
PyNaCl

Release 1.4.0.dev1

Dec 25, 2018

Contents

1	Features	3
2	Contents	5
2.1	Public Key Encryption	5
2.2	Secret Key Encryption	9
2.3	Digital Signatures	12
2.4	Hashing	16
2.5	Password hashing	19
3	Support Features	25
3.1	Encoders	25
3.2	Exceptions	26
3.3	Utilities	27
3.4	nacl.hash	27
3.5	nacl.pwhash	28
3.6	nacl.hashlib	33
3.7	Installation	34
3.8	Doing A Release	35
3.9	Reference vectors	35
3.10	Changelog	49
3.11	Indices and tables	51
	Bibliography	53
	Python Module Index	55

PyNaCl is a Python binding to [libsodium](#), which is a fork of the [Networking and Cryptography library](#). These libraries have a stated goal of improving usability, security and speed. It supports Python 2.7 and 3.4+ as well as PyPy 2.6+.

CHAPTER 1

Features

- Digital signatures
- Secret-key encryption
- Public-key encryption
- Hashing and message authentication
- Password based key derivation and password hashing

2.1 Public Key Encryption

Imagine Alice wants something valuable shipped to her. Because it's valuable, she wants to make sure it arrives securely (i.e. hasn't been opened or tampered with) and that it's not a forgery (i.e. it's actually from the sender she's expecting it to be from and nobody's pulling the old switcheroo).

One way she can do this is by providing the sender (let's call him Bob) with a high-security box of her choosing. She provides Bob with this box, and something else: a padlock, but a padlock without a key. Alice is keeping that key all to herself. Bob can put items in the box then put the padlock onto it. But once the padlock snaps shut, the box cannot be opened by anyone who doesn't have Alice's private key.

Here's the twist though: Bob also puts a padlock onto the box. This padlock uses a key Bob has published to the world, such that if you have one of Bob's keys, you know a box came from him because Bob's keys will open Bob's padlocks (let's imagine a world where padlocks cannot be forged even if you know the key). Bob then sends the box to Alice.

In order for Alice to open the box, she needs two keys: her private key that opens her own padlock, and Bob's well-known key. If Bob's key doesn't open the second padlock, then Alice knows that this is not the box she was expecting from Bob, it's a forgery.

This bidirectional guarantee around identity is known as mutual authentication.

2.1.1 Examples

nacl.public.Box

The `Box` class uses the given public and private (secret) keys to derive a shared key, which is used with the nonce given to encrypt the given messages and to decrypt the given ciphertexts. The same shared key will be generated from both pairing of keys, so given two keypairs belonging to Alice (`pkalice`, `skalice`) and Bob (`pkbob`, `skbob`), the key derived from (`pkalice`, `skbob`) will equal that from (`pkbob`, `skalice`).

This is how the system works:

```
import nacl.utils
from nacl.public import PrivateKey, Box

# Generate Bob's private key, which must be kept secret
skbob = PrivateKey.generate()

# Bob's public key can be given to anyone wishing to send
#   Bob an encrypted message
pkbob = skbob.public_key

# Alice does the same and then Alice and Bob exchange public keys
skalice = PrivateKey.generate()
pkalice = skalice.public_key

# Bob wishes to send Alice an encrypted message so Bob must make a Box with
#   his private key and Alice's public key
bob_box = Box(skbob, pkalice)

# This is our message to send, it must be a bytestring as Box will treat it
#   as just a binary blob of data.
message = b"Kill all humans"
```

PyNaCl can automatically generate a random nonce for us, making the encryption very simple:

```
# Encrypt our message, it will be exactly 40 bytes longer than the
#   original message as it stores authentication information and the
#   nonce alongside it.
encrypted = bob_box.encrypt(message)
```

However, if we need to use an explicit nonce, it can be passed along with the message:

```
# This is a nonce, it MUST only be used once, but it is not considered
#   secret and can be transmitted or stored alongside the ciphertext. A
#   good source of nonces are just sequences of 24 random bytes.
nonce = nacl.utils.random(Box.NONCE_SIZE)

encrypted = bob_box.encrypt(message, nonce)
```

Finally, the message is decrypted (regardless of how the nonce was generated):

```
# Alice creates a second box with her private key to decrypt the message
alice_box = Box(skalice, pkbob)

# Decrypt our message, an exception will be raised if the encryption was
#   tampered with or there was otherwise an error.
plaintext = alice_box.decrypt(encrypted)
print(plaintext.decode('utf-8'))
```

```
Kill all humans
```

nacl.public.SealedBox

The `SealedBox` class encrypts messages addressed to a specified key-pair by using ephemeral sender's keypairs, which will be discarded just after encrypting a single plaintext message.

This kind of construction allows sending messages, which only the recipient can decrypt without providing any kind of cryptographic proof of sender's authorship.

Warning: By design, the recipient will have no means to trace the ciphertext to a known author, since the sending keypair itself is not bound to any sender's identity, and the sender herself will not be able to decrypt the ciphertext she just created, since the private part of the key cannot be recovered after use.

This is how the system works:

```
import nacl.utils
from nacl.public import PrivateKey, SealedBox

# Generate Bob's private key, as we've done in the Box example
skbob = PrivateKey.generate()
pkbob = skbob.public_key

# Alice wishes to send an encrypted message to Bob,
# but prefers the message to be untraceable
sealed_box = SealedBox(pkbob)

# This is Alice's message
message = b"Kill all kittens"

# Encrypt the message, it will carry the ephemeral key public part
# to let Bob decrypt it
encrypted = sealed_box.encrypt(message)
```

Now, Bob wants to read the secret message he just received; therefore he must create a `SealedBox` using his own private key:

```
unseal_box = SealedBox(skbob)
# decrypt the received message
plaintext = unseal_box.decrypt(encrypted)
print(plaintext.decode('utf-8'))
```

```
Kill all kittens
```

2.1.2 Reference

class `nacl.public.PublicKey` (*public_key, encoder*)

The public key counterpart to an `Curve25519 PrivateKey` for encrypting messages.

Parameters

- **public_key** (*bytes*) – Encoded `Curve25519` public key.
- **encoder** – A class that is able to decode the `public_key`.

class `nacl.public.PrivateKey` (*private_key, encoder*)

Private key for decrypting messages using the `Curve25519` algorithm.

Warning: This **must** be protected and remain secret. Anyone who knows the value of your `PrivateKey` can decrypt any message encrypted by the corresponding `PublicKey`

Parameters

- **private_key** (*bytes*) – The private key used to decrypt messages.
- **encoder** – A class that is able to decode the *private_key*.

public_key

An instance of *PublicKey* that corresponds with the private key.

classmethod generate ()

Generates a random *PrivateKey* object

Returns An instance of *PrivateKey*.

class nacl.public.Box (private_key, public_key)

The Box class boxes and unboxes messages between a pair of keys

The ciphertexts generated by *Box* include a 16 byte authenticator which is checked as part of the decryption. An invalid authenticator will cause the decrypt function to raise an exception. The authenticator is not a signature. Once you've decrypted the message you've demonstrated the ability to create arbitrary valid message, so messages you send are repudiable. For non-repudiable messages, sign them after encryption.

Parameters

- **private_key** – An instance of *PrivateKey* used to encrypt and decrypt messages
- **public_key** – An instance of *PublicKey* used to encrypt and decrypt messages

classmethod decode (encoded, encoder)

Decodes a serialized *Box*.

Returns An instance of *Box*.

encrypt (plaintext, nonce, encoder)

Encrypts the plaintext message using the given *nonce* (or generates one randomly if omitted) and returns the ciphertext encoded with the encoder.

Warning: It is **VITALLY** important that the nonce is a nonce, i.e. it is a number used only once for any given key. If you fail to do this, you compromise the privacy of the messages encrypted.

Parameters

- **plaintext** (*bytes*) – The plaintext message to encrypt.
- **nonce** (*bytes*) – The nonce to use in the encryption.
- **encoder** – A class that is able to decode the ciphertext.

Returns An instance of *EncryptedMessage*.

decrypt (ciphertext, nonce, encoder)

Decrypts the ciphertext using the *nonce* (explicitly, when passed as a parameter or implicitly, when omitted, as part of the ciphertext) and returns the plaintext message.

Parameters

- **ciphertext** (*bytes*) – The encrypted message to decrypt.
- **nonce** (*bytes*) – The nonce to use in the decryption.
- **encoder** – A class that is able to decode the plaintext.

Return bytes The decrypted plaintext.

shared_key()

Returns the Curve25519 shared secret, that can then be used as a key in other symmetric ciphers.

Warning: It is **VITALLY** important that you use a nonce with your symmetric cipher. If you fail to do this, you compromise the privacy of the messages encrypted. Ensure that the key length of your cipher is 32 bytes.

Return bytes The shared secret.

class nacl.public.**SealedBox**(*receiver_key*)

The SealedBox class can box and unbox messages sent to a receiver key using an ephemeral sending keypair.

encrypt(*plaintext*, *encoder*)

Encrypt the message using a Box constructed from an ephemeral key-pair and the receiver key.

The public part of the ephemeral key-pair will be enclosed in the returned ciphertext.

The private part of the ephemeral key-pair will be scrubbed before returning the ciphertext, therefore, the sender will not be able to decrypt the message.

Parameters

- **plaintext** (*bytes*) – The plaintext message to encrypt.
- **encoder** – A class that is able to decode the ciphertext.

Return bytes The public part of the ephemeral keypair, followed by the encrypted ciphertext

decrypt(*ciphertext*, *encoder*)

Decrypt the message using a Box constructed from the receiver key and the ephemeral key enclosed in the ciphertext.

Parameters

- **ciphertext** (*bytes*) – The ciphertext message to decrypt.
- **encoder** – A class that is able to decode the ciphertext.

Return bytes The decrypted message

Algorithm

- **Public Keys:** Curve25519 high-speed elliptic curve cryptography

2.2 Secret Key Encryption

Secret key encryption (also called symmetric key encryption) is analogous to a safe. You can store something secret through it and anyone who has the key can open it and view the contents. *SecretBox* functions as just such a safe, and like any good safe any attempts to tamper with the contents are easily detected.

Secret key encryption allows you to store or transmit data over insecure channels without leaking the contents of that message, nor anything about it other than the length.

2.2.1 Example

```
import nacl.secret
import nacl.utils

# This must be kept secret, this is the combination to your safe
key = nacl.utils.random(nacl.secret.SecretBox.KEY_SIZE)

# This is your safe, you can use it to encrypt or decrypt messages
box = nacl.secret.SecretBox(key)

# This is our message to send, it must be a bytestring as SecretBox will
# treat it as just a binary blob of data.
message = b"The president will be exiting through the lower levels"
```

PyNaCl can automatically generate a random nonce for us, making the encryption very simple:

```
# Encrypt our message, it will be exactly 40 bytes longer than the
# original message as it stores authentication information and the
# nonce alongside it.
encrypted = box.encrypt(message)

assert len(encrypted) == len(message) + box.NONCE_SIZE + box.MACBYTES
```

However, if we need to use an explicit nonce, it can be passed along with the message:

```
# This is a nonce, it MUST only be used once, but it is not considered
# secret and can be transmitted or stored alongside the ciphertext. A
# good source of nonces are just sequences of 24 random bytes.
nonce = nacl.utils.random(nacl.secret.SecretBox.NONCE_SIZE)

encrypted = box.encrypt(message, nonce)
```

If you need to get the ciphertext and the authentication data without the nonce, you can get the *ciphertext* attribute of the *EncryptedMessage* instance returned by *encrypt()*:

```
nonce = nacl.utils.random(nacl.secret.SecretBox.NONCE_SIZE)

encrypted = box.encrypt(message, nonce)

# since we are transmitting the nonce by some other means,
# we just need to get the ciphertext and authentication data

ctext = encrypted.ciphertext

# ctext is just nacl.secret.SecretBox.MACBYTES longer
# than the original message

assert len(ctext) == len(message) + box.MACBYTES
```

Finally, the message is decrypted (regardless of how the nonce was generated):

```
# Decrypt our message, an exception will be raised if the encryption was
# tampered with or there was otherwise an error.
plaintext = box.decrypt(encrypted)
print(plaintext.decode('utf-8'))
```

The president will be exiting through the lower levels

2.2.2 Requirements

Key

The 32 bytes key given to `SecretBox` must be kept secret. It is the combination to your “safe” and anyone with this key will be able to decrypt the data, or encrypt new data.

Nonce

The 24-byte nonce (Number used once) given to `encrypt()` and `decrypt()` must **NEVER** be reused for a particular key. Reusing a nonce may give an attacker enough information to decrypt or forge other messages. A nonce is not considered secret and may be freely transmitted or stored in plaintext alongside the ciphertext.

A nonce does not need to be random or unpredictable, nor does the method of generating them need to be secret. A nonce could simply be a counter incremented with each message encrypted, which can be useful in connection-oriented protocols to reject duplicate messages (“replay attacks”). A bidirectional connection could use the same key for both directions, as long as their nonces never overlap (e.g. one direction always sets the high bit to “1”, the other always sets it to “0”).

If you use a counter-based nonce along with a key that is persisted from one session to another (e.g. saved to disk), you must store the counter along with the key, to avoid accidental nonce reuse on the next session. For this reason, many protocols derive a new key for each session, reset the counter to zero with each new key, and never store the derived key or the counter.

You can safely generate random nonces by calling `random()` with `SecretBox.NONCE_SIZE`.

2.2.3 Reference

class `nacl.secret.SecretBox` (*key, encoder*)

The `SecretBox` class encrypts and decrypts messages using the given secret key.

The ciphertexts generated by `SecretBox` include a 16 byte authenticator which is checked as part of the decryption. An invalid authenticator will cause the decrypt function to raise an exception. The authenticator is not a signature. Once you’ve decrypted the message you’ve demonstrated the ability to create arbitrary valid message, so messages you send are repudiable. For non-repudiable messages, sign them after encryption.

Parameters

- **key** (*bytes*) – The secret key used to encrypt and decrypt messages.
- **encoder** – A class that is able to decode the *key*.

encrypt (*plaintext, nonce, encoder*)

Encrypts the plaintext message using the given *nonce* (or generates one randomly if omitted) and returns the ciphertext encoded with the encoder.

Warning: It is **VITALLY** important that the nonce is a nonce, i.e. it is a number used only once for any given key. If you fail to do this, you compromise the privacy of the messages encrypted. Give your nonces a different prefix, or have one side use an odd counter and one an even counter. Just make sure they are different.

Parameters

- **plaintext** (*bytes*) – The plaintext message to encrypt.
- **nonce** (*bytes*) – The nonce to use in the encryption.
- **encoder** – A class that is able to decode the ciphertext.

Returns An instance of *EncryptedMessage*.

decrypt (*ciphertext, nonce, encoder*)

Decrypts the ciphertext using the *nonce* (explicitly, when passed as a parameter or implicitly, when omitted, as part of the ciphertext) and returns the plaintext message.

Parameters

- **ciphertext** (*bytes*) – The encrypted message to decrypt.
- **nonce** (*bytes*) – The nonce to use in the decryption.
- **encoder** – A class that is able to decode the plaintext.

Return bytes The decrypted plaintext.

2.2.4 Algorithm details

Encryption Salsa20 stream cipher

Authentication Poly1305 MAC

2.3 Digital Signatures

You can use a digital signature for many of the same reasons that you might sign a paper document. A valid digital signature gives a recipient reason to believe that the message was created by a known sender such that they cannot deny sending it (authentication and non-repudiation) and that the message was not altered in transit (integrity).

Digital signatures allow you to publish a public key, and then you can use your private signing key to sign messages. Others who have your public key can then use it to validate that your messages are actually authentic.

2.3.1 Example

Signer's perspective (*SigningKey*)

```
import nacl.encoding
import nacl.signing

# Generate a new random signing key
signing_key = nacl.signing.SigningKey.generate()

# Sign a message with the signing key
signed = signing_key.sign(b"Attack at Dawn")

# Obtain the verify key for a given signing key
verify_key = signing_key.verify_key

# Serialize the verify key to send it to a third party
verify_key_hex = verify_key.encode(encoder=nacl.encoding.HexEncoder)
```


Verifier's perspective (*VerifyKey*)

```
import nacl.signing

# Create a VerifyKey object from a hex serialized public key
verify_key = nacl.signing.VerifyKey(verify_key_hex,
                                    encoder=nacl.encoding.HexEncoder)

# Check the validity of a message's signature
# The message and the signature can either be passed separately or
# concatenated together. These are equivalent:
verify_key.verify(signed)
verify_key.verify(signed.message, signed.signature)

# Alter the signed message text
forged = signed[:-1] + bytes([int(signed[-1]) ^ 1])
# Will raise nacl.exceptions.BadSignatureError, since the signature check
# is failing
verify_key.verify(forged)
```

```
Traceback (most recent call last):
...
nacl.exceptions.BadSignatureError: Signature was forged or corrupt
```

2.3.2 Reference

class `nacl.signing.SigningKey` (*seed*, *encoder*)

Private key for producing digital signatures using the Ed25519 algorithm.

Signing keys are produced from a 32-byte (256-bit) random seed value. This value can be passed into the *SigningKey* as a `bytes()` whose length is 32.

Warning: This **must** be protected and remain secret. Anyone who knows the value of your *SigningKey* or its seed can masquerade as you.

Parameters

- **seed** (*bytes*) – Random 32-byte value (i.e. private key).
- **encoder** – A class that is able to decode the *seed*.

verify_key

An instance of *VerifyKey* (i.e. public key) that corresponds with the signing key.

classmethod generate ()

Generates a random *SigningKey* object

Returns An instance of *SigningKey*.

sign (*message*, *encoder*)

Sign a message using this key.

Parameters

- **message** (*bytes*) – The data to be signed.
- **encoder** – A class that is able to decode the signed message.

Returns An instance of *SignedMessage*.

class `nacl.signing.VerifyKey(key, encoder)`

The public key counterpart to an Ed25519 *SigningKey* for producing digital signatures.

Parameters

- **key** (*bytes*) – A serialized Ed25519 public key.
- **encoder** – A class that is able to decode the *key*.

verify (*smessage, signature, encoder*)

Verifies the signature of a signed message.

Parameters

- **smessage** (*bytes*) – The signed message to verify. This is either the original message or the concated signature and message.
- **signature** (*bytes*) – The signature of the message to verify against. If the value of *smessage* is the concated signature and message, this parameter can be `None`.
- **encoder** – A class that is able to decode the secret message and signature.

Return bytes The message if successfully verified.

Raises `nacl.exceptions.BadSignatureError` – This is raised if the signature is invalid.

class `nacl.signing.SignedMessage`

A bytes subclass that holds a messaged that has been signed by a *SigningKey*.

signature

The signature contained within the *SignedMessage*.

message

The message contained within the *SignedMessage*.

2.3.3 Ed25519

Ed25519 is a public-key signature system with several attractive features:

- **Fast single-signature verification:** Ed25519 takes only 273364 cycles to verify a signature on Intel’s widely deployed Nehalem/Westmere lines of CPUs. (This performance measurement is for short messages; for very long messages, verification time is dominated by hashing time.) Nehalem and Westmere include all Core i7, i5, and i3 CPUs released between 2008 and 2010, and most Xeon CPUs released in the same period.
- **Even faster batch verification:** Ed25519 performs a batch of 64 separate signature verifications (verifying 64 signatures of 64 messages under 64 public keys) in only 8.55 million cycles, i.e., under 134000 cycles per signature. Ed25519 fits easily into L1 cache, so contention between cores is negligible: a quad-core 2.4GHz Westmere verifies 71000 signatures per second, while keeping the maximum verification latency below 4 milliseconds.
- **Very fast signing:** Ed25519 takes only 87548 cycles to sign a message. A quad-core 2.4GHz Westmere signs 109000 messages per second.
- **Fast key generation:** Key generation is almost as fast as signing. There is a slight penalty for key generation to obtain a secure random number from the operating system; `/dev/urandom` under Linux costs about 6000 cycles.
- **High security level:** This system has a 2^{128} security target; breaking it has similar difficulty to breaking NIST P-256, RSA with ~3000-bit keys, strong 128-bit block ciphers, etc. The best attacks known actually cost more than 2^{140} bit operations on average, and degrade quadratically in success probability as the number of bit operations drops.


- **Collision resilience:** Hash-function collisions do not break this system. This adds a layer of defense against the possibility of weakness in the selected hash function.
- **No secret array indices:** Ed25519 never reads or writes data from secret addresses in RAM; the pattern of addresses is completely predictable. Ed25519 is therefore immune to cache-timing attacks, hyperthreading attacks, and other side-channel attacks that rely on leakage of addresses through the CPU cache.
- **No secret branch conditions:** Ed25519 never performs conditional branches based on secret data; the pattern of jumps is completely predictable. Ed25519 is therefore immune to side-channel attacks that rely on leakage of information through the branch-prediction unit.
- **Small signatures:** Ed25519 signatures are only 512-bits (64 bytes), one of the smallest signature sizes available.
- **Small keys:** Ed25519 keys are only 256-bits (32 bytes), making them small enough to easily copy and paste. Ed25519 also allows the public key to be derived from the private key, meaning that it doesn't need to be included in a serialized private key in cases you want both.
- **Deterministic:** Unlike (EC)DSA, Ed25519 does not rely on an entropy source when signing messages (which has led to [catastrophic private key compromises](#)), but instead computes signature nonces from a combination of a hash of the signing key's "seed" and the message to be signed. This avoids using an entropy source for nonces, which can be a potential attack vector if the entropy source is not generating good random numbers. Even a single reused nonce can lead to a complete disclosure of the private key in these schemes, which Ed25519 avoids entirely by being deterministic instead of tied to an entropy source.

The numbers 87548 and 273364 shown above are official [eBATS](#) reports for a Westmere CPU (Intel Xeon E5620, hydra2).

Ed25519 signatures are elliptic-curve signatures, carefully engineered at several levels of design and implementation to achieve very high speeds without compromising security.

Algorithm

- **Signatures:** [Ed25519 digital signature system](#)



k Ed25519 private key (passed into *SigningKey*)

A Ed25519 public key derived from **k**

M message to be signed

R a deterministic nonce value calculated from a combination of private key data **RH** and the message **M**

S Ed25519 signature

2.4 Hashing

Cryptographic secure hash functions are irreversible transforms of input data to a fixed length *digest*.

The standard properties of a cryptographic hash make these functions useful both for standalone usage as data integrity checkers, as well as `black-box` building blocks of other kind of algorithms and data structures.

All of the hash functions exposed in `nacl.hash` can be used as data integrity checkers.

(continued from previous page)

```
MSG = 'Digest of {0} message {1} original digest'

for chk in (('original', orig_dgs),
            ('truncated', shrt_dgs),
            ('modified', mdfd_dgs)):

    print(MSG.format(chk[0], eq_chk(dgst, chk[1])))
```

```
Digest of original message equals original digest
Digest of truncated message is different from original digest
Digest of modified message is different from original digest
```

2.4.2 Additional hashing usages for blake2b

As already hinted above, traditional cryptographic hash functions can be used as building blocks for other uses, typically combining a secret-key with the message via some construct like the HMAC one.

The *blake2b* hash function can be used directly both for message authentication and key derivation, replacing the HMAC construct and the HKDF one by setting the additional parameters *key*, *salt* and *person*.

Please note that **key stretching procedures** like HKDF or the one outlined in *Key derivation* are **not** suited to derive a *cryptographically-strong* key from a *low-entropy input* like a plain-text password or to compute a strong *long-term stored* hash used as password verifier. See the *Password hashing* section for some more informations and usage examples of the password hashing constructs provided in *pwhash*.

2.4.3 Message authentication

To authenticate a message, using a secret key, the *blake2b* function must be called as in the following example.

Message authentication example

```
import nacl.encoding
import nacl.utils
from nacl.hash import blake2b

msg = 16*b'256 BytesMessage'
msg2 = 16*b'256 bytesMessage'

auth_key = nacl.utils.random(size=64)
# the simplest way to get a cryptographic quality auth_key
# is to generate it with a cryptographic quality
# random number generator

auth1_key = nacl.utils.random(size=64)
# generate a different key, just to show the mac is changed
# both with changing messages and with changing keys

mac0 = blake2b(msg, key=auth_key, encoder=nacl.encoding.HexEncoder)
mac1 = blake2b(msg, key=auth1_key, encoder=nacl.encoding.HexEncoder)
mac2 = blake2b(msg2, key=auth_key, encoder=nacl.encoding.HexEncoder)
```

(continues on next page)

(continued from previous page)

```
for i, mac in enumerate((mac0, mac1, mac2)):
    print('Mac{0} is: {1}'.format(i, mac))
```

```
Mac0 is: b'...'
Mac1 is: b'...'
Mac2 is: b'...'.
```

2.4.4 Key derivation

The blake2b algorithm can replace a key derivation function by following the lines of:

Key derivation example

```
import nacl.encoding
import nacl.utils
from nacl.hash import blake2b

master_key = nacl.utils.random(64)

derivation_salt = nacl.utils.random(16)

personalization = b'<DK usage>'

derived = blake2b(b'', key=master_key, salt=derivation_salt,
                 person=personalization,
                 encoder=nacl.encoding.RawEncoder)
```

By repeating the key derivation procedure before encrypting our messages, and sending the `derivation_salt` along with the encrypted message, we can expect to never reuse a key, drastically reducing the risks which ensue from such a reuse.

2.5 Password hashing

Password hashing and password based key derivation mechanisms in actual use are all based on the idea of iterating a hash function many times on a combination of the password and a random `salt`, which is stored along with the hash, and allows verifying a proposed password while avoiding clear-text storage.

The latest developments in password hashing have been *memory-hard* and *tunable* mechanisms, pioneered by `scrypt` [SD2012], and followed-on by the schemes submitted to the **Password Hashing Competition** [PHC].

The `nacl.pwhash` module exposes both the PHC recommended partially data dependent `argon2id` and the data independent `argon2i` mechanisms alongside to the `scrypt` one.

In the case of password storage, it's usually suggested to give preference to data dependent mechanisms, therefore the default mechanism suggested by `libsodium` since version 1.0.15, and therefore by `PyNaCl` since version 1.2 is `argon2id`.

If you think in your use-case the risk of potential timing-attacks stemming from data-dependency is greater than the potential time/memory trade-offs stemming out of data-independency, you should prefer `argon2i` to `argon2id` or `scrypt`

2.5.1 Hashers and parameters

PyNaCl exposes the functions and the associated parameters needed to exploit the password hashing constructions in a uniform way in the modules `argon2id`, `argon2i` and `scrypt`, therefore, if you need to change your choice of construction, you simply need to replace one module name with another in the example below.

Further, if you just want to use a default choosen construction, you can directly call `nacl.pwhash.str()` or `nacl.pwhash.kdf()` to use the preferred construct in modular crypt password hashing or key derivation mode.

2.5.2 Password storage and verification

All implementations of the modular crypt hasher `str` function internally generate a random salt, and return a hash encoded in ascii modular crypt format, which can be stored in a shadow-like file

```
>>> import nacl.pwhash
>>> password = b'my password'
>>> for i in range(4):
...     print(nacl.pwhash.str(password))
...
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
>>>
>>> # if needed, each hasher is exposed
... # in just the same way:
... for i in range(4):
...     print(nacl.pwhash.scrypt.str(password))
...
b'$7$c6..../...'
b'$7$c6..../...'
b'$7$c6..../...'
b'$7$c6..../...'
>>>
>>> for i in range(4):
...     print(nacl.pwhash.argon2i.str(password))
...
b'$argon2i$v=19$m=32768,t=4,p=1$...'
b'$argon2i$v=19$m=32768,t=4,p=1$...'
b'$argon2i$v=19$m=32768,t=4,p=1$...'
b'$argon2i$v=19$m=32768,t=4,p=1$...'
>>>
>>> # and
...
>>> for i in range(4):
...     print(nacl.pwhash.argon2id.str(password))
...
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
b'$argon2id$v=19$m=65536,t=2,p=1$...'
>>>
```

To verify a user-proposed password, the `verify()` function checks the stored hash prefix, and dispatches verification to the correct checker, which in turn extracts the used salt, memory and operation count parameters from the modular format string and checks the compliance of the proposed password with the stored hash


```

>>> import nacl.pwhash
>>> hashed = (b'$7$C6..../....qv5tF9KG2WbuMeU0a0TCoqwLHQ8s0TjQdSagne'
...          b'9NvU0$3d218uChMvdvN6EwSvKHMASKZIG51XPISZQDcktKyN7'
...          )
>>> correct = b'my password'
>>> wrong = b'My password'
>>> # while the result will be True on password match,
... # on mismatch an exception will be raised
... res = nacl.pwhash.verify(hashed, correct)
>>> print(res)
True
>>>
>>> res2 = nacl.pwhash.verify_scryptsalsa208sha256(hashed, wrong)
Traceback (most recent call last):
...
nacl.exceptions.InvalidkeyError: Wrong password
>>> # the verify function raises an exception
... # also when it is run against a password hash
... # starting with a prefix it doesn't know
... wrong_hash = (b'?$7$C6..../....qv5tF9KG2WbuMeU0a0TCoqwLHQ8s0TjQdSagne'
...               b'9NvU0$3d218uChMvdvN6EwSvKHMASKZIG51XPISZQDcktKyN7'
...               )
>>> res = nacl.pwhash.verify(wrong_hash, correct)
Traceback (most recent call last):
...
nacl.exceptions.InvalidkeyError: given password_hash is not in a supported format

```

2.5.3 Key derivation

Alice needs to send a secret message to Bob, using a shared password to protect the content. She generates a random salt, combines it with the password using one of the *kdf* functions and sends the message along with the salt and key derivation parameters.

```

from nacl import pwhash, secret, utils

password = b'password shared between Alice and Bob'
message = b"This is a message for Bob's eyes only"

kdf = pwhash.argon2i.kdf
salt = utils.random(pwhash.argon2i.SALTBYTES)
ops = pwhash.argon2i.OPSLIMIT_SENSITIVE
mem = pwhash.argon2i.MEMLIMIT_SENSITIVE

# or, if there is a need to use scrypt:
# kdf = pwhash.scrypt.kdf
# salt = utils.random(pwhash.scrypt.SALTBYTES)
# ops = pwhash.scrypt.OPSLIMIT_SENSITIVE
# mem = pwhash.scrypt.MEMLIMIT_SENSITIVE

Alices_key = kdf(secret.SecretBox.KEY_SIZE, password, salt,
                 opslimit=ops, memlimit=mem)
Alices_box = secret.SecretBox(Alices_key)
nonce = utils.random(secret.SecretBox.NONCE_SIZE)

```

(continues on next page)

(continued from previous page)

```

encrypted = Alices_box.encrypt(message, nonce)

# now Alice must send to Bob both the encrypted message
# and the KDF parameters: salt, opslimit and memlimit;
# using the same kdf mechanism, parameters **and password**
# Bob is able to derive the correct key to decrypt the message

Bobs_key = kdf(secret.SecretBox.KEY_SIZE, password, salt,
              opslimit=ops, memlimit=mem)
Bobs_box = secret.SecretBox(Bobs_key)
received = Bobs_box.decrypt(encrypted)
print(received.decode('utf-8'))

```

```
This is a message for Bob's eyes only
```

if Eve manages to get the encrypted message, and tries to decrypt it with a incorrect password, even if she does know all of the key derivation parameters, she would derive a different key. Therefore the decryption would fail and an exception would be raised

```

>>> # ops, mem and salt are the same used by Alice
...
>>>
>>> guessed_pw = b'I think Alice shared this password with Bob'
>>>
>>> Eves_key = pwhash.argon2i.kdf(secret.SecretBox.KEY_SIZE,
...                               guessed_pw, salt,
...                               opslimit=ops, memlimit=mem)
>>> Eves_box = secret.SecretBox(Eves_key)
>>> intercepted = Eves_box.decrypt(encrypted)
Traceback (most recent call last):
...
nacl.exceptions.CryptoError: Decryption failed. Ciphertext failed ...

```

Contrary to the hashed password storage case where a serialization format is well-defined, in the raw key derivation case the library user must take care to store (and retrieve) both a reference to the kdf used to derive the secret key and all the derivation parameters. These parameters are needed to later generate the same secret key from the password.

Module level constants for operation and memory cost tweaking

To help in selecting the correct values for the tweaking parameters for the used construction, all the implementation modules provide suggested values for the *opslimit* and *memlimit* parameters with the names:

- *OPSLIMIT_INTERACTIVE*
- *MEMLIMIT_INTERACTIVE*
- *OPSLIMIT_SENSITIVE*
- *MEMLIMIT_SENSITIVE*
- *OPSLIMIT_MODERATE*
- *MEMLIMIT_MODERATE*

and the corresponding minimum and maximum allowed values in:

- *OPSLIMIT_MIN*

- *MEMLIMIT_MIN*
- *OPSLIMIT_MAX*
- *MEMLIMIT_MAX*

Further, for each construction, pwhash modules expose the following constants:

- *STRPREFIX*
- *PWHASH_SIZE*
- *SALTBYTES*
- *BYTES_MIN*
- *BYTES_MAX*

In general, the `_INTERACTIVE` values are recommended in the case of hashes stored for interactive password checking, and lead to a sub-second password verification time, with a memory consumption in the tens of megabytes range, while the `_SENSITIVE` values are meant to store hashes for password protecting sensitive data, and lead to hashing times exceeding one second, with memory consumption in the hundred of megabytes range. The `_MODERATE` values, suggested for `argon2` mechanisms are meant to run the construct at a runtime and memory cost intermediate between `_INTERACTIVE` and `_SENSITIVE`.

3.1 Encoders

PyNaCl supports a simple method of encoding and decoding messages in different formats. Encoders are simple classes with static methods that encode/decode and are typically passed as a keyword argument *encoder* to various methods.

For example you can generate a signing key and encode it in hex with:

```
hex_key = nacl.signing.SigningKey.generate().encode(encoder=nacl.encoding.HexEncoder)
```

Then you can later decode it from hex:

```
signing_key = nacl.signing.SigningKey(hex_key, encoder=nacl.encoding.HexEncoder)
```

3.1.1 Built in Encoders

```
class nacl.encoding.RawEncoder
class nacl.encoding.HexEncoder
class nacl.encoding.Base16Encoder
class nacl.encoding.Base32Encoder
class nacl.encoding.Base64Encoder
class nacl.encoding.URLSafeBase64Encoder
```

3.1.2 Defining your own Encoder

Defining your own encoder is easy. Each encoder is simply a class with 2 static methods. For example here is the hex encoder:

```
import binascii

class HexEncoder(object):

    @staticmethod
    def encode(data):
        return binascii.hexlify(data)

    @staticmethod
    def decode(data):
        return binascii.unhexlify(data)
```

3.2 Exceptions

All of the exceptions raised from PyNaCl-exposed methods/functions are subclasses of `nacl.exceptions.CryptoError`. This means downstream users can just wrap cryptographic operations inside a

```
try:
    # cryptographic operations
except nacl.exceptions.CryptoError:
    # cleanup after any kind of exception
    # raised from cryptographic-related operations
```

These are the exceptions implemented in `nacl.exceptions`:

3.2.1 PyNaCl specific exceptions

```
class CryptoError
    Base exception for all nacl related errors

class BadSignatureError
    Raised when the signature was forged or otherwise corrupt.

class InvalidkeyError
    Raised on password/key verification mismatch
```

3.2.2 PyNaCl exceptions mixing-in standard library ones

Both for clarity and for compatibility with previous releases of the PyNaCl, the following exceptions mix-in the same-named standard library exception to `CryptoError`.

```
class RuntimeError
    is a subclass of both CryptoError and standard library's RuntimeError, raised for internal library errors

class AssertionError
    is a subclass of both CryptoError and standard library's AssertionError, raised by default from ensure() when the checked condition is False

class TypeError
    is a subclass of both CryptoError and standard library's TypeError

class ValueError
    is a subclass of both CryptoError and standard library's ValueError
```

3.3 Utilities

class `nacl.utils.EncryptedMessage`

A `bytes` subclass that holds a message that has been encrypted by a `SecretBox` or `Box`. The full content of the `bytes` object is the combined nonce and ciphertext.

nonce

The nonce used during the encryption of the `EncryptedMessage`.

ciphertext

The ciphertext contained within the `EncryptedMessage`.

`nacl.utils.random(size=32)`

Returns a random bytestring with the given `size`.

Parameters `size` (*bytes*) – The size of the random bytestring.

Return bytes The random bytestring.

`nacl.utils.ensure(cond, *args, raising=nacl.exceptions.AssertionError)`

Returns if a condition is true, otherwise raise a caller-configurable `Exception`

Parameters

- **cond** (*bool*) – the condition to be checked
- **args** (*sequence*) – the arguments to be passed to the exception's constructor
- **raising** (*exception*) – the exception to be raised if `cond` is `False`

3.4 nacl.hash

`nacl.hash.sha256(message, encoder)`

Hashes `message` with SHA256.

Parameters

- **message** (*bytes*) – The message to hash.
- **encoder** – A class that is able to encode the hashed message.

Return bytes The hashed message.

`nacl.hash.sha512(message, encoder)`

Hashes `message` with SHA512.

Parameters

- **message** (*bytes*) – The message to hash.
- **encoder** – A class that is able to encode the hashed message.

Return bytes The hashed message.

`nacl.hash.blake2b(data, digest_size=BLAKE2B_BYTES, key=b'', salt=b'', person=b'', encoder=nacl.encoding.HexEncoder)`

One-shot blake2b digest

Parameters

- **data** (*bytes*) – the digest input byte sequence

- **digest_size** (*int*) – the requested digest size; must be at most `BLAKE2B_BYTES_MAX`; the default digest size is `BLAKE2B_BYTES`
- **key** (*bytes*) – the key to be set for keyed MAC/PRF usage; if set, the key must be at most `BLAKE2B_KEYBYTES_MAX` long
- **salt** (*bytes*) – an initialization salt at most `BLAKE2B_SALTBYTES` long; it will be zero-padded if needed
- **person** (*bytes*) – a personalization string at most `BLAKE2B_PERSONALBYTES` long; it will be zero-padded if needed
- **encoder** (*class*) – the encoder to use on returned digest

Returns encoded bytes data

Return type the return type of the chosen encoder

`nacl.hash.siphash24` (*message*, *key=b*”, *encoder=nacl.encoding.HexEncoder*)

Computes a keyed MAC of *message* using siphash-2-4

Parameters

- **message** (*bytes*) – The message to hash.
- **key** (*bytes(SIPHASH_KEYBYTES)*) – the message authentication key to be used It must be a `SIPHASH_KEYBYTES` long bytes sequence
- **encoder** – A class that is able to encode the hashed message.

Returns The hashed message.

Return type `bytes(SIPHASH_BYTES)` long bytes sequence

`nacl.hash.siphashx24` (*message*, *key=b*”, *encoder=nacl.encoding.HexEncoder*)

New in version 1.2.

Computes a keyed MAC of *message* using the extended output length variant of siphash-2-4

Parameters

- **message** (*bytes*) – The message to hash.
- **key** (*bytes(SIPHASHX_KEYBYTES)*) – the message authentication key to be used It must be a `SIPHASHX_KEYBYTES` long bytes sequence
- **encoder** – A class that is able to encode the hashed message.

Returns The hashed message.

Return type `bytes(SIPHASHX_BYTES)` long bytes sequence

3.5 nacl.pwhash

The package `pwhash` provides implementations of modern *memory-hard* password hashing construction exposing modules with a uniform API.

3.5.1 Functions exposed at top level

The top level module only provides the functions implementing ascii encoded hashing and verification.

`nacl.pwhash.str` (*password*, *opslimit=OPSLIMIT_INTERACTIVE*, *memlimit=MEMLIMIT_INTERACTIVE*)
Returns a password verifier hash, generated with the password hasher chosen as a default by libsodium.

Parameters

- **password** (*bytes*) – password used to seed the key derivation procedure; its length must be between `PASSWD_MIN` and `PASSWD_MAX`
- **opslimit** (*int*) – the time component (operation count) of the key derivation procedure’s computational cost; it must be between `OPSLIMIT_MIN` and `OPSLIMIT_MAX`
- **memlimit** (*int*) – the memory occupation component of the key derivation procedure’s computational cost; it must be between `MEMLIMIT_MIN` and `MEMLIMIT_MAX`

Returns the ascii encoded password hash along with a prefix encoding the used hashing construct, the random generated salt and the operation and memory limits used to generate the password hash

Return type `bytes`

As of PyNaCl version 1.2 this is `nacl.pwhash.argon2id.str()`.

New in version 1.2.

`nacl.pwhash.verify` (*password_hash*, *password*)

This function checks if hashing the proposed password, with the same construction and parameters encoded in the password hash would generate the same encoded string, thus verifying the correct password has been proposed in an authentication attempt.

New in version 1.2.

Module level constants

The top level module defines the constants related to the `str()` hashing construct and its corresponding `verify()` password verifier.

`nacl.pwhash.PASSWD_MIN`

`nacl.pwhash.PASSWD_MAX`
minimum and maximum length of the password to hash

`nacl.pwhash.PWHASH_SIZE`
maximum size of the encoded hash

`nacl.pwhash.OPSLIMIT_MIN`

`nacl.pwhash.OPSLIMIT_MAX`
minimum and maximum operation count for the hashing construct

`nacl.pwhash.MEMLIMIT_MIN`

`nacl.pwhash.MEMLIMIT_MAX`
minimum and maximum memory occupation for the hashing construct

and the recommended values for the `opslimit` and `memlimit` parameters

`nacl.pwhash.MEMLIMIT_INTERACTIVE`

`nacl.pwhash.OPSLIMIT_INTERACTIVE`
recommended values for the interactive user authentication password check case, leading to a sub-second hashing time

`nacl.pwhash.MEMLIMIT_SENSITIVE`

`nacl.pwhash.OPSLIMIT_SENSITIVE`

recommended values for generating a password hash/derived key meant to protect sensitive data, leading to a multi-second hashing time

`nacl.pwhash.MEMLIMIT_MODERATE`

`nacl.pwhash.OPSLIMIT_MODERATE`

values leading to a hashing time and memory cost intermediate between the interactive and the sensitive cases

3.5.2 Per-mechanism password hashing implementation modules

Along with the respective `str()` and `verify()` functions, the modules implementing named password hashing constructs expose also a `kdf()` function returning a raw pseudo-random bytes sequence derived from the input parameters

`nacl.pwhash.argon2id`

`nacl.pwhash.argon2id.kdf` (*size*, *password*, *salt*, *opslimit*=`OPSLIMIT_SENSITIVE`, *memlimit*=`MEMLIMIT_SENSITIVE`, *encoder*=`nacl.encoding.RawEncoder`)

Derive a *size* bytes long key from a caller-supplied *password* and *salt* pair using the `argon2id` partially data dependent memory-hard construct.

Parameters

- **size** (*int*) – derived key size, must be between `BYTES_MIN` and `BYTES_MAX`
- **password** (*bytes*) – password used to seed the key derivation procedure; its length must be between `PASSWD_MIN` and `PASSWD_MAX`
- **salt** (*bytes*) – **RANDOM** salt used in the key derivation procedure; its length must be exactly `SALTBYTES`
- **opslimit** (*int*) – the time component (operation count) of the key derivation procedure’s computational cost; it must be between `OPSLIMIT_MIN` and `OPSLIMIT_MAX`
- **memlimit** (*int*) – the memory occupation component of the key derivation procedure’s computational cost; it must be between `MEMLIMIT_MIN` and `MEMLIMIT_MAX`

Return type `bytes`

The default settings for `opslimit` and `memlimit` are those deemed correct for generating a key, which can be used to protect sensitive data for a long time, leading to a multi-second hashing time.

New in version 1.2.

`nacl.pwhash.argon2id.str` (*password*, *opslimit*=`OPSLIMIT_INTERACTIVE`, *memlimit*=`MEMLIMIT_INTERACTIVE`)

Returns a password verifier hash, generated with the `argon2id` password hasher.

See: `nacl.pwhash.str()` for the general API.

New in version 1.2.

`nacl.pwhash.argon2id.verify` (*password_hash*, *password*)

This function verifies the proposed *password*, using *password_hash* as a password verifier.

See: `nacl.pwhash.verify()` for the general API.

New in version 1.2.

Module level constants

The module defines the constants related to the *kdf()* raw hashing construct

`nacl.pwhash.argon2id.SALTBYTES`
 the length of the random bytes sequence passed in as a salt to the *kdf()*

`nacl.pwhash.argon2id.BYTES_MIN`

`nacl.pwhash.argon2id.BYTES_MAX`
 the minimum and maximum allowed values for the `size` parameter of the *kdf()*

The meaning of each of the constants

`nacl.pwhash.argon2id.PASSWD_MIN`

`nacl.pwhash.argon2id.PASSWD_MAX`

`nacl.pwhash.argon2id.PWHASH_SIZE`

`nacl.pwhash.argon2id.OPSLIMIT_MIN`

`nacl.pwhash.argon2id.OPSLIMIT_MAX`

`nacl.pwhash.argon2id.MEMLIMIT_MIN`

`nacl.pwhash.argon2id.MEMLIMIT_MAX`

`nacl.pwhash.argon2id.MEMLIMIT_INTERACTIVE`

`nacl.pwhash.argon2id.OPSLIMIT_INTERACTIVE`

`nacl.pwhash.argon2id.MEMLIMIT_SENSITIVE`

`nacl.pwhash.argon2id.OPSLIMIT_SENSITIVE`

`nacl.pwhash.argon2id.MEMLIMIT_MODERATE`

`nacl.pwhash.argon2id.OPSLIMIT_MODERATE`
 is the same as in `nacl.hash`.

`nacl.pwhash.argon2i`

`nacl.pwhash.argon2i.kdf`(*size*, *password*, *salt*, *opslimit=OPSLIMIT_SENSITIVE*, *mem-limit=MEMLIMIT_SENSITIVE*, *encoder=nacl.encoding.RawEncoder*)
 Derive a *size* bytes long key from a caller-supplied *password* and *salt* pair using the `argon2i` data independent memory-hard construct.

See: `py:func:nacl.pwhash.argon2id.kdf` for the general API.

New in version 1.2.

`nacl.pwhash.argon2i.str`(*password*, *opslimit=OPSLIMIT_INTERACTIVE*, *mem-limit=MEMLIMIT_INTERACTIVE*)
 Returns a password verifier hash, generated with the `argon2i` password hasher.

See: `nacl.pwhash.str()` for the general API.

New in version 1.2.

`nacl.pwhash.argon2i.verify`(*password_hash*, *password*)
 This function verifies the proposed *password*, using *password_hash* as a password verifier.

See: `nacl.pwhash.verify()` for the general API.

New in version 1.2.

Module level constants

The meaning of each of the constants

`nacl.pwhash.argon2i.PASSWD_MIN`
`nacl.pwhash.argon2i.PASSWD_MAX`
`nacl.pwhash.argon2i.PWHASH_SIZE`
`nacl.pwhash.argon2i.SALTBYTES`
`nacl.pwhash.argon2i.BYTES_MIN`
`nacl.pwhash.argon2i.BYTES_MAX`
`nacl.pwhash.argon2i.OPSLIMIT_MIN`
`nacl.pwhash.argon2i.OPSLIMIT_MAX`
`nacl.pwhash.argon2i.MEMLIMIT_MIN`
`nacl.pwhash.argon2i.MEMLIMIT_MAX`
`nacl.pwhash.argon2i.MEMLIMIT_INTERACTIVE`
`nacl.pwhash.argon2i.OPSLIMIT_INTERACTIVE`
`nacl.pwhash.argon2i.MEMLIMIT_SENSITIVE`
`nacl.pwhash.argon2i.OPSLIMIT_SENSITIVE`
`nacl.pwhash.argon2i.MEMLIMIT_MODERATE`
`nacl.pwhash.argon2i.OPSLIMIT_MODERATE`
is the same as in `nacl.pwhash` and `nacl.pwhash.argon2id`

`nacl.pwhash.scrypt`

`nacl.pwhash.scrypt.kdf` (*size*, *password*, *salt*, *opslimit=OPSLIMIT_SENSITIVE*, *mem-limit=MEMLIMIT_SENSITIVE*, *encoder=nacl.encoding.RawEncoder*)
Derive a *size* bytes long key from a caller-supplied *password* and *salt* pair using the `scrypt` data dependent memory-hard construct.

See: `nacl.pwhash.argon2id.kdf()` for the general API.

New in version 1.2.

`nacl.pwhash.scrypt.str` (*password*, *opslimit=OPSLIMIT_INTERACTIVE*, *mem-limit=MEMLIMIT_INTERACTIVE*)
Returns a password verifier hash, generated with the `scrypt` password hasher.

See: `nacl.pwhash.str()` for the general API.

New in version 1.2.

`nacl.pwhash.scrypt.verify` (*password_hash*, *password*)
This function verifies the proposed *password*, using *password_hash* as a password verifier.

See: `py:func:nacl.pwhash.verify` for the general API.

New in version 1.2.

Module level constants

The meaning of each of the constants

```

nacl.pwhash.scrypt.PASSWD_MIN
nacl.pwhash.scrypt.PASSWD_MAX
nacl.pwhash.scrypt.PWHASH_SIZE
nacl.pwhash.scrypt.SALTBYTES
nacl.pwhash.scrypt.BYTES_MIN
nacl.pwhash.scrypt.BYTES_MAX
nacl.pwhash.scrypt.OPSLIMIT_MIN
nacl.pwhash.scrypt.OPSLIMIT_MAX
nacl.pwhash.scrypt.MEMLIMIT_MIN
nacl.pwhash.scrypt.MEMLIMIT_MAX
nacl.pwhash.scrypt.MEMLIMIT_INTERACTIVE
nacl.pwhash.scrypt.OPSLIMIT_INTERACTIVE
nacl.pwhash.scrypt.MEMLIMIT_SENSITIVE
nacl.pwhash.scrypt.OPSLIMIT_SENSITIVE
nacl.pwhash.scrypt.MEMLIMIT_MODERATE
nacl.pwhash.scrypt.OPSLIMIT_MODERATE
    is the same as in nacl.pwhash and nacl.pwhash.argon2id

```

3.6 nacl.hashlib

The `nacl.hashlib` module exposes directly usable implementations of raw constructs which `libsodium` exposes with simplified APIs, like the ones in `nacl.hash` and in `nacl.pwhash`.

The `blake2b` and `scrypt()` implementations are as API compatible as possible with the corresponding ones added to cpython standard library's `hashlib` module in cpython's version 3.6.

class `nacl.hashlib.blake2b` (*data=b", digest_size=BYTES, key=b", salt=b", person=b"*)

Returns an hash object which exposes an API mostly compatible to python3.6's `hashlib.blake2b` (the only difference being missing support for tree hashing parameters in the constructor)

The methods `update()`, `copy()`, `digest()` and `hexdigest()` have the same semantics as described in `hashlib` documentation.

Each instance exposes the `digest_size`, `block_size` name properties as required by `hashlib` API.

MAX_DIGEST_SIZE

the maximum allowed value of the requested `digest_size`

MAX_KEY_SIZE

the maximum allowed size of the password parameter

PERSON_SIZE

the maximum size of the personalization

SALT_SIZE

the maximum size of the salt

`nacl.hashlib.scrypt` (*password*, *salt*=”, *n*=2**20, *r*=8, *p*=1, *maxmem*=2**25, *dklen*=64)

Derive a raw cryptographic key using the scrypt KDF.

Parameters

- **password** (*bytes*) – the input password
- **salt** (*bytes*) – a cryptographically-strong random salt
- **n** (*int*) – CPU/Memory cost factor
- **r** (*int*) – block size multiplier: the used block size will be $128 * r$
- **p** (*int*) – requested parallelism: the number of independently running scrypt constructs which will contribute to the final key generation
- **maxmem** (*int*) – maximum memory the whole scrypt construct will be entitled to use
- **dklen** (*int*) – length of the derived key

Returns a buffer *dklen* bytes long containing the derived key

Implements the same signature as the `hashlib.scrypt` implemented in cpython version 3.6

The recommended values for *n*, *r*, *p* in 2012 were $n = 2^{14}$, $r = 8$, $p = 1$; as of 2016, libsodium suggests using $n = 2^{14}$, $r = 8$, $p = 1$ in a “interactive” setting and $n = 2^{20}$, $r = 8$, $p = 1$ in a “sensitive” setting.

The total memory usage will respectively be a little greater than 16MB in the “interactive” setting, and a little greater than 1GB in the “sensitive” setting.

3.7 Installation

3.7.1 Binary wheel install

PyNaCl ships as a binary wheel on OS X, Windows and Linux [manylinux1](#)¹, so all dependencies are included. Make sure you have an up-to-date pip and run:

```
$ pip install pynacl
```

3.7.2 Linux source build

PyNaCl relies on [libsodium](#), a portable C library. A copy is bundled with PyNaCl so to install you can run:

```
$ pip install pynacl
```

If you’d prefer to use the version of `libsodium` provided by your distribution, you can disable the bundled copy during install by running:

```
$ SODIUM_INSTALL=system pip install pynacl
```

¹ [manylinux1](#) wheels are built on a baseline linux environment based on Centos 5.11 and should work on most x86 and x86_64 glibc based linux environments.

Warning: Usage of the legacy `easy_install` command provided by `setuptools` is generally discouraged, and is completely unsupported in PyNaCl's case.

3.8 Doing A Release

To run a PyNaCl release follow these steps:

- Update the version number in `src/nacl/__init__.py`.
- Update `README.rst` changelog section with the date of the release.
- Send a pull request with these items and wait for it to be merged.
- Run `invoke release {version}`

Once the release script completes you can verify that the `sdist` and `wheels` are present on PyPI and then send a new PR to bump the version to the next major version (e.g. `1.2.0.dev1`).

3.9 Reference vectors

In addition to the policy of keeping any code path in PyNaCl covered by unit tests, the output from cryptographic primitives and constructions must be verified as being conformant to the reference implementations or standards.

3.9.1 Imported reference vectors

Wherever possible it is the PyNaCl project's policy to use existing reference vectors for primitives or constructions. These vectors should ideally be in their original format, but it is acceptable to make minimal changes to ease parsing at our discretion.

Box construction

The reference vector for testing the `nacl.public.Box` implementation come from `libsodium's test/default/box.c` and `test/default/box2.c` and the corresponding expected outputs in `test/default/box.exp` and `test/default/box2.exp`

SecretBox construction

The reference vector for testing the `nacl.secret.SecretBox` implementation come from `libsodium's test/default/secretbox.c` and the corresponding expected outputs in `test/default/secretbox.exp`

chacha20poly1305

The reference vectors for both the legacy `draft-agl-tls-chacha20poly1305` and the IETF ratified `rfc7539` `chacha20poly1305` constructions are taken from `libressl version 2.5.5 tests/aeadttests.txt`, excluding the shortened authentication tag vectors, since `libsodium` only supports full sized tags.

xchacha20poly1305

The reference vector for the xchacha20poly1305 construction is taken from the first test in libsodium's test/default/aead_xchacha20poly1305.c, which defines the parameters, and the corresponding expected output from aead_xchacha20poly1305.exp.

siphash24 and siphashx24

The reference vectors for both the original and the 128 bit variants of the siphash-2-4 construction are taken from the reference code sources. In particular, the original expected results come from siphash's vectors.h, while the key and the input messages have been generated following the respective definitions in siphash's test.c.

3.9.2 Custom generated reference vectors

In cases where there are no standardized test vectors, or the available ones are not applicable to libsodium's implementation, test vectors are custom generated.

Argon2 constructs reference vectors

Since libsodium implements a different API for argon2 constructs than the one exposed by the reference implementation available at *The password hash Argon2...* <<https://github.com/P-H-C/phc-winner-argon2/>>, the `kats` data provided along to the reference implementation sources cannot be directly used as test vectors in PyNaCl tests.

Therefore, we are using a python driver for the command line `argon2`, which can be built following the instruction in the reference implementation sources.

Vector generation

The `argondriver.py` requires setting, via the command line option `-x`, the path to the `argon2` executable; and as a default generates hex-encoded raw hash data on standard output.

Setting the `-e` option on the command line allows generating modular crypt formatted hashes.

The other command line options influence the minimum and maximum sizes of generated parameters as shown in the driver's command line help, which is printed by inserting the `-h` option in the command line.

To generate vector data files in `tests/data`, the `argondriver.py` have been called to generate password hashes with parameters compatible with libsodium's implementation; in particular, the minimum operations count must be 3 for `argon2i` and 1 for `argon2id`, and the salt length must be 16 for raw hashes, and can vary for modular crypt formatted hashes.

The full command lines used in generating the vectors are:

for raw argon2i

```
python3 docs/vectors/python/argondriver.py \
    -x ~/phc-winner-argon2/argon2 \
    -c argon2i \
    -s 16 -S 16 -p 8 -P 16 -m 14 -M 18 \
    -l 18 -L 32 -t 3 -T 5 -n 10 \
    -w tests/data/raw_argon2id_hashes.json
```

for raw argon2id


```
python3 docs/vectors/python/argondriver.py \
-x ~/phc-winner-argon2/argon2 \
-c argon2id \
-s 16 -S 16 -p 8 -P 16 -m 14 -M 18 \
-l 18 -L 32 -t 1 -T 5 -n 10 \
-w tests/data/raw_argon2id_hashes.json
```

for modular crypt argon2i

```
python3 docs/vectors/python/argondriver.py \
-x ~/phc-winner-argon2/argon2 \
-c argon2i -e \
-s 8 -S 16 -p 8 -P 16 -m 14 -M 18 \
-l 18 -L 32 -t 3 -T 5 -n 10 \
-w tests/data/modular_crypt_argon2id_hashes.json
```

for modular crypt argon2id

```
python3 docs/vectors/python/argondriver.py \
-x ~/phc-winner-argon2/argon2 \
-c argon2id -e \
-s 8 -S 16 -p 8 -P 16 -m 14 -M 18 \
-l 18 -L 32 -t 1 -T 5 -n 10 \
-w tests/data/modular_crypt_argon2id_hashes.json
```

Code for the vector generator driver

The code for `argondriver.py` is available inside the `docs/vectors/python` directory of PyNaCl distribution and can also be directly downloaded from `argondriver.py`.

Listing 1: `argondriver.py`

```
#!/usr/bin/python
#
from __future__ import division, print_function

import argparse
import json
import random
import string
import subprocess
import sys

class argonRunner(object):
    GOODCHARS = string.ascii_letters + string.digits

    def __init__(self, args):
        self.exe = args.exe
        self.mnsaltlen = args.mnsaltlen
        self.mnpwlen = args.mnpwlen
        self.mndgstlen = args.mndgstlen
        self.mnmemb = args.mnmemb
        self.mniters = args.mniters
        self.mxsaltlen = args.mxsaltlen
```

(continues on next page)

(continued from previous page)

```

self.mxpwlen = args.mxpwlen
self.mxdgstlen = args.mxdgstlen
self.mxmem = args.mxmem
self.mxiters = args.mxiters
self.encoded = args.encoded
self.rng = random.SystemRandom()
self.version = args.version
self.construct = args.construct
self.maxcount = args.n
self.count = 0

def _runOnce(self, passwd, salt, dgst_len, maxmem, iters):
    """
    """
    argv = [self.exe, salt.encode('ascii'),
            '-t', '{0:2d}'.format(iters),
            '-m', '{0:2d}'.format(maxmem),
            '-l', '{0:3d}'.format(dgst_len),
            '-v', self.version,
            ]

    if self.encoded:
        argv.append('-e')
        mode = 'crypt'
    else:
        argv.append('-r')
        mode = 'raw'
    if self.construct == 'argon2i':
        argv.append('-i')
    elif self.construct == 'argon2d':
        argv.append('-d')
    elif self.construct == 'argon2id':
        argv.append('-id')
    p = subprocess.Popen(argv, stdin=subprocess.PIPE,
                        stdout=subprocess.PIPE)
    out, err = p.communicate(passwd.encode('ascii'))
    return dict(passwd=passwd, salt=salt, dgst_len=dgst_len,
               maxmem=2 ** maxmem, iters=iters, mode=mode,
               pwhash=out.decode('ascii').rstrip(),
               construct=self.construct,
               )

def _genSalt(self):
    sltln = self.rng.randint(self.mnsaltlen, self.mxsaltlen)
    chrs = [self.rng.choice(self.GOODCHARS) for x in range(sltln)]
    return ''.join(chrs)

def _genPw(self):
    pwln = self.rng.randint(self.mnpwlen, self.mxpwlen)
    chrs = [self.rng.choice(self.GOODCHARS) for x in range(pwln)]
    return ''.join(chrs)

def __next__(self):
    if self.count >= self.maxcount:
        raise StopIteration
    psw = self._genPw()
    slt = self._genSalt()

```

(continues on next page)

(continued from previous page)

```

    mem = self.rng.randint(self.mnmem, self.mxmem)
    iters = self.rng.randint(self.mnitters, self.mxitters)
    dgstln = self.rng.randint(self.mndgstln, self.mxdgstln)
    rs = self._runOnce(psw, slt, dgstln, mem, iters)
    self.count += 1
    return rs

def __iter__(self):
    return self

next = __next__

if __name__ == '__main__':

    p = argparse.ArgumentParser()
    p.add_argument('-x', '--executable', dest='exe', required=True)
    p.add_argument('-c', '--construction', dest='construct',
                   type=str, default='argon2i')
    p.add_argument('-v', '--version', dest='version',
                   type=str, default='13')
    p.add_argument('-e', '--encoded', dest='encoded', default=False,
                   action='store_true',)
    p.add_argument('-s', '--min-salt-len', dest='mnsaltlen', type=int,
                   default=8)
    p.add_argument('-S', '--max-salt-len', dest='mxsaltlen', type=int,
                   default=8)
    p.add_argument('-p', '--min-password-len', dest='mnpwlen',
                   type=int, default=16)
    p.add_argument('-P', '--max-password-len', dest='mxpwlen',
                   type=int, default=16)
    p.add_argument('-l', '--min-digest-len', dest='mndgstln',
                   type=int, default=64)
    p.add_argument('-L', '--max-digest-len', dest='mxdgstln',
                   type=int, default=64)
    p.add_argument('-m', '--min-memory-exponent', dest='mnmem',
                   type=int, default=16)
    p.add_argument('-M', '--max-memory-exponent', dest='mxmem',
                   type=int, default=16)
    p.add_argument('-t', '--min-time-opscount', dest='mnitters',
                   type=int, default=3)
    p.add_argument('-T', '--max-time-opscount', dest='mxitters',
                   type=int, default=3)
    p.add_argument('-n', '--count', dest='n', type=int, default=10)
    p.add_argument('-w', '--output', dest='outfile',
                   default=sys.stdout, type=argparse.FileType('w'))

    args = p.parse_args()

    res = [x for x in argonRunner(args)]

    json.dump(res, args.outfile, indent=2, separators=(',', ': '))

```

Blake2b reference vectors

While the blake2b construction is a keyed hash and variable output length algorithm which can optionally be initialized with limited size salt and personalization parameters, the `known answers` json file in the reference `blake2` sources just provides vectors for default length hash with empty salt and personalization.

To fill this test gap, we used both the `pyblake` and the `libsodium` implemented generators provided by `crypto test vectors` for the `blake2b` mechanism to generate twenty vectors in each of `test/data/crypto-test-vectors-blake2-nosalt-nopersonalization.txt` and `test/data/crypto-test-vectors-blake2-salt-personalization.txt`

Vector generation

After cloning the github project with

```
$ git clone https://github.com/jedisct1/crypto-test-vectors.git
```

the needed source files will be available in the `nosalt-nopersonalization` and `salt-personalization` subdirectories of `crypto-test-vectors/crypto/hash/blake2/blake2b/`.

To run the python generators, after ensuring the needed `pyblake2` module is available in the python environment, it will be enough to run the following commands at the shell prompt:

```
$ BLAKE="${PWD}/crypto-test-vectors/crypto/hash/blake2/blake2b"
$ NOPERS="${BLAKE}/nosalt-nopersonalization/generators"
$ PERSON="${BLAKE}/salt-personalization/generators"
$ python "${NOPERS}/pyblake2/generator.py" 10 > py_nopers_vectors
$ python "${PERSON}/pyblake2/generator.py" 10 > py_pers_vectors
```

On linux systems, after installing the required `libsodium` development package, the C-language generators, can get built by running:

```
$ BLAKE="${PWD}/crypto-test-vectors/crypto/hash/blake2/blake2b"
$ NOPERS="${BLAKE}/nosalt-nopersonalization/generators"
$ PERSON="${BLAKE}/salt-personalization/generators"
$ for i in "${NOPERS}/libsodium" "${PERSON}/libsodium"; do (cd "${i}" && make); done
```

and then run by executing:

```
$ BLAKE="${PWD}/crypto-test-vectors/crypto/hash/blake2/blake2b"
$ NOPERS="${BLAKE}/nosalt-nopersonalization/generators"
$ PERSON="${BLAKE}/salt-personalization/generators"
$ "${NOPERS}/libsodium/generator" 10 > py_nopers_vectors_c
$ "${PERSON}/libsodium/generator" 10 > py_pers_vectors_c
```

scrypt reference vectors

`libsodium` exposes both a simplified `scrypt` KDF/password storage API which parametrizes the CPU and memory load in term of a `opslimit` parameter and a `memlimit` one, and a “traditional” low-level API parametrized in terms of a (N, r, p) triple.

While we used the vectors from [RFC 7914](#) to test the traditional API, the simplified API is only implemented by `libsodium`, and therefore we just added a KDF generation check using the `ascii` encoded passphrase “The quick brown fox jumps over the lazy dog.”, and verified the results were the same we could get from the version of `hashlib.scrypt`, as provided in python version 3.6 `stdlib`.

```

>>> import hashlib
>>> import nacl
>>> import nacl.bindings
>>> import nacl.pwhash.scrypt
>>> pick_scrypt_params = nacl.bindings.nacl_bindings_pick_scrypt_params
>>> nacl.pwhash.scrypt.kdf(32,
...                         b'The quick brown fox jumps over the lazy dog.',
...                         b"ef537f25c895bfa782526529a9b63d97",
...                         opslimit=20000, memlimit=100 * (2 ** 20))
b'\x10e>\xc8A8\x11\xde\x07\xf1\x0f\x98EG\xe6)VJ\xd4yN\xae\xd3P\x87yP\x1b\xc7+n*'
>>> n_log2, r, p = pick_scrypt_params(20000, 100 * (2 ** 20))
>>> print(2 ** n_log2, r, p)
1024 8 1
>>> hashlib.scrypt(b'The quick brown fox jumps over the lazy dog.',
...                salt=b"ef537f25c895bfa782526529a9b63d97",
...                n=1024, r=8, p=1, dklen=32)
b'\x10e>\xc8A8\x11\xde\x07\xf1\x0f\x98EG\xe6)VJ\xd4yN\xae\xd3P\x87yP\x1b\xc7+n*'

```

SealedBox reference vectors

Since libsodium's tests do not provide reference data for the SealedBox construction, the implementation is verified with a `sealbox_test_vectors` utility program that produces and checks custom test vectors by making specific calls to libsodium API.

To build the `sealbox_test_vectors` you need a C language compiler, a prebuilt libsodium library more recent than version 1.0.3 and the corresponding include headers.

In a UNIX-like programming environment you should then execute:

```
$ cc -o sealbox_test_vectors sealbox_test_vectors.c -lsodium -lc
```

If you prefer using a locally compiled installation of the bundled sources, refer to *Building the bundled library* and then run:

```
$ cc -o sealbox_test_vectors sealbox_test_vectors.c \
  ${SODIUMINCL} ${SODIUMLIB} -lsodium -lc
```

Vector generation

When called with one or more command line arguments, `sealbox_test_vectors` will generate the number of hex-encoded vectors requested as first argument, with the optional second and third arguments influencing the length of the randomly generated messages:

```
$ ./sealbox_test_vectors 1
XXXX...      XXXX...      <len>:XXXX...      <len>:XXXX...
```

The second argument, if present, sets both a minimum and a maximum length on generated messages, overriding the default 128 bytes values respectively with the supplied value and with twice the supplied value.

The third argument, if present, sets the maximum length of generated messages.

Vector test

When called without command line arguments, `sealbox_test_vectors` will parse and hex-decode the lines given as standard input and check if decoding the encrypted message will return the original message. A “OK”/“FAIL” tag will be appended to the input line to signify success/failure of the test.

To check correct “round-trip” behavior, you can run `sealbox_test_vectors` as a test vector generator against itself:

```
$ ./sealbox_test_vectors 1 | ./sealbox_test_vectors
XXXX...      XXXX... <len>:XXXX...  <len>:XXXX...  OK
```

If you want to check the vectors distributed with PyNaCl’s sources, after setting the environment variable `PYNACL_BASE` to the directory where the unpacked source for PyNaCl has been extracted/cloned, you could run:

```
$ ./sealbox_test_vectors < ${PYNACL_BASE}/tests/data/sealed_box_ref.txt
77076d ... 8c86  OK
```

Source code for the vector checker utility

The source code for `sealbox_test_vectors` is available inside the `docs/vectors/c-source` directory of PyNaCl distribution and can also be directly downloaded from `sealbox_test_vectors.c`.

Listing 2: `sealbox_test_vectors.c`

```
/*
 * Copyright 2017 Donald Stuftt and individual contributors
 *
 * Licensed under the Apache License, Version 2.0 (the "License");
 * you may not use this file except in compliance with the License.
 * You may obtain a copy of the License at
 *
 * http://www.apache.org/licenses/LICENSE-2.0
 *
 * Unless required by applicable law or agreed to in writing, software
 * distributed under the License is distributed on an "AS IS" BASIS,
 * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
 * See the License for the specific language governing permissions and
 * limitations under the License.
 *
 * Test vector generator/checker for libsodium's box_seal APIs
 * to build in a unix-like environment, use a command line like
 * $ cc sealbox_test_vectors.c -I${IPATH} -L${LPATH} -lsodium -o sealbox_test_vectors
 * with IPATH and LPATH defined to respectively point to libsodium's include path
 * and to the directory containing the link library libsodium.a or libsodium.o
 */
#include <stdio.h>
#include <string.h>
#include <sodium.h>

int checkone (char *hxsecret, char *hxpub, size_t pflen, char *hxplaintext,
              size_t crlen, char *hxencrypted) {

    int pklen = crypto_box_PUBLICKEYBYTES;
    int sklen = crypto_box_SECRETKEYBYTES;
```

(continues on next page)

(continued from previous page)

```

char *skr = sodium_malloc (sklen);
char *pub = sodium_malloc (pklen);
char *txt = sodium_malloc (ptlen);
char *crpt = sodium_malloc (crlen);
char *outp = sodium_malloc (ptlen);

int rs = sodium_hex2bin (skr, sklen, hxsecret, 2 * sklen,
                        NULL, NULL, NULL);
rs |= sodium_hex2bin (pub, pklen, hxpub, 2 * pklen, NULL, NULL, NULL);
rs |= sodium_hex2bin (txt, ptlen, hxplaintext, strlen (hxplaintext),
                    NULL, NULL, NULL);
rs |= sodium_hex2bin (crpt, crlen, hxencrypted, strlen (hxencrypted),
                    NULL, NULL, NULL);

if (rs == 0)
    rs = crypto_box_seal_open (outp, crpt, crlen, pub, skr);
if (rs == 0)
    rs = sodium_memcmp (outp, txt, ptlen);

sodium_free (crpt);
sodium_free (txt);
sodium_free (skr);
sodium_free (pub);

return rs;
}

void gentestline (int minmsglen, int maxmsglen) {

int pklen = crypto_box_PUBLICKEYBYTES;
int sklen = crypto_box_SECRETKEYBYTES;
size_t txtlen = minmsglen + randombytes_uniform (maxmsglen - minmsglen + 1);
size_t encrlen = txtlen + crypto_box_SEALBYTES;

char *skr = sodium_malloc (sklen);
char *pub = sodium_malloc (pklen);
char *txt = sodium_malloc (txtlen);
char *crpt = sodium_malloc (encrlen);

crypto_box_keypair (pub, skr);
randombytes_buf (txt, txtlen);

crypto_box_seal (crpt, txt, txtlen, pub);

char *hskr = sodium_malloc (sklen * 2 + 1);
char *hpub = sodium_malloc (pklen * 2 + 1);
char *htxt = sodium_malloc (txtlen * 2 + 1);
char *hkrp = sodium_malloc (encrlen * 2 + 1);

sodium_bin2hex (hskr, sklen * 2 + 1, skr, sklen);
sodium_bin2hex (hpub, pklen * 2 + 1, pub, pklen);
sodium_bin2hex (htxt, txtlen * 2 + 1, txt, txtlen);
sodium_bin2hex (hkrp, encrlen * 2 + 1, crpt, encrlen);

printf ("%s\t%s\t%zu:%s\t%zu:%s\n", hskr, hpub, txtlen, htxt, encrlen, hkrp);
}

```

(continues on next page)

(continued from previous page)

```

int main (int argc, char **argv) {
/*
 * If called without any argument, the resulting executable will
 * read and hex decode the secret and public part of the receiver key,
 * the original plaintext and the ciphertext, and then
 * check if the message resulting from decrypting ciphertext with
 * the secret key is equal to the given plaintext
 *
 * If called with a sequence of integer arguments, sealbox_test_vectors
 * will generate the requested number of reference lines, encrypting
 * random messages.
 *
 */
    if (sodium_init () == -1) {
        exit (1);
    }

    if (argc == 1) {
        size_t lsz = 0;
        char *line = NULL;
        ssize_t lln = 0;
        int res;
        char hxsecret[2 * crypto_box_SECRETKEYBYTES + 1];
        char hxpub[2 * crypto_box_PUBLICKEYBYTES + 1];
        char hxplaintext[2048 + 1];
        char hxencrypted[2048 + 2 * crypto_box_SEALBYTES + 1];
        char cmpplaintext[5 + 2048 + 1];
        char cmpencrypted[5 + 2048 + 2 * crypto_box_SEALBYTES + 1];
        size_t ptlen = 0;
        size_t crlen = 0;

        while (lln = getline (&line, &lsz, stdin) > 0) {
            if (lln > 0) {
                if (strncmp (line, "#", 1) == 0 ||
                    strncmp (line, "\n", 1) == 0 ||
                    strncmp (line, "\r", 1) == 0)
                    continue;

                sscanf (line, "%s%s%s%s",
                        hxsecret, hxpub, cmpplaintext, cmpencrypted);
                sscanf (cmpplaintext, "%zu:%s",
                        &ptlen, hxplaintext);
                sscanf (cmpencrypted, "%zu:%s",
                        &crlen, hxencrypted);
                if (ptlen == 0)
                    memset (hxplaintext, 0, sizeof (hxplaintext));
                if (crlen == 0)
                    memset (hxencrypted, 0, sizeof (hxencrypted));
                res = checkone (hxsecret, hxpub, ptlen, hxplaintext, crlen,
                ↪hxencrypted);
                char *rsstr = (res == 0) ? "OK" : "FAIL";
                printf ("%s\t%s\t%zu:%s\t%zu:%s\t%s\n",
                        hxsecret, hxpub, ptlen, hxplaintext, crlen, hxencrypted,
                ↪rsstr);
            }
            free (line);
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

        line = NULL;
    }
} else {
    int nlines = atoi (argv[1]);
    int minmsgl = 128;
    int maxmsgl = 128;
    if (argc == 3) {
        minmsgl = atoi (argv[2]);
        maxmsgl = atoi (argv[2]) * 2;
    } else if (argc == 4) {
        minmsgl = atoi (argv[2]);
        maxmsgl = atoi (argv[3]);
    }
    for (int i = 0; i < nlines; i++) {
        gentestline (minmsgl, maxmsgl);
    }
}
}

```

secretstream reference vectors

Since libsodium's tests do not provide reference data for the secretstream construction, the implementation is verified with a `secretstream_test_vector` utility program that produces custom test vectors by making specific calls to the libsodium API.

To build the `secretstream_test_vector` you need a C language compiler, a prebuilt libsodium library more recent than version 1.0.14 and the corresponding include headers.

In a UNIX-like programming environment you should then execute:

```
$ cc -o secretstream_test_vector secretstream_test_vector.c -lsodium -lc
```

If you prefer using a locally compiled installation of the bundled sources, refer to *Building the bundled library* and then run:

```
$ cc -o secretstream_test_vector secretstream_test_vector.c \
    ${SODIUMINCL} ${SODIUMLIB} -lsodium -lc
```

Vector generation

```
$ ./secretstream_test_vector -h
Usage: secretstream_test_vector [-c num_chunks] [-r]
```

When called, the program will output a JSON dictionary containing `key`, `header`, and `chunks`. The `chunks` is a list of individual messages passed to `crypto_secretstream_xchacha20poly1305_push` containing `tag`, `message`, `ad` and `ciphertext` keys.

Source code for the vector checker utility

The source code for `secretstream_test_vector` is available inside the `docs/vectors/c-source` directory of PyNaCl distribution and can also be directly downloaded from `secretstream_test_vector.c`.

Listing 3: secretstream_test_vector.c

```

/*
 * Copyright 2018 Donald Stufft and individual contributors
 *
 * Licensed under the Apache License, Version 2.0 (the "License");
 * you may not use this file except in compliance with the License.
 * You may obtain a copy of the License at
 *
 * http://www.apache.org/licenses/LICENSE-2.0
 *
 * Unless required by applicable law or agreed to in writing, software
 * distributed under the License is distributed on an "AS IS" BASIS,
 * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
 * See the License for the specific language governing permissions and
 * limitations under the License.
 *
 * Test vector generator/checker for libsodium's crypto_secretstream APIs
 * to build in a unix-like environment, use a command line like
 * $ cc secretstream_test_vector.c \
 *     -I${IPATH} -L${LPATH} -lsodium \
 *     -o secretstream_test_vector
 * with IPATH and LPATH defined to respectively point to libsodium's include path
 * and to the directory containing the link library libsodium.a or libsodium.o
 */
#include <stdio.h>
#include <string.h>
#include <sodium.h>
#include <unistd.h>

#define MAX_AD_SIZE 32
#define MAX_CHUNK_SIZE 512
#define CHK(cmd) \
    do { if ((rc = (cmd)) != 0) { \
        fprintf(stderr, "api call failed, code=%d", rc); \
        exit(1); \
    }} while(0)

int usage(int argc, char **argv) {
    fprintf(stderr, "Usage: %s [-c num_chunks] [-r]\n", argv[0]);
    return 1;
}

int main (int argc, char **argv) {
    int c, rc;
    int num_chunks = 1;
    int rekey = 0;

    crypto_secretstream_xchacha20poly1305_state state;
    unsigned char header[crypto_secretstream_xchacha20poly1305_HEADERBYTES];
    unsigned char key[crypto_secretstream_xchacha20poly1305_KEYBYTES];
    unsigned char m[MAX_CHUNK_SIZE];
    unsigned char ad[MAX_AD_SIZE];
    unsigned char ct[MAX_CHUNK_SIZE + crypto_secretstream_xchacha20poly1305_ABYTES];
    char key_hex[sizeof(key) * 2 + 1];
    char header_hex[sizeof(header) * 2 + 1];

```

(continues on next page)

(continued from previous page)

```

char m_hex[sizeof(m) * 2 + 1];
char ad_hex[sizeof(ad) * 2 + 1];
char ct_hex[sizeof(ct) * 2 + 1];
unsigned long long m_len, ad_len, ct_len;
unsigned char tag;

while ((c = getopt(argc, argv, "hc:r")) != -1) {
    switch (c) {
        case 'c':
            num_chunks = atoi(optarg);
            break;
        case 'r':
            rekey = 1;
            break;
        case 'h':
            return usage(argc, argv);
        default:
            return 1;
    }
}
if (optind < argc) return usage(argc, argv);

if (sodium_init() == -1) {
    exit(1);
}

/* output format:
 * {
 *   "key": "hex",
 *   "header": "hex",
 *   "chunks": [
 *     {
 *       "tag": 0,
 *       "ad": "hex",
 *       "message": "hex",
 *       "ciphertext": "hex"
 *     },
 *     ...
 *   ]
 * }
 */

crypto_secretstream_xchacha20poly1305_keygen(key);
CHK(crypto_secretstream_xchacha20poly1305_init_push(&state, header, key));
sodium_bin2hex(key_hex, sizeof key_hex, key, sizeof key);
sodium_bin2hex(header_hex, sizeof header_hex, header, sizeof header);
printf("{\n  \"key\": \"%s\", \n  \"header\": \"%s\", \n  \"chunks\": [\n",
        key_hex, header_hex);
for (c = 1 ; c <= num_chunks ; ++c) {
    tag =
        c == num_chunks ? crypto_secretstream_xchacha20poly1305_TAG_FINAL
        : rekey ? crypto_secretstream_xchacha20poly1305_TAG_REKEY
        : crypto_secretstream_xchacha20poly1305_TAG_MESSAGE;
    ad_len = randombytes_uniform(MAX_AD_SIZE);
    m_len = randombytes_uniform(MAX_CHUNK_SIZE - 1) + 1;
    randombytes_buf(m, m_len);
    randombytes_buf(ad, ad_len);
}
}

```

(continues on next page)

(continued from previous page)

```

CHK(crypto_secretstream_xchacha20poly1305_push(
    &state, ct, &ct_len, m, m_len, ad, ad_len, tag));
sodium_bin2hex(m_hex, m_len * 2 + 1, m, m_len);
if (ad_len > 0) {
    sodium_bin2hex(ad_hex, ad_len * 2 + 1, ad, ad_len);
}
sodium_bin2hex(ct_hex, ct_len * 2 + 1, ct, ct_len);
printf("    {\n"
       "        \"tag\": %d,\n          \"ad\": %s%s%s,\n"
       "        \"message\": \"%s\",\n        \"ciphertext\": \"%s\"\n"
       "    }%s\n",
       tag,
       ad_len > 0 ? "\"" : "",
       ad_len > 0 ? ad_hex : "null",
       ad_len > 0 ? "\"" : "",
       m_hex,
       ct_hex,
       c < num_chunks ? ", " : "");
}
printf(" ]\n}\n");

return 0;
}

```

Building the bundled library

If you want to avoid a system-wide installation of libsodium's development files just for compiling and running the tests, you can instead install the library and header files inside PyNaCl's sources.

Linux systems

On Linux (and presumably other UNIX-like systems), after entering the PyNaCl source directory you must execute the following commands:

```

$ mkdir -p build/libsodium
$ cd build/libsodium
$ ../../src/libsodium/configure --prefix=$PWD --disable-shared
$ make
$ make install
$ cd ../../

```

If all went well,

```
$ ls build/libsodium/{lib,include}
```

should generate something like the following output:

```

build/libsodium/include:
sodium sodium.h

build/libsodium/lib:
libsodium.a libsodium.la pkgconfig

```

If you now define and export the

```
$ SODIUMINCL="-I${PWD}/build/libsodium/include"
$ export SODIUMINCL
$ SODIUMLIB="-L${PWD}/build/libsodium/lib"
$ export SODIUMLIB
```

environment variables, you can instruct the compiler to use the just-installed library by simply dereferencing the path flags on the c compiler command line

```
$ cc ${SODIUMINCL} ${SODIUMLIB}
```

3.10 Changelog

3.10.1 1.4.0 (UNRELEASED)

- Add low level bindings for extracting the seed and the public key from `crypto_sign_ed25519` secret key
- Add `wheel` and `setuptools` `setup_requirements` in `setup.py` (#485)
- Fix checks on very slow builders (#481, #495)

3.10.2 1.3.0 2018-09-26

- Added support for Python 3.7.
- Update `libsodium` to 1.0.16.
- Run and test all code examples in PyNaCl docs through sphinx's doctest builder.
- Add low-level bindings for chacha20-poly1305 AEAD constructions.
- Add low-level bindings for the chacha20-poly1305 secretstream constructions.
- Add low-level bindings for ed25519ph pre-hashed signing construction.
- Add low-level bindings for constant-time increment and addition on fixed-precision big integers represented as little-endian byte sequences.
- Add low-level bindings for the ISO/IEC 7816-4 compatible padding API.
- Add low-level bindings for `libsodium`'s `crypto_kx...` key exchange construction.
- Set hypothesis deadline to `None` in `tests/test_pwhash.py` to avoid incorrect test failures on slower processor architectures. GitHub issue #370

3.10.3 1.2.1 - 2017-12-04

- Update hypothesis minimum allowed version.
- Infrastructure: add proper configuration for `readthedocs` builder runtime environment.

3.10.4 1.2.0 - 2017-11-01

- Update `libsodium` to 1.0.15.
- Infrastructure: add `jenkins` support for automatic build of `manylinux1` binary wheels

- Added support for `SealedBox` construction.
- Added support for `argon2i` and `argon2id` password hashing constructs and restructured high-level password hashing implementation to expose the same interface for all hashers.
- Added support for 128 bit `siphashx24` variant of `siphash24`.
- Added support for `from_seed` APIs for X25519 keypair generation.
- Dropped support for Python 3.3.

3.10.5 1.1.2 - 2017-03-31

- reorder link time library search path when using bundled libsodium

3.10.6 1.1.1 - 2017-03-15

- Fixed a circular import bug in `nacl.utils`.

3.10.7 1.1.0 - 2017-03-14

- Dropped support for Python 2.6.
- Added `shared_key()` method on `Box`.
- You can now pass `None` to `nonce` when encrypting with `Box` or `SecretBox` and it will automatically generate a random nonce.
- Added support for `siphash24`.
- Added support for `blake2b`.
- Added support for `scrypt`.
- Update `libsodium` to 1.0.11.
- Default to the bundled `libsodium` when compiling.
- All raised exceptions are defined mixing-in `nacl.exceptions.CryptoError`

3.10.8 1.0.1 - 2016-01-24

- Fix an issue with absolute paths that prevented the creation of wheels.

3.10.9 1.0 - 2016-01-23

- PyNaCl has been ported to use the new APIs available in `ffi 1.0+`. Due to this change we no longer support PyPy releases older than 2.6.
- Python 3.2 support has been dropped.
- Functions to convert between `Ed25519` and `Curve25519` keys have been added.

3.10.10 0.3.0 - 2015-03-04

- The low-level API (*nacl.c.**) has been changed to match the upstream NaCl C/C++ conventions (as well as those of other NaCl bindings). The order of arguments and return values has changed significantly. To avoid silent failures, *nacl.c* has been removed, and replaced with *nacl.bindings* (with the new argument ordering). If you have code which calls these functions (e.g. *nacl.c.crypto_box_keypair()*), you must review the new docstrings and update your code/imports to match the new conventions.

3.11 Indices and tables

- [genindex](#)
- [modindex](#)
- [search](#)

Bibliography

- [SD2012] A nice overview of password hashing history is available in Solar Designer's presentation [Password security: past, present, future](#)
- [PHC] The Argon2 recommendation is prominently shown in the [Password Hashing Competition](#) site, along to the special recognition shortlist and the original call for submissions.

n

`nacl.pwhash`, 28

`nacl.pwhash.argon2i`, 31

`nacl.pwhash.argon2id`, 30

`nacl.pwhash.scrypt`, 32

A

AssertionError (*built-in class*), 26

B

BadSignatureError (*built-in class*), 26
 Base16Encoder (*class in nacl.encoding*), 25
 Base32Encoder (*class in nacl.encoding*), 25
 Base64Encoder (*class in nacl.encoding*), 25
 blake2b (*class in nacl.hashlib*), 33
 blake2b () (*in module nacl.hash*), 27
 Box (*class in nacl.public*), 8
 BYTES_MAX (*in module nacl.pwhash.argon2i*), 32
 BYTES_MAX (*in module nacl.pwhash.argon2id*), 31
 BYTES_MAX (*in module nacl.pwhash.scrypt*), 33
 BYTES_MIN (*in module nacl.pwhash.argon2i*), 32
 BYTES_MIN (*in module nacl.pwhash.argon2id*), 31
 BYTES_MIN (*in module nacl.pwhash.scrypt*), 33

C

ciphertext (*nacl.utils.EncryptedMessage attribute*),
 27
 CryptoError (*built-in class*), 26

D

decode () (*nacl.public.Box class method*), 8
 decrypt () (*nacl.public.Box method*), 8
 decrypt () (*nacl.public.SealedBox method*), 9
 decrypt () (*nacl.secret.SecretBox method*), 12

E

encrypt () (*nacl.public.Box method*), 8
 encrypt () (*nacl.public.SealedBox method*), 9
 encrypt () (*nacl.secret.SecretBox method*), 11
 EncryptedMessage (*class in nacl.utils*), 27
 ensure () (*in module nacl.utils*), 27

G

generate () (*nacl.public.PrivateKey class method*), 8

generate () (*nacl.signing.SigningKey class method*),
 13

H

HexEncoder (*class in nacl.encoding*), 25

I

InvalidkeyError (*built-in class*), 26

K

kdf () (*in module nacl.pwhash.argon2i*), 31
 kdf () (*in module nacl.pwhash.argon2id*), 30
 kdf () (*in module nacl.pwhash.scrypt*), 32

M

MAX_DIGEST_SIZE (*nacl.hashlib.blake2b attribute*),
 33
 MAX_KEY_SIZE (*nacl.hashlib.blake2b attribute*), 33
 MEMLIMIT_INTERACTIVE (*in module nacl.pwhash*),
 29
 MEMLIMIT_INTERACTIVE (*in module*
nacl.pwhash.argon2i), 32
 MEMLIMIT_INTERACTIVE (*in module*
nacl.pwhash.argon2id), 31
 MEMLIMIT_INTERACTIVE (*in module*
nacl.pwhash.scrypt), 33
 MEMLIMIT_MAX (*in module nacl.pwhash*), 29
 MEMLIMIT_MAX (*in module nacl.pwhash.argon2i*), 32
 MEMLIMIT_MAX (*in module nacl.pwhash.argon2id*), 31
 MEMLIMIT_MAX (*in module nacl.pwhash.scrypt*), 33
 MEMLIMIT_MIN (*in module nacl.pwhash*), 29
 MEMLIMIT_MIN (*in module nacl.pwhash.argon2i*), 32
 MEMLIMIT_MIN (*in module nacl.pwhash.argon2id*), 31
 MEMLIMIT_MIN (*in module nacl.pwhash.scrypt*), 33
 MEMLIMIT_MODERATE (*in module nacl.pwhash*), 30
 MEMLIMIT_MODERATE (*in module*
nacl.pwhash.argon2i), 32
 MEMLIMIT_MODERATE (*in module*
nacl.pwhash.argon2id), 31

MEMLIMIT_MODERATE (in module *nacl.pwhash.scrypt*), 33
 MEMLIMIT_SENSITIVE (in module *nacl.pwhash*), 29
 MEMLIMIT_SENSITIVE (in module *nacl.pwhash.argon2i*), 32
 MEMLIMIT_SENSITIVE (in module *nacl.pwhash.argon2id*), 31
 MEMLIMIT_SENSITIVE (in module *nacl.pwhash.scrypt*), 33
 message (*nacl.signing.SignedMessage* attribute), 14

N

nacl.pwhash (module), 28
nacl.pwhash.argon2i (module), 31
nacl.pwhash.argon2id (module), 30
nacl.pwhash.scrypt (module), 32
 nonce (*nacl.utils.EncryptedMessage* attribute), 27

O

OPSLIMIT_INTERACTIVE (in module *nacl.pwhash*), 29
 OPSLIMIT_INTERACTIVE (in module *nacl.pwhash.argon2i*), 32
 OPSLIMIT_INTERACTIVE (in module *nacl.pwhash.argon2id*), 31
 OPSLIMIT_INTERACTIVE (in module *nacl.pwhash.scrypt*), 33
 OPSLIMIT_MAX (in module *nacl.pwhash*), 29
 OPSLIMIT_MAX (in module *nacl.pwhash.argon2i*), 32
 OPSLIMIT_MAX (in module *nacl.pwhash.argon2id*), 31
 OPSLIMIT_MAX (in module *nacl.pwhash.scrypt*), 33
 OPSLIMIT_MIN (in module *nacl.pwhash*), 29
 OPSLIMIT_MIN (in module *nacl.pwhash.argon2i*), 32
 OPSLIMIT_MIN (in module *nacl.pwhash.argon2id*), 31
 OPSLIMIT_MIN (in module *nacl.pwhash.scrypt*), 33
 OPSLIMIT_MODERATE (in module *nacl.pwhash*), 30
 OPSLIMIT_MODERATE (in module *nacl.pwhash.argon2i*), 32
 OPSLIMIT_MODERATE (in module *nacl.pwhash.argon2id*), 31
 OPSLIMIT_MODERATE (in module *nacl.pwhash.scrypt*), 33
 OPSLIMIT_SENSITIVE (in module *nacl.pwhash*), 29
 OPSLIMIT_SENSITIVE (in module *nacl.pwhash.argon2i*), 32
 OPSLIMIT_SENSITIVE (in module *nacl.pwhash.argon2id*), 31
 OPSLIMIT_SENSITIVE (in module *nacl.pwhash.scrypt*), 33

P

PASSWD_MAX (in module *nacl.pwhash*), 29
 PASSWD_MAX (in module *nacl.pwhash.argon2i*), 32
 PASSWD_MAX (in module *nacl.pwhash.argon2id*), 31

PASSWD_MAX (in module *nacl.pwhash.scrypt*), 33
 PASSWD_MIN (in module *nacl.pwhash*), 29
 PASSWD_MIN (in module *nacl.pwhash.argon2i*), 32
 PASSWD_MIN (in module *nacl.pwhash.argon2id*), 31
 PASSWD_MIN (in module *nacl.pwhash.scrypt*), 33
 PERSON_SIZE (*nacl.hashlib.blake2b* attribute), 33
 PrivateKey (class in *nacl.public*), 7
 public_key (*nacl.public.PrivateKey* attribute), 8
 PublicKey (class in *nacl.public*), 7
 PWHASH_SIZE (in module *nacl.pwhash*), 29
 PWHASH_SIZE (in module *nacl.pwhash.argon2i*), 32
 PWHASH_SIZE (in module *nacl.pwhash.argon2id*), 31
 PWHASH_SIZE (in module *nacl.pwhash.scrypt*), 33

R

random () (in module *nacl.utils*), 27
 RawEncoder (class in *nacl.encoding*), 25
 RuntimeError (built-in class), 26

S

SALT_SIZE (*nacl.hashlib.blake2b* attribute), 33
 SALTBYTES (in module *nacl.pwhash.argon2i*), 32
 SALTBYTES (in module *nacl.pwhash.argon2id*), 31
 SALTBYTES (in module *nacl.pwhash.scrypt*), 33
 script () (in module *nacl.hashlib*), 34
 SealedBox (class in *nacl.public*), 9
 SecretBox (class in *nacl.secret*), 11
 sha256 () (in module *nacl.hash*), 27
 sha512 () (in module *nacl.hash*), 27
 shared_key () (*nacl.public.Box* method), 8
 sign () (*nacl.signing.SigningKey* method), 13
 signature (*nacl.signing.SignedMessage* attribute), 14
 SignedMessage (class in *nacl.signing*), 14
 SigningKey (class in *nacl.signing*), 13
 siphash24 () (in module *nacl.hash*), 28
 siphashx24 () (in module *nacl.hash*), 28
 str () (in module *nacl.pwhash*), 28
 str () (in module *nacl.pwhash.argon2i*), 31
 str () (in module *nacl.pwhash.argon2id*), 30
 str () (in module *nacl.pwhash.scrypt*), 32

T

TypeError (built-in class), 26

U

URLSafeBase64Encoder (class in *nacl.encoding*), 25

V

ValueError (built-in class), 26
 verify () (in module *nacl.pwhash*), 29
 verify () (in module *nacl.pwhash.argon2i*), 31
 verify () (in module *nacl.pwhash.argon2id*), 30

`verify()` (in module `nacl.pwhash.scrypt`), 32
`verify()` (`nacl.signing.VerifyKey` method), 14
`verify_key` (`nacl.signing.SigningKey` attribute), 13
`VerifyKey` (class in `nacl.signing`), 14