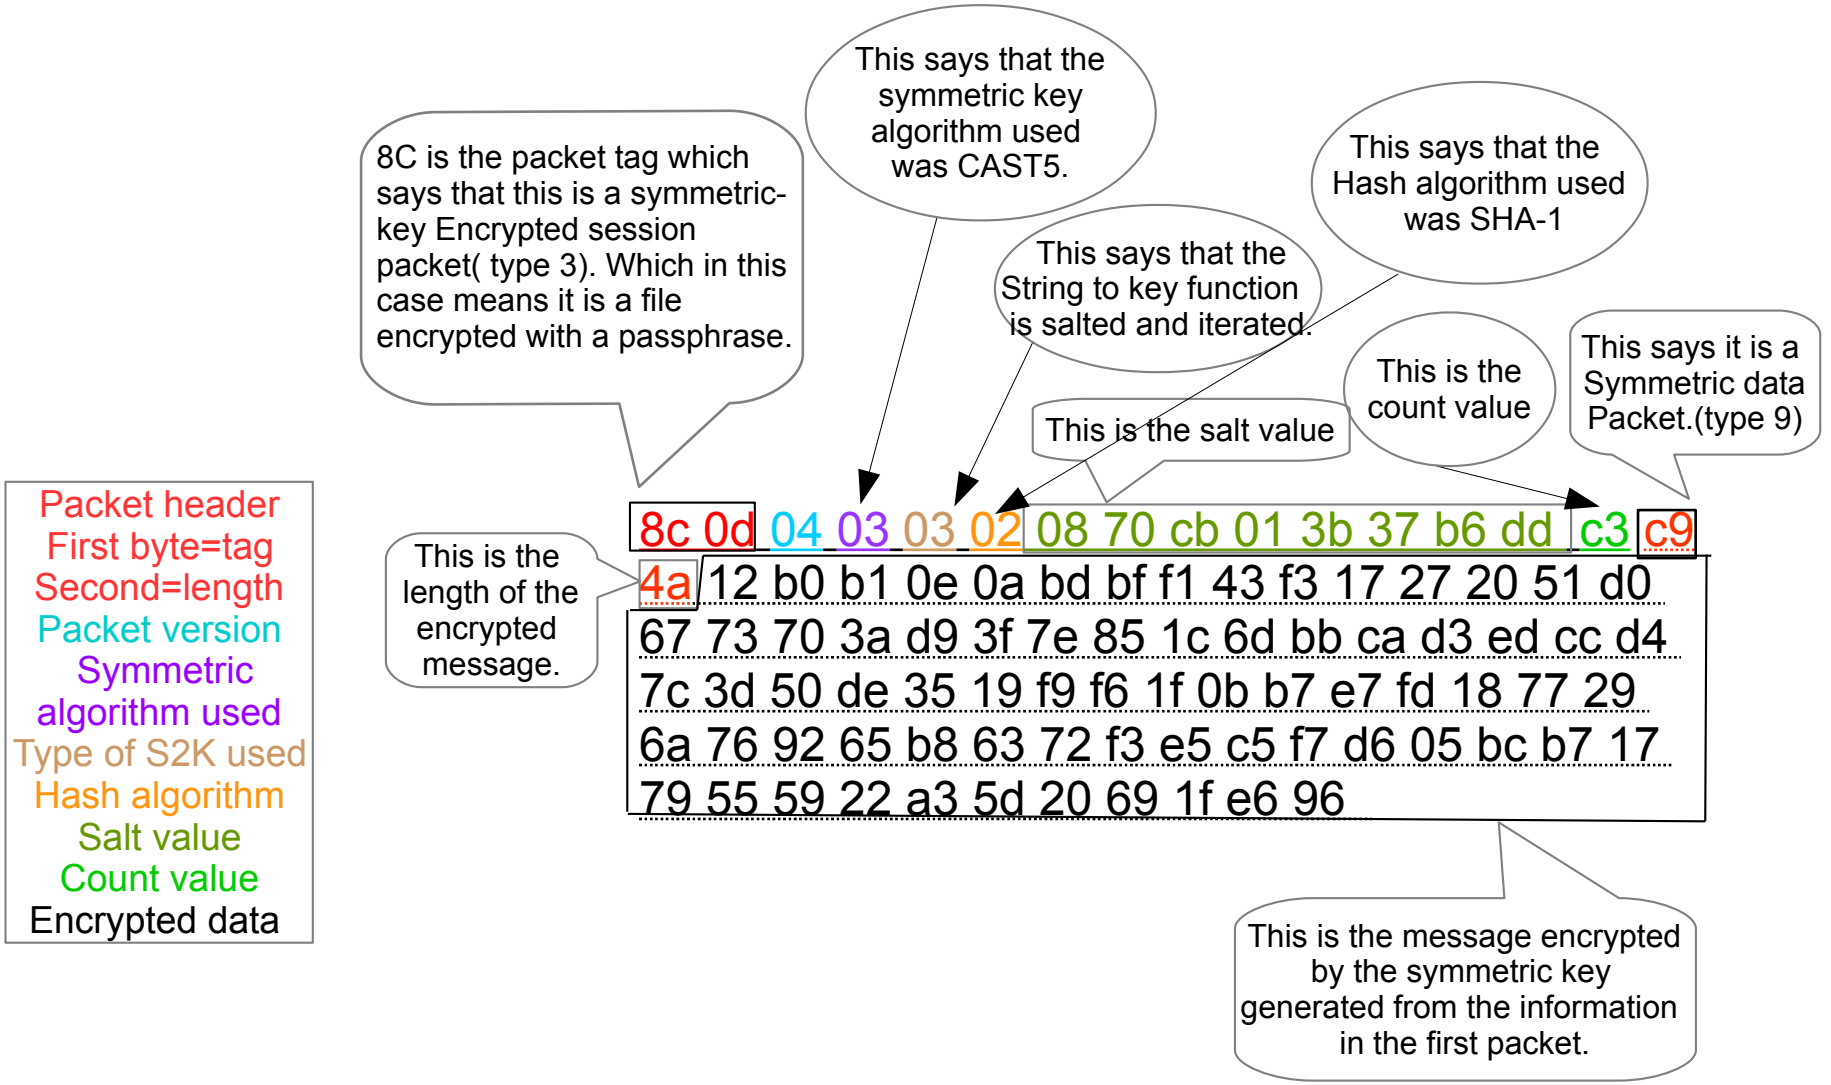

Metadata Protection

Metadata that might be useful to (some) attackers:

- Magic: this is an OpenPGP file! Suspicious!
- Cipher: is it worth trying to crack?
- Passphrase: worth trying password cracker?
- Recipient key-IDs: where to point rubber hose?
- # of recipients: aha, it's *that* group of dissidents!

Metadata Protection

Set goal to “encrypted every bit”?

- Produce Uniform Random Blobs (URBs)

Technical+usability challenges:

- How does recipient find, decrypt session key?
  - Obviously *requires* “trial decryptions”; fast enough?
- How to efficiently handle *multiple*
  - Passphrases
  - Receipient public-keys
  - Public-key schemes, curves

Claim: all are manageable. But how worthwhile?

What about Padding?

Encrypted file length leaks metadata too!

- Straw-man: pad all encrypted files to same size
  - Reduces information leakage to zero, yay!
- Wood-man: pad to next power of two
  - Reduces leakage from O(log L) to O(log log L)
  - “Best possible” while tolerating constant-factor waste
- Brick-man: pad a bit more intelligently
  - Still reduce leakage from O(log L) to O(log log L)
  - But limit waste to 12% max, decreasing with file size
  - Details in draft-in-progress, will share on request

Encrypted file size vs padding waste

Length	Length bits	Leak bits	Length inc	Max waste
1	1	0	1	0.00%
2	2	1	1	0.00%
4	3	2	1	0.00%
8	4	2	2	11.11%
16	5	3	2	5.88%
32	6	3	4	9.09%
64	7	3	8	10.77%
128	8	3	16	11.63%
256	9	4	16	5.84%
512	10	4	32	6.04%
1024	11	4	64	6.15%
2048	12	4	128	6.20%
4096	13	4	256	6.22%
8192	14	4	512	6.24%
16384	15	4	1024	6.24%
32768	16	4	2048	6.25%
65536	17	5	2048	3.12%
131072	18	5	4096	3.12%
262144	19	5	8192	3.12%
524288	20	5	16384	3.12%
1048576	21	5	32768	3.12%
2097152	22	5	65536	3.12%
4194304	23	5	131072	3.12%
8388608	24	5	262144	3.12%

Partial-File Integrity Protection

Brought up by DKG, discussed on CFRG list.

Two motivating use-cases: (orthogonal?)

- Streaming-mode decryption (restore backup)
  - Check bytes before they leave pipe, w/o storing it all
  - Need incremental MAC+signature per chunk?
- Random-access decryption (ala Tahoe-LAFS)
  - Encryptor builds Merkle tree, stores in trailer
  - Decryptor uses to decrypt, check individual chunks
- If we support, are they mandatory? Optional?
  - Simplicity vs power vs usability?