

Workgroup: openpgp
Internet-Draft:
draft-dkg-openpgp-stateless-cli-02
Published: 6 March 2020
Intended Status: Informational
Expires: 7 September 2020
Authors: D.K. Gillmor
 ACLU

Stateless OpenPGP Command Line Interface

Abstract

This document defines a generic stateless command-line interface for dealing with OpenPGP messages, known as sop. It aims for a minimal, well-structured API covering OpenPGP object security.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 7 September 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Requirements Language](#)
 - [1.2. Terminology](#)
 - [1.3. Using sop in a Test Suite](#)
- [2. Examples](#)
- [3. Subcommands](#)
 - [3.1. version: Version Information](#)
 - [3.2. generate-key: Generate a Secret Key](#)
 - [3.3. extract-cert: Extract a Certificate from a Secret Key](#)
 - [3.4. sign: Create Detached Signatures](#)
 - [3.5. verify: Verify Detached Signatures](#)
 - [3.6. encrypt: Encrypt a Message](#)
 - [3.7. decrypt: Decrypt a Message](#)
 - [3.8. armor: Convert binary to ASCII](#)
 - [3.9. dearmor: Convert ASCII to binary](#)
 - [3.10. detach-inband-signature-and-message: split a clearsigned message](#)
- [4. Input String Types](#)
 - [4.1. DATE](#)
 - [4.2. USERID](#)
- [5. Input/Output Indirect Types](#)
 - [5.1. Special Designators for Indirect Types](#)
 - [5.2. CERTS](#)
 - [5.3. KEY](#)
 - [5.4. CIPHERTEXT](#)

- [5.5. SIGNATURES](#)
- [5.6. SESSIONKEY](#)
- [5.7. PASSWORD](#)
- [5.8. VERIFICATIONS](#)
- [5.9. DATA](#)
- [6. Failure Modes](#)
- [7. Guidance for Implementers](#)
 - [7.1. One OpenPGP Message at a Time](#)
 - [7.2. Simplified Subset of OpenPGP Message](#)
 - [7.3. Validate Signatures Only from Known Signers](#)
 - [7.4. OpenPGP inputs can be either Binary or ASCII-armored](#)
 - [7.5. Detached Signatures](#)
 - [7.6. Reliance on Supplied Certs and Keys](#)
 - [7.7. Text is always UTF-8](#)
 - [7.8. Passwords are Human-Readable](#)
 - [7.9. Be careful with Special Designators](#)
- [8. Guidance for Consumers](#)
 - [8.1. Choosing between -as=text and -as=binary](#)
 - [8.2. Special Designators and Unusual Filenames](#)
- [9. Security Considerations](#)
 - [9.1. Signature Verification](#)
 - [9.2. Compression](#)
- [10. Privacy Considerations](#)
 - [10.1. Object Security vs. Transport Security](#)
- [11. Document Considerations](#)
 - [11.1. Document History](#)

[11.2. Future Work](#)

[12. Acknowledgements](#)

[13. References](#)

[13.1. Normative References](#)

[13.2. Informative References](#)

[Author's Address](#)

1. Introduction

Different OpenPGP implementations have many different requirements, which typically break down in two main categories: key/certificate management and object security.

The purpose of this document is to provide a "stateless" interface that primarily handles the object security side of things, and assumes that secret key management and certificate management will be handled some other way.

Isolating object security from key/certificate management should make it easier to provide interoperability testing for the object security side of OpenPGP implementations, as described in [Section 1.3](#).

This document defines a generic stateless command-line interface for dealing with OpenPGP messages, known here by the placeholder `sop`. It aims for a minimal, well-structured API.

An OpenPGP implementation should not name its executable `sop` to implement this specification. It just needs to provide a program that conforms to this interface.

A `sop` implementation should leave no trace on the system, and its behavior should not be affected by anything other than command-line arguments and input.

Obviously, the user will need to manage their secret keys (and their peers' certificates) somehow, but the goal of this interface is to separate out that task from the task of interacting with OpenPGP messages.

While this document identifies a command-line interface, the rough outlines of this interface should also be amenable to relatively straightforward library implementations in different languages.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

This document uses the term "key" to refer exclusively to OpenPGP Transferable Secret Keys (see section 11.2 of [[RFC4880](#)]).

It uses the term "certificate" to refer to OpenPGP Transferable Public Key (see section 11.1 of [[RFC4880](#)]).

"Stateless" in "Stateless OpenPGP" means avoiding secret key and certificate state. The user is responsible for managing all OpenPGP certificates and secret keys themselves, and passing them to sop as needed. The user should also not be concerned that any state could affect the underlying operations.

OpenPGP revocations can have "Reason for Revocation" (section 5.2.3.23 of [[RFC4880](#)]), which can be either "soft" or "hard". The set of "soft" reasons is: "Key is superseded" and "Key is retired and no longer used". All other reasons (and revocations that do not state a reason) are "hard" revocations.

1.3. Using sop in a Test Suite

If an OpenPGP implementation provides a sop interface, it can be used to test interoperability (e.g., [[OpenPGP-Interoperability-Test-Suite](#)]).

Such an interop test suite can, for example, use custom code (*not* sop) to generate a new OpenPGP object that incorporates new primitives, and feed that object to a stable of sop implementations, to determine whether those implementations can consume the new form.

Or, the test suite can drive each sop implementation with a simple input, and observe which cryptographic primitives each implementation chooses to use as it produces output.

2. Examples

These examples show no error checking, but give a flavor of how sop might be used in practice from a shell.

The key and certificate files described in them (e.g. `alice.sec`) could be for example those found in [[I-D.draft-bre-openpgp-samples-00](#)].

```
sop generate-key "Alice Lovelace <alice@openpgp.example>" > alice.sec
sop extract-cert < alice.sec > alice.pgp

sop sign --as=text alice.sec < statement.txt > statement.txt.asc
sop verify announcement.txt.asc alice.pgp < announcement.txt

sop encrypt --sign-with=alice.sec --as=mime bob.pgp < msg.eml > encrypte
sop decrypt alice.sec < ciphertext.asc > cleartext.out
```

See [Section 6](#) for more information about errors and error handling.

3. Subcommands

`sop` uses a subcommand interface, similar to those popularized by systems like `git` and `svn`.

If the user supplies a subcommand that `sop` does not implement, it fails with `UNSUPPORTED_SUBCOMMAND`. If a `sop` implementation does not handle a supplied option for a given subcommand, it fails with `UNSUPPORTED_OPTION`.

All subcommands that produce OpenPGP material on standard output produce ASCII-armored (section 6 of [[I-D.ietf-openpgp-rfc4880bis](#)]) objects by default (except for `sop dearmor`). These subcommands have a `--no-armor` option, which causes them to produce binary OpenPGP material instead.

All subcommands that accept OpenPGP material on input should be able to accept either ASCII-armored or binary inputs (see [Section 7.4](#)) and behave accordingly.

See [Section 5](#) for details about how various forms of OpenPGP material are expected to be structured.

3.1. version: Version Information

```
sop version
```

```
*Standard Input: ignored
```

```
*Standard Output: version string
```

The version string emitted should contain the name of the `sop` implementation, followed by a single space, followed by the version number. A `sop` implementation should use a version number that

respects an established standard that is easily comparable and parsable, like [\[SEMVER\]](#).

Example:

```
$ sop version
ExampleSop 0.2.1
$
```

3.2. generate-key: Generate a Secret Key

```
sop generate-key [--no-armor] [--] [USERID...]
```

*Standard Input: ignored

*Standard Output: KEY ([Section 5.3](#))

Generate a single default OpenPGP key with zero or more User IDs.

The generated secret key SHOULD be usable for as much of the sop functionality as possible. In particular:

*It should be possible to extract an OpenPGP certificate from the KEY with `sop extract-cert`.

*The KEY should be able to create signatures (with `sop sign`) that are verifiable by using `sop verify` with the extracted certificate.

*The KEY should be able to decrypt messages (with `sop decrypt`) that are encrypted by using `sop encrypt` with the extracted certificate.

The detailed internal structure of the certificate is left to the discretion of the sop implementation.

Example:

```
$ sop generate-key 'Alice Lovelace <alice@openpgp.example>' > alice.sec
$ head -n1 < alice.sec
-----BEGIN PGP PRIVATE KEY BLOCK-----
$
```

3.3. extract-cert: Extract a Certificate from a Secret Key

```
sop extract-cert [--no-armor]
```

*Standard Input: KEY ([Section 5.3](#))

*Standard Output: CERTS ([Section 5.2](#))

Note that the resultant CERTS object will only ever contain one OpenPGP certificate, since KEY contains exactly one OpenPGP Transferable Secret Key.

Example:

```
$ sop extract-cert < alice.sec > alice.pgp
$ head -n1 < alice.pgp
-----BEGIN PGP PUBLIC KEY BLOCK-----
$
```

3.4. sign: Create Detached Signatures

```
sop sign [--no-armor]
      [--as={binary|text}] [--] KEY [KEY...]
```

*Standard Input: DATA ([Section 5.9](#))

*Standard Output: SIGNATURES ([Section 5.5](#))

Exactly one signature will be made by each supplied KEY.

--as defaults to binary. If --as=text and the input DATA is not valid UTF-8 ([Section 7.7](#)), sop sign fails with EXPECTED_TEXT.

--as=binary SHOULD result in an OpenPGP signature of type 0x00 ("Signature of a binary document"). --as=text SHOULD result in an OpenPGP signature of type 0x01 ("Signature of a canonical text document"). See section 5.2.1 of [[RFC4880](#)] for more details.

sop sign MUST NOT produce any extra signatures beyond those from KEY objects supplied on the command line.

Example:

```
$ sop sign --as=text alice.sec < message.txt > message.txt.asc
$ head -n1 < message.txt.asc
-----BEGIN PGP SIGNATURE-----
$
```

3.5. verify: Verify Detached Signatures

```
sop verify [--not-before=DATE] [--not-after=DATE]
      [--] SIGNATURES CERTS [CERTS...]
```

*Standard Input: DATA ([Section 5.9](#))

*Standard Output: VERIFICATIONS ([Section 5.8](#))

--not-before and --not-after indicate that signatures with dates outside certain range MUST NOT be considered valid.

--not-before defaults to the beginning of time. Accepts the special value - to indicate the beginning of time (i.e. no lower boundary).

--not-after defaults to the current system time (now). Accepts the special value - to indicate the end of time (i.e. no upper boundary).

sop verify only returns OK if at least one certificate included in any CERTS object made a valid signature in the range over the DATA supplied.

For details about the valid signatures, the user MUST inspect the VERIFICATIONS output.

If no CERTS are supplied, sop verify fails with MISSING_ARG.

If no valid signatures are found, sop verify fails with NO_SIGNATURE.

See [Section 9.1](#) for more details about signature verification.

Example:

(In this example, we see signature verification succeed first, and then fail on a modified version of the message.)

```
$ sop verify message.txt.asc alice.pgp < message.txt
2019-10-29T18:36:45Z EB85BB5FA33A75E15E944E63F231550C4F47E38E EB85BB5FA3
$ echo $?
0
$ tr a-z A-Z < message.txt | sop verify message.txt.asc alice.pgp
$ echo $?
3
$
```

3.6. encrypt: Encrypt a Message

```
sop encrypt [--as={binary|text|mime}]
  [--no-armor]
  [--with-password=PASSWORD...]
  [--sign-with=KEY...]
  [--] [CERTS...]
```

*Standard Input: DATA ([Section 5.9](#))

*Standard Output: CIPHERTEXT ([Section 5.4](#))

--as defaults to binary. The setting of --as corresponds to the one octet format field found in the Literal Data packet at the core of the output CIPHERTEXT. If --as is set to binary, the octet is b (0x62). If it is text, the format octet is u (0x75). If it is mime, the format octet is m (0x6d).

--with-password enables symmetric encryption (and can be used multiple times if multiple passwords are desired). If sop encrypt encounters a PASSWORD which is not a valid UTF-8 string ([Section 7.7](#)), or is otherwise not robust in its representation to humans, it fails with PASSWORD_NOT_HUMAN_READABLE. If sop encrypt sees trailing whitespace at the end of a PASSWORD, it will trim the trailing whitespace before using the password. See [Section 7.8](#) for more discussion about passwords.

--sign-with creates exactly one signature by the identified secret key (and can be used multiple times if signatures from multiple keys are desired).

If --as is set to binary, then --sign-with will sign as a binary document (OpenPGP signature type 0x00).

If --as is set to text, then --sign-with will sign as a canonical text document (OpenPGP signature type 0x01). In this case, if the input DATA is not valid UTF-8 ([Section 7.7](#)), sop encrypt fails with EXPECTED_TEXT.

sop should only be invoked with --as=mime when the input DATA is a MIME message ([RFC2045](#)). If --sign-with is supplied for such a message, then if the input data is valid UTF-8, sop SHOULD sign as a canonical text document (OpenPGP signature type 0x01). However, a MIME message itself might not be valid UTF-8, for example, if a MIME subpart contains a raw binary object. If --sign-with is supplied for input DATA that is not valid UTF-8, sop encrypt MAY sign as a binary document (OpenPGP signature type 0x00).

sop encrypt MUST NOT produce any extra signatures beyond those from KEY objects identified by --sign-with.

The resulting CIPHERTEXT should be decryptable by the secret keys corresponding to every certificate included in all CERTS, as well as each password given with --with-password.

If no CERTS or --with-password options are present, sop encrypt fails with MISSING_ARG.

If at least one of the identified certificates requires encryption to an unsupported asymmetric algorithm, sop encrypt fails with UNSUPPORTED_ASYMMETRIC_ALGO.

If at least one of the identified certificates is not encryption-capable (e.g., revoked, expired, no encryption-capable flags on primary key and valid subkeys), sop encrypt fails with CERT_CANNOT_ENCRYPT.

If sop encrypt fails for any reason, it emits no CIPHERTEXT.

Example:

(In this example, bob.bin is a file containing Bob's binary-formatted OpenPGP certificate. Alice is encrypting a message to both herself and Bob.)

```
$ sop encrypt --as=mime --sign-with=alice.key alice.asc bob.bin < messag
$ head -n1 encrypted.asc
-----BEGIN PGP MESSAGE-----
$
```

3.7. decrypt: Decrypt a Message

```
sop decrypt [--session-key-out=SESSIONKEY]
  [--with-session-key=SESSIONKEY...]
  [--with-password=PASSWORD...]
  [--verify-out=VERIFICATIONS
  [--verify-with=CERTS...]
  [--verify-not-before=DATE]
  [--verify-not-after=DATE] ]
  [--] [KEY...]
```

*Standard Input: CIPHERTEXT ([Section 5.4](#))

*Standard Output: DATA ([Section 5.9](#))

The caller can ask sop for the session key discovered during decryption by supplying the --session-key-out option. If the specified file already exists in the filesystem, sop decrypt will fail with OUTPUT_EXISTS. When decryption is successful, sop decrypt writes the discovered session key to the specified file.

--with-session-key enables decryption of the CIPHERTEXT using the session key directly against the SEIPD packet. This option can be used multiple times if several possible session keys should be tried.

--with-password enables decryption based on any SKESK (section 5.3 of [[I-D.ietf-openpgp-rfc4880bis](#)]) packets in the CIPHERTEXT. This option can be used multiple times if the user wants to try more than one password.

If `sop decrypt` tries and fails to use a supplied `PASSWORD`, and it observes that there is trailing UTF-8 whitespace at the end of the `PASSWORD`, it will retry with the trailing whitespace stripped. See [Section 7.8](#) for more discussion about passwords.

`--verify-out` produces signature verification status to the designated file. If the designated file already exists in the filesystem, `sop decrypt` will fail with `OUTPUT_EXISTS`.

The return code of `sop decrypt` is not affected by the results of signature verification. The caller **MUST** check the returned `VERIFICATIONS` to confirm signature status. An empty `VERIFICATIONS` output indicates that no valid signatures were found.

`--verify-with` identifies a set of certificates whose signatures would be acceptable for signatures over this message.

If the caller is interested in signature verification, both `--verify-out` and at least one `--verify-with` must be supplied. If only one of these arguments is supplied, `sop decrypt` fails with `INCOMPLETE_VERIFICATION`.

`--verify-not-before` and `--verify-not-after` provide a date range for acceptable signatures, by analogy with the options for `sop verify` (see [Section 3.5](#)). They should only be supplied when doing signature verification.

See [Section 9.1](#) for more details about signature verification.

If no `KEY` or `--with-password` or `--with-session-key` options are present, `sop decrypt` fails with `MISSING_ARG`.

If unable to decrypt, `sop decrypt` fails with `CANNOT_DECRYPT`.

`sop decrypt` only emits cleartext to Standard Output that was successfully decrypted.

Example:

(In this example, Alice stashes and re-uses the session key of an encrypted message.)

```
$ sop decrypt --session-key-out=session.key alice.sec < ciphertext.asc >
$ ls -l ciphertext.asc cleartext.out
-rw-r--r-- 1 user user   321 Oct 28 01:34 ciphertext.asc
-rw-r--r-- 1 user user   285 Oct 28 01:34 cleartext.out
$ sop decrypt --with-session-key=session.key < ciphertext.asc > cleartex
$ diff cleartext.out cleartext2.out
$
```

3.8. armor: Convert binary to ASCII

sop armor [--label={auto|sig|key|cert|message}]

*Standard Input: OpenPGP material (SIGNATURES, KEY, CERTS, or CIPHERTEXT)

*Standard Output: the same material with ASCII-armoring added, if not already present

The user can choose to specify the label used in the header and tail of the armoring.

The default for --label is auto, in which case, sop inspects the input and chooses the label appropriately, based on the type of the first OpenPGP packet. If the type of the first OpenPGP packet is:

*0x02 (Signature), the packet stream should be parsed as a SIGNATURES input (with Armor Header BEGIN PGP SIGNATURE).

*0x05 (Secret-Key), the packet stream should be parsed as a KEY input (with Armor Header BEGIN PGP PRIVATE KEY BLOCK).

*0x06 (Public-Key), the packet stream should be parsed as a CERTS input (with Armor Header BEGIN PGP PUBLIC KEY BLOCK).

*0x01 (Public-key Encrypted Session Key) or 0x03 (Symmetric-key Encrypted Session Key), the packet stream should be parsed as a CIPHERTEXT input (with Armor Header BEGIN PGP MESSAGE).

If the input packet stream does not match the expected sequence of packet types, sop armor fails with BAD_DATA.

Since sop armor accepts ASCII-armored input as well as binary input, this operation is idempotent on well-structured data. A caller can use this subcommand blindly ensure that any well-formed OpenPGP packet stream is 7-bit clean.

Example:

```
$ sop armor < bob.bin > bob.pgp
$ head -n1 bob.pgp
-----BEGIN PGP PUBLIC KEY BLOCK-----
$
```

3.9. dearmor: Convert ASCII to binary

sop dearmor

*Standard Input: OpenPGP material (SIGNATURES, KEY, CERTS, or CIPHERTEXT)

*Standard Output: the same material with any ASCII-armoring removed

If the input packet stream does not match any of the the expected sequence of packet types, sop dearmor fails with BAD_DATA. See also [Section 7.4](#).

Since sop dearmor accepts binary-formatted input as well as ASCII-armored input, this operation is idempotent on well-structured data. A caller can use this subcommand blindly ensure that any well-formed OpenPGP packet stream is in its standard binary representation.

Example:

```
$ sop dearmor < message.txt.asc > message.txt.sig
$
```

3.10. detach-inband-signature-and-message: split a clearsigned message

sop detach-inband-signature-and-message --signatures-out=SIGNATURES

*Standard Input: DATA (clearsigned message)

*Standard Output: DATA (the message without the cleartext signature framework)

In some contexts, the user may encounter a clearsigned ("inline PGP") message (section 7 of [RFC4880](#)) rather than a message and its detached signature. This subcommand takes such a clearsigned message on standard input, and splits it into:

*the potentially signed material on standard output, and

*a detached signature block to the destination identified by --signatures-to

Note that no cryptographic verification of the signatures is done by this subcommand. Once the clearsigned message is separated, verification of the detached signature can be done with sop verify.

If no --signatures-to is supplied, sop detach-inband-signature-and-message fails with MISSING_ARG.

Note that the signature block in a clearsigned message may contain multiple signatures. All signatures found in the signature block will be emitted to the `--signatures-to` destination.

The message body in the clearsigned message will be dash-escaped on standard input (see section 7.1 of [[RFC4880](#)]). The output of `sop detach-inband-signature-and-message` will have dash-escaping removed.

If the input DATA contains no clearsigned message, `sop detach-inband-signature-and-message` fails with `BAD_DATA`. If the input DATA contains more than one clearsigned message, `sop detach-inband-signature-and-message` also fails with `BAD_DATA`. A `sop` implementation MAY accept (and discard) leading and trailing data around the inline PGP clearsigned message.

If the file designated by `--signatures-to` already exists in the filesystem, `sop detach-inband-signature-and-message` will fail with `OUTPUT_EXISTS`.

Example:

```
$ sop detach-inband-signature-and-message --signature-out=Release.pgp <
$ sop verify Release.pgp archive-keyring.pgp < Release
$
```

4. Input String Types

Some material is passed to `sop` directly as a string on the command line.

4.1. DATE

An ISO-8601 formatted timestamp with time zone, or the special value `now` to indicate the current system time.

Examples:

```
now
2019-10-29T12:11:04+00:00
2019-10-24T23:48:29Z
20191029T121104Z
```

In some cases where used to specify lower and upper boundaries, a DATE value can be set to `-` to indicate "no time limit".

A flexible implementation of `sop` MAY accept date inputs in other unambiguous forms.

Note that whenever `sop` emits a timestamp (e.g. in [Section 5.8](#)) it MUST produce only a UTC-based ISO-8601 compliant representation.

4.2. USERID

This is an arbitrary UTF-8 string ([Section 7.7](#)). By convention, most User IDs are of the form Display Name <email.address@example.com>, but they do not need to be.

5. Input/Output Indirect Types

Some material is passed to sop indirectly, typically by referring to a filename containing the data in question. This type of data may also be passed to sop on Standard Input, or delivered by sop to Standard Output.

If any input data is specified explicitly to be read from a file that does not exist, sop will fail with MISSING_INPUT.

If any input data does not meet the requirements described below, sop will fail with BAD_DATA.

5.1. Special Designators for Indirect Types

An indirect argument or parameter that starts with "@" (COMMERCIAL AT, U+0040) is not treated as a filename, but is reserved for special handling, based on the prefix that follows the @. We describe two of those prefixes (@ENV: and @FD:) here. A sop implementation that receives such a special designator but does not know how to handle a given prefix in that context MUST fail with UNSUPPORTED_SPECIAL_PREFIX.

If the filename for any indirect material used as input has the special form @ENV:xxx, then contents of environment variable \$xxx is used instead of looking in the filesystem. @ENV is for input only: if the prefix @ENV: is used for any output argument, sop fails with UNSUPPORTED_SPECIAL_PREFIX.

If the filename for any indirect material used as either input or output has the special form @FD:nnn where nnn is a decimal integer, then the associated data is read from file descriptor nnn.

See [Section 7.9](#) for more details about safe handling of these special designators.

5.2. CERTS

One or more OpenPGP certificates (section 11.1 of [[I-D.ietf-openpgp-rfc4880bis](#)]), aka "Transferable Public Key". May be armored (see [Section 7.4](#)).

Although some existing workflows may prefer to use one CERTS object with multiple certificates in it (a "keyring"), supplying exactly

one certificate per CERTS input will make error reporting clearer and easier.

5.3. KEY

Exactly one OpenPGP Transferable Secret Key (section 11.2 of [[I-D.ietf-openpgp-rfc4880bis](#)]). May be armored (see [Section 7.4](#)).

Secret key material should be in cleartext (that is, it should not be locked with a password). If the secret key material is locked with a password, sop may fail with error KEY_IS_PROTECTED.

5.4. CIPHERTEXT

sop accepts only a restricted subset of the arbitrarily-nested grammar allowed by the OpenPGP Messages definition (section 11.3 of [[I-D.ietf-openpgp-rfc4880bis](#)]).

In particular, it accepts and generates only:

An OpenPGP message, consisting of a sequence of PKESKs (section 5.1 of [[I-D.ietf-openpgp-rfc4880bis](#)]) and SKESKs (section 5.3 of [[I-D.ietf-openpgp-rfc4880bis](#)]), followed by one SEIPD (section 5.14 of [[I-D.ietf-openpgp-rfc4880bis](#)]).

The SEIPD can decrypt into one of two things:

- *"Maybe Signed Data" (see below), or

- *Compressed data packet that contains "Maybe Signed Data"

"Maybe Signed Data" is a sequence of:

- *N (zero or more) one-pass signature packets, followed by

- *zero or more signature packets, followed by

- *one Literal data packet, followed by

- *N signature packets (corresponding to the outer one-pass signatures packets)

FIXME: does any tool do compression inside signing? Do we need to handle that?

May be armored (see [Section 7.4](#)).

5.5. SIGNATURES

One or more OpenPGP Signature packets. May be armored (see [Section 7.4](#)).

5.6. SESSIONKEY

This documentation uses the GnuPG defacto ASCII representation:

ALGONUM:HEXKEY

where ALGONUM is the decimal value associated with the OpenPGP Symmetric Key Algorithms (section 9.3 of [[I-D.ietf-openpgp-rfc4880bis](#)]) and HEXKEY is the hexadecimal representation of the binary key.

Example AES-256 session key:

9:FCA4BEAF687F48059CACC14FB019125CD57392BAB7037C707835925CBF9F7BCD

5.7. PASSWORD

This is expected to be a UTF-8 string ([Section 7.7](#)), but for sop decrypt, any bytestring that the user supplies will be accepted. Note the details in sop encrypt and sop decrypt about trailing whitespace!

See also [Section 7.8](#) for more discussion.

5.8. VERIFICATIONS

One line per successful signature verification. Each line has three structured fields delimited by a single space, followed by arbitrary text to the end of the line that forms a message describing the verification.

*ISO-8601 UTC datestamp

*Fingerprint of the signing key (may be a subkey)

*Fingerprint of primary key of signing certificate (if signed by primary key, same as the previous field)

*message describing the verification (free form)

Note that while [Section 4.1](#) permits a sop implementation to accept other unambiguous date representations, its date output here MUST be a strict ISO-8601 UTC date timestamp. In particular:

- *the date and time fields MUST be separated by T, not by whitespace, since whitespace is used as a delimiter

- *the time MUST be emitted in UTC, with the explicit suffix Z

Example:

```
2019-10-24T23:48:29Z C90E6D36200A1B922A1509E77618196529AE5FF8 C4BC2DDB38
```

5.9. DATA

Cleartext, arbitrary data. This is either a bytestream or UTF-8 text.

It MUST only be UTF-8 text in the case of input supplied to sop sign --as=text or sop encrypt --as={mime|text}. If sop receives DATA containing non-UTF-8 octets in this case, it will fail (see [Section 7.7](#)) with EXPECTED_TEXT.

6. Failure Modes

sop return codes have both mnemonics and numeric values.

When sop succeeds, it will return 0 (OK) and emit nothing to Standard Error. When sop fails, it fails with a non-zero return code, and emits one or more warning messages on Standard Error. Known return codes include:

Value	Mnemonic	Meaning
0	OK	Success
3	NO_SIGNATURE	No acceptable signatures found (sop verify)
13	UNSUPPORTED_ASYMMETRIC_ALGO	Asymmetric algorithm unsupported (sop encrypt)
17	CERT_CANNOT_ENCRYPT	Certificate not encryption-capable (e.g., expired, revoked, unacceptable usage flags) (sop encrypt)
19	MISSING_ARG	Missing required argument
23	INCOMPLETE_VERIFICATION	Incomplete verification instructions (sop decrypt)
29	CANNOT_DECRYPT	Unable to decrypt (sop decrypt)
31	PASSWORD_NOT_HUMAN_READABLE	Non-UTF-8 or otherwise unreliable password (sop encrypt)
37	UNSUPPORTED_OPTION	Unsupported option

Value	Mnemonic	Meaning
41	BAD_DATA	Invalid data type (no secret key where KEY expected, etc)
53	EXPECTED_TEXT	Non-text input where text expected
59	OUTPUT_EXISTS	Output file already exists
61	MISSING_INPUT	Input file does not exist
67	KEY_IS_PROTECTED	A KEY input is protected (locked) with a password, and sop cannot unlock it
69	UNSUPPORTED_SUBCOMMAND	Unsupported subcommand
71	UNSUPPORTED_SPECIAL_PREFIX	An indirect parameter is a special designator (it starts with @) but sop does not know how to handle the prefix
73	AMBIGUOUS_INPUT	A indirect input parameter is a special designator (it starts with @), and a filename matching the designator is actually present

Table 1

If a sop implementation fails in some way not contemplated by this document, it MAY return any non-zero error code, not only those listed above.

7. Guidance for Implementers

sop uses a few assumptions that implementers might want to consider.

7.1. One OpenPGP Message at a Time

sop is intended to be a simple tool that operates on one OpenPGP object at a time. It should be composable, if you want to use it to deal with multiple OpenPGP objects.

FIXME: discuss what this means for streaming. The stdio interface doesn't necessarily imply streamed output.

7.2. Simplified Subset of OpenPGP Message

While the formal grammar for OpenPGP Message is arbitrarily nestable, sop constrains itself to what it sees as a single "layer" (see [Section 5.4](#)).

This is a deliberate choice, because it is what most consumers expect. Also, if an arbitrarily-nested structure is parsed with a recursive algorithm, this risks a denial of service vulnerability.

sop intends to be implementable with a parser that defensively declines to do recursive descent into an OpenPGP Message.

Note that an implementation of sop decrypt MAY choose to handle more complex structures, but if it does, it should document the other structures it handles and why it chooses to do so. We can use such documentation to improve future versions of this spec.

7.3. Validate Signatures Only from Known Signers

There are generally only a few signers who are relevant for a given OpenPGP message. When verifying signatures, sop expects that the caller can identify those relevant signers ahead of time.

7.4. OpenPGP inputs can be either Binary or ASCII-armored

OpenPGP material on input can be in either ASCII-armored or binary form. This is a deliberate choice because there are typical scenarios where the program can't predict which form will appear. Expecting the caller of sop to detect the form and adjust accordingly seems both redundant and error-prone.

The simple way to detect possible ASCII-armoring is to see whether the high bit of the first octet is set: section 4.2 of [\[RFC4880\]](#) indicates that bit 7 is always one in the first octet of an OpenPGP packet. In standard ASCII-armor, the first character is "-" (HYPHEN-MINUS, U+002D), so the high bit should be cleared.

When considering an input as ASCII-armored OpenPGP material, sop MAY reject an input based on any of the following variations (see section 6.2 of [\[RFC4880\]](#) for precise definitions):

- *An unknown Armor Header Line
- *Any text before the Armor Header Line
- *Malformed lines in the Armor Headers section
- *Any non-whitespace data after the Armor Tail
- *Any Radix-64 encoded line with more than 76 characters
- *Invalid characters in the Radix-64-encoded data
- *An invalid Armor Checksum
- *A mismatch between the Armor Header Line and the Armor Tail

For robustness, sop SHOULD be willing to ignore whitespace after the Armor Tail.

When considering OpenPGP material as input, regardless of whether it is ASCII-armored or binary, sop SHOULD reject any material that doesn't produce a valid stream of OpenPGP packets. For example, sop SHOULD raise an error if an OpenPGP packet header is malformed, or if there is trailing garbage after the end of a packet.

For a given type of OpenPGP input material (i.e., SIGNATURES, CERTS, KEY, or CIPHERTEXT), sop SHOULD also reject any input that does not conform to the expected packet stream. See [Section 5](#) for the expected packet stream for different types.

7.5. Detached Signatures

sop deals with detached signatures as the baseline form of OpenPGP signatures.

The primary alternative to detached signatures is inline signatures, but handling an inline signature requires parsing to delimit the multiple parts of the document, including at least:

- *any preamble before the message
- *the inline message header (delimiter line, OpenPGP headers)
- *the message itself
- *the divider between the message and the signature (including any OpenPGP headers there)
- *the signature
- *the divider that terminates the signature
- *any suffix after the signature

Note also that the preamble or the suffix might be arbitrary text, and might themselves contain OpenPGP messages (whether signatures or otherwise).

If the parser that does this split differs in any way from the parser that does the verification, or parts of the message are confused, it would be possible to produce a verification status and an actual signed message that don't correspond to one another.

Blurred boundary problems like this can produce ugly attacks similar to those found in [\[EFAIL\]](#).

7.6. Reliance on Supplied Certs and Keys

A truly stateless implementation may find that it spends more time validating the internal consistency of certificates and keys than it does on the actual object security operations.

For performance reasons, an implementation may choose to ignore validation on certificate and key material supplied to it. The security implications of doing so depend on how the certs and keys are managed outside of sop.

7.7. Text is always UTF-8

Various places in this specification require UTF-8 [[RFC3629](#)] when encoding text. sop implementations SHOULD NOT consider textual data in any other character encoding.

OpenPGP Implementations MUST already handle UTF-8, because various parts of [[RFC4880](#)] require it, including:

- *User ID
- *Notation name
- *Reason for revocation
- *ASCII-armor Comment: header

Dealing with messages in other charsets leads to weird security failures like [[Charset-Switching](#)], especially when the charset indication is not covered by any sort of cryptographic integrity check. Restricting textual data to UTF-8 universally across the OpenPGP ecosystem eliminates any such risk without losing functionality, since UTF-8 can encode all known characters.

7.8. Passwords are Human-Readable

Passwords are generally expected to be human-readable, as they are typically recorded and transmitted as human-visible, human-transferable strings. However, they are used in the OpenPGP protocol as bytestrings, so ensuring that there is a reliable bidirectional mapping between strings and bytes. The maximally robust behavior here is for sop encrypt to constrain the choice of passwords to strings that have such a mapping, and for sop decrypt to try multiple plausible versions of any supplied PASSWORD.

When generating material based on a password, sop encrypt enforces that the password is actually meaningfully human-transferable (requiring UTF-8, trimming trailing whitespace). Some sop encrypt

implementations may make even more strict requirements on input to ensure that they are transferable between humans in a robust way.

For example, a more strict sop encrypt MAY also:

- *forbid leading whitespace

- *forbid non-printing characters other than SPACE (U+0020), such as ZERO WIDTH NON-JOINER (U+200C) or TAB (U+0009)

- *require the password to be in Unicode Normal Form C ([\[UNICODE-NORMALIZATION\]](#))

Violations of these more-strict policies SHOULD result in an error of `PASSWORD_NOT_HUMAN_READABLE`.

A sop encrypt implementation typically SHOULD NOT attempt enforce a minimum "password strength", but in the event that some implementation does, it MUST NOT represent a weak password with `PASSWORD_NOT_HUMAN_READABLE`.

When sop decrypt receives a `PASSWORD` input, it sees it as a bytestring. If the bytestring fails to work as a password, but ends in UTF-8 whitespace, it will try again with the trailing whitespace removed. This handles a common pattern of using a file with a final newline, for example. The pattern here is one of robustness in the face of typical errors in human-transferred textual data.

A more robust sop decrypt implementation that finds neither of the above two attempts work for a given `PASSWORD` MAY try additional variations if they produce a different bytestring, such as:

- *trimming any leading whitespace, if discovered

- *trimming any internal non-printable characters other than SPACE (U+0020)

- *converting the supplied `PASSWORD` into Unicode Normal Form C ([\[UNICODE-NORMALIZATION\]](#))

A sop decrypt implementation that stages multiple decryption attempts like this SHOULD consider the computational resources consumed by each attempt, to avoid presenting an attack surface for resource exhaustion in the face of a non-standard `PASSWORD` input.

7.9. Be careful with Special Designators

As documented in [Section 5.1](#), special designators for indirect inputs like `@ENV:` and `@FD:` (and indirect outputs using `@FD:`) warrant some special/cautious handling.

For one thing, it's conceivable that the filesystem could contain a file with these literal names. If sop receives an indirect output parameter that starts with an "@" (COMMERCIAL AT, U+0040) it MUST NOT write to the filesystem for that parameter. A sop implementation that receives such a parameter as input MAY test for the presence of such a file in the filesystem and fail with `AMBIGUOUS_INPUT` to warn the user of the ambiguity and possible confusion.

These special designators are likely to be used to pass sensitive data (like secret key material or passwords) so that it doesn't need to touch the filesystem. Given this sensitivity, sop should be careful with such an input, and minimize its leakage to other processes. In particular, sop SHOULD NOT leak any environment variable identified by `@ENV:` or file descriptor identified by `@FD:` to any subprocess unless the subprocess specifically needs access to that data.

8. Guidance for Consumers

While sop is originally conceived of as an interface for interoperability testing, it's conceivable that an application that uses OpenPGP for object security would want to use it.

FIXME: more guidance for how to use such a tool safely and efficiently goes here.

FIXME: if an encrypted OpenPGP message arrives without metadata, it is difficult to know which signers to consider when decrypting. How do we do this efficiently without invoking `sop decrypt` twice, once without `--verify-*` and again with the expected identity material?

8.1. Choosing between `-as=text` and `-as=binary`

A program that invokes sop to generate an OpenPGP signature typically needs to decide whether it is making a text or binary signature.

By default, sop will make a binary signature. The caller of `sop sign` should choose `--as=text` only when it knows that: - the data being signed is in fact textual, and encoded in UTF-8, and - the signed data might be transmitted to the recipient (the verifier of the signature) over a channel that has the propensity to transform line-endings.

Examples of such channels include FTP ([[RFC0959](#)]) and SMTP ([[RFC5321](#)]).

8.2. Special Designators and Unusual Filenames

In some cases, a user of `sop` might want to pass all the files in a given directory as positional parameters (e.g., a list of CERTS files to test a signature against).

If one of the files has a name that starts with `--`, it might be confused by `sop` for an option. If one of the files has a name that starts with `@`, it might be confused by `sop` as a special designator ([Section 5.1](#)).

If the user wants to deliberately refer to such an ambiguously-named file in the filesystem, they should prefix the filename with `./` or use an absolute path.

Any specific `@FD:` special designator SHOULD NOT be supplied more than once to an invocation of `sop`. If a `sop` invocation sees multiple copies of a specific `@FD:n` input (e.g., `sop sign @FD:3 @FD:3`), it MAY fail with `MISSING_INPUT` even if file descriptor 3 contains a valid KEY, because the bytestream for the KEY was consumed by the first argument. Doubling up on the same `@FD:` for output (e.g., `sop decrypt --session-key-out=@FD:3 --verify-out=@FD:3`) also results in an ambiguous data stream.

9. Security Considerations

The OpenPGP object security model is typically used for confidentiality and authenticity purposes.

9.1. Signature Verification

In many contexts, an OpenPGP signature is verified to prove the origin and integrity of an underlying object.

When `sop` checks a signature (e.g. via `sop verify` or `sop decrypt --verify-with`), it MUST NOT consider it to be verified unless all of these conditions are met:

- *The signature must be made by a signing-capable public key that is present in one of the supplied certificates
- *The certificate and signing subkey must have been created before or at the signature time
- *The certificate and signing subkey must not have been expired at the signature time
- *The certificate and signing subkey must not be revoked with a "hard" revocation

- *If the certificate or signing subkey is revoked with a "soft" revocation, then the signature time must predate the revocation
- *The signing subkey must be properly bound to the primary key, and cross-signed
- *The signature (and any dependent signature, such as the cross-sig or subkey binding signatures) must be made with strong cryptographic algorithms (e.g., not MD5 or a 1024-bit RSA key)

Implementers MAY also consider other factors in addition to the origin and authenticity, including application-specific information.

For example, consider the application domain of checking software updates. If software package Foo version 13.3.2 was signed on 2019-10-04, and the user receives a copy of Foo version 12.4.8 that was signed on 2019-10-16, it may be authentic and have a more recent signature date. But it is not an upgrade ($12.4.8 < 13.3.2$), and therefore it should not be applied automatically.

In such cases, it is critical that the application confirms that the other information verified is *also* protected by the relevant OpenPGP signature.

Signature validity is a complex topic (see for example the discussion at [[DISPLAYING-SIGNATURES](#)]), and this documentation cannot list all possible details.

9.2. Compression

The interface as currently specified does not allow for control of compression. Compressing and encrypting data that may contain both attacker-supplied material and sensitive material could leak information about the sensitive material (see the CRIME attack).

Unless an application knows for sure that no attacker-supplied material is present in the input, it should not compress during encryption.

10. Privacy Considerations

Material produced by `sop encrypt` may be placed on an untrusted machine (e.g., sent through the public SMTP network). That material may contain metadata that leaks associational information (e.g., recipient identifiers in PKESK packets (section 5.1 of [[I-D.ietf-openpgp-rfc4880bis](#)])). **FIXME:** document things like PURBs and --hidden-recipient)

10.1. Object Security vs. Transport Security

OpenPGP offers an object security model, but says little to nothing about how the secured objects get to the relevant parties.

When sending or receiving OpenPGP material, the implementer should consider what privacy leakage is implicit with the transport.

11. Document Considerations

[RFC Editor: please remove this section before publication]

This document is currently edited as markdown. Minor editorial changes can be suggested via merge requests at <https://gitlab.com/dkg/openpgp-stateless-cli> or by e-mail to the authors. Please direct all significant commentary to the public IETF OpenPGP mailing list: openpgp@ietf.org

11.1. Document History

substantive changes between -01 and -02:

- *Added mnemonics for return codes
- *decrypt should fail when asked to output to a pre-existing file
- *Removed superfluous --armor option
- *Much more specific about what armor --label=auto should do
- *armor and dearmor are now fully idempotent, but work only well-formed OpenPGP streams
- *Dropped armor --allow-nested
- *Specified what encrypt --as= means
- *New error code: KEY_IS_PROTECTED
- *Documented expectations around human-readable, human-transferable passwords
- *New subcommand: detach-inband-signature-and-message
- *More specific guidance about special designators like @FD: and @ENV:, including new error codes UNSUPPORTED_SPECIAL_PREFIX and AMBIGUOUS_INPUT

substantive changes between -00 and -01:

- *Changed generate subcommand to generate-key

- *Changed convert subcommand to extract-cert
- *Added "Input String Types" section as distinct from indirect I/O
- *Made implicit arguments potentially explicit (e.g. sop armor --label=auto)
- *Added --allow-nested to sop armor to make it idempotent by default
- *Added fingerprint of signing (sub)key to VERIFICATIONS output
- *Dropped --mode and --session-key arguments for sop encrypt (no plausible use, not needed for interop)
- *Added --with-session-key argument to sop decrypt to allow for session-key-based decryption
- *Added examples to each subcommand
- *More detailed error codes for sop encrypt
- *Move from CERT to CERTS (each CERTS argument might contain multiple certificates)

11.2. Future Work

- *certificate transformation into popular publication forms:
 - WKD
 - DANE OPENPGPKEY
 - Autocrypt
- *sop encrypt - specify compression? (see [Section 9.2](#))
- *sop encrypt - specify padding policy/mechanism?
- *sop decrypt - how can it more safely handle zip bombs?
- *sop decrypt - what should it do when encountering weakly-encrypted (or unencrypted) input?
- *sop encrypt - minimize metadata (e.g. --throw-keyids)?
- *handling secret keys that are locked with passwords?
- *specify an error if a DATE arrives as input without a time zone?

*add considerations about what it means for armored CERTS to contain multiple certificates - multiple armorings? one big blob?

*do we need an interface or option (for performance?) with the semantics that sop doesn't validate certificates internally, it just accepts whatever's given as legit data? (see [Section 7.6](#))

*do we need to be able to assemble a clearsigned message? I'd rather not, given the additional complications.

12. Acknowledgements

This work was inspired by Justus Winter's [[OpenPGP-Interoperability-Test-Suite](#)].

The following people contributed helpful feedback and considerations to this draft, but are not responsible for its problems:

*Allan Nordhoej

*Antoine Beaupre

*Edwin Taylor

*Jameson Rollins

*Justus Winter

*Vincent Breitmoser

13. References

13.1. Normative References

[I-D.ietf-openpgp-rfc4880bis]

Koch, W., carlson, b., Tse, R., Atkins, D., and D. Gillmor, "OpenPGP Message Format", Work in Progress, Internet-Draft, draft-ietf-openpgp-rfc4880bis-08, 6 September 2019, <<http://www.ietf.org/internet-drafts/draft-ietf-openpgp-rfc4880bis-08.txt>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.

[RFC4880]

Callas, J., Donnerhackle, L., Finney, H., Shaw, D., and R. Thayer, "OpenPGP Message Format", RFC 4880, DOI 10.17487/RFC4880, November 2007, <<https://www.rfc-editor.org/info/rfc4880>>.

[RFC8174]

Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

13.2. Informative References

[Charset-Switching]

Gillmor, D.K., "Inline PGP Considered Harmful", 24 February 2014, <<https://dkg.fifthhorseman.net/notes/inline-pgp-harmful/>>.

[DISPLAYING-SIGNATURES]

Brunschwig, P., "On Displaying Signatures", n.d., <https://admin.hostpoint.ch/pipermail/enigmail-users_enigmail.net/2017-November/004683.html>.

[EFAIL]

Poddebniak, D. and C. Dresen, "Efail: Breaking S/MIME and OpenPGP Email Encryption using Exfiltration Channels", n.d., <<https://efail.de>>.

[I-D.draft-bre-openpgp-samples-00]

Einarsson, B., juga, j., and D. Gillmor, "OpenPGP Example Keys and Certificates", Work in Progress, Internet-Draft, draft-bre-openpgp-samples-00, 15 October 2019, <<http://www.ietf.org/internet-drafts/draft-bre-openpgp-samples-00.txt>>.

[OpenPGP-Interoperability-Test-Suite]

"OpenPGP Interoperability Test Suite", 28 October 2019, <<https://tests.sequoia-pgp.org/>>.

[RFC0959]

Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, DOI 10.17487/RFC0959, October 1985, <<https://www.rfc-editor.org/info/rfc959>>.

[RFC2045]

Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, DOI 10.17487/RFC2045, November 1996, <<https://www.rfc-editor.org/info/rfc2045>>.

[RFC5321]

Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <<https://www.rfc-editor.org/info/rfc5321>>.

[SEMVER]

Preston-Werner, T., "Semantic Versioning 2.0.0", 18 June 2013, <<https://semver.org/>>.

[UNICODE-NORMALIZATION]

Whistler, K., "Unicode Normalization Forms",
4 February 2019, <<https://unicode.org/reports/tr15/>>.

Author's Address

Daniel Kahn Gillmor
American Civil Liberties Union
125 Broad St.
New York, NY, 10004
United States of America

Email: dkg@fifthhorseman.net