

---

# **kafka-python Documentation**

***Release 1.2.5***

**Dana Powers**

July 17, 2016



<b>1</b>	<b>KafkaConsumer</b>	<b>3</b>
<b>2</b>	<b>KafkaProducer</b>	<b>5</b>
<b>3</b>	<b>Compression</b>	<b>7</b>
<b>4</b>	<b>Protocol</b>	<b>9</b>
<b>5</b>	<b>Low-level</b>	<b>11</b>
5.1	Usage . . . . .	11
5.2	kafka-python API . . . . .	13
5.3	Simple APIs (DEPRECATED) . . . . .	26
5.4	Install . . . . .	29
5.5	Tests . . . . .	30
5.6	Compatibility . . . . .	31
5.7	Support . . . . .	31
5.8	License . . . . .	32
5.9	Changelog . . . . .	32



Python client for the Apache Kafka distributed stream processing system. kafka-python is designed to function much like the official java client, with a sprinkling of pythonic interfaces (e.g., consumer iterators).

kafka-python is best used with newer brokers (0.10 or 0.9), but is backwards-compatible with older versions (to 0.8.0). Some features will only be enabled on newer brokers, however; for example, fully coordinated consumer groups – i.e., dynamic partition assignment to multiple consumers in the same group – requires use of 0.9 kafka brokers. Supporting this feature for earlier broker releases would require writing and maintaining custom leadership election and membership / health check code (perhaps using zookeeper or consul). For older brokers, you can achieve something similar by manually assigning different partitions to each consumer instance with config management tools like chef, ansible, etc. This approach will work fine, though it does not support rebalancing on failures. See [Compatibility](#) for more details.

Please note that the master branch may contain unreleased features. For release documentation, please see [readthedocs](#) and/or python's inline help.

```
>>> pip install kafka-python
```



---

## KafkaConsumer

---

`KafkaConsumer` is a high-level message consumer, intended to operate as similarly as possible to the official java client. Full support for coordinated consumer groups requires use of kafka brokers that support the Group APIs: kafka v0.9+.

See `KafkaConsumer` for API and configuration details.

The consumer iterator returns `ConsumerRecords`, which are simple `namedtuples` that expose basic message attributes: topic, partition, offset, key, and value:

```
>>> from kafka import KafkaConsumer
>>> consumer = KafkaConsumer('my_favorite_topic')
>>> for msg in consumer:
...     print (msg)
```

```
>>> # manually assign the partition list for the consumer
>>> from kafka import TopicPartition
>>> consumer = KafkaConsumer(bootstrap_servers='localhost:1234')
>>> consumer.assign([TopicPartition('foobar', 2)])
>>> msg = next(consumer)
```

```
>>> # Deserialize msgpack-encoded values
>>> consumer = KafkaConsumer(value_deserializer=msgpack.loads)
>>> consumer.subscribe(['msgpackfoo'])
>>> for msg in consumer:
...     assert isinstance(msg.value, dict)
```





---

## KafkaProducer

---

`KafkaProducer` is a high-level, asynchronous message producer. The class is intended to operate as similarly as possible to the official java client. See `KafkaProducer` for more details.

```
>>> from kafka import KafkaProducer
>>> producer = KafkaProducer(bootstrap_servers='localhost:1234')
>>> for _ in range(100):
...     producer.send('foobar', b'some_message_bytes')
```

```
>>> # Block until all pending messages are sent
>>> producer.flush()
```

```
>>> # Block until a single message is sent (or timeout)
>>> producer.send('foobar', b'another_message').get(timeout=60)
```

```
>>> # Use a key for hashed-partitioning
>>> producer.send('foobar', key=b'foo', value=b'bar')
```

```
>>> # Serialize json messages
>>> import json
>>> producer = KafkaProducer(value_serializer=lambda v: json.dumps(v).encode('utf-8'))
>>> producer.send('fizzbuzz', {'foo': 'bar'})
```

```
>>> # Serialize string keys
>>> producer = KafkaProducer(key_serializer=str.encode)
>>> producer.send('flipflap', key='ping', value=b'1234')
```

```
>>> # Compress messages
>>> producer = KafkaProducer(compression_type='gzip')
>>> for i in range(1000):
...     producer.send('foobar', b'msg %d' % i)
```



---

## **Compression**

---

kafka-python supports gzip compression/decompression natively. To produce or consume lz4 compressed messages, you must install lz4tools and xxhash (modules may not work on python2.6). To enable snappy, install python-snappy (also requires snappy library). See Installation for more information.



---

### Protocol

---

A secondary goal of kafka-python is to provide an easy-to-use protocol layer for interacting with kafka brokers via the python repl. This is useful for testing, probing, and general experimentation. The protocol support is leveraged to enable a `check_version()` method that probes a kafka broker and attempts to identify which version it is running (0.8.0 to 0.10).



Legacy support is maintained for low-level consumer and producer classes, SimpleConsumer and SimpleProducer.

## 5.1 Usage

### 5.1.1 KafkaConsumer

```
from kafka import KafkaConsumer

# To consume latest messages and auto-commit offsets
consumer = KafkaConsumer('my-topic',
                          group_id='my-group',
                          bootstrap_servers=['localhost:9092'])

for message in consumer:
    # message value and key are raw bytes -- decode if necessary!
    # e.g., for unicode: `message.value.decode('utf-8')`
    print ("%s:%d:%d: key=%s value=%s" % (message.topic, message.partition,
                                          message.offset, message.key,
                                          message.value))

# consume earliest available messages, dont commit offsets
KafkaConsumer(auto_offset_reset='earliest', enable_auto_commit=False)

# consume json messages
KafkaConsumer(value_deserializer=lambda m: json.loads(m.decode('ascii'))))

# consume msgpack
KafkaConsumer(value_deserializer=msgpack.unpackb)

# StopIteration if no message after 1sec
KafkaConsumer(consumer_timeout_ms=1000)

# Subscribe to a regex topic pattern
consumer = KafkaConsumer()
consumer.subscribe(pattern='^awesome.*')

# Use multiple consumers in parallel w/ 0.9 kafka brokers
# typically you would run each on a different server / process / CPU
consumer1 = KafkaConsumer('my-topic',
                           group_id='my-group',
                           bootstrap_servers='my.server.com')
```

```
consumer2 = KafkaConsumer('my-topic',
                           group_id='my-group',
                           bootstrap_servers='my.server.com')
```

There are many configuration options for the consumer class. See `KafkaConsumer` API documentation for more details.

## 5.1.2 KafkaProducer

```
from kafka import KafkaProducer
from kafka.errors import KafkaError

producer = KafkaProducer(bootstrap_servers=['broker1:1234'])

# Asynchronous by default
future = producer.send('my-topic', b'raw_bytes')

# Block for 'synchronous' sends
try:
    record_metadata = future.get(timeout=10)
except KafkaError:
    # Decide what to do if produce request failed...
    log.exception()
    pass

# Successful result returns assigned partition and offset
print (record_metadata.topic)
print (record_metadata.partition)
print (record_metadata.offset)

# produce keyed messages to enable hashed partitioning
producer.send('my-topic', key=b'foo', value=b'bar')

# encode objects via msgpack
producer = KafkaProducer(value_serializer=msgpack.dumps)
producer.send('msgpack-topic', {'key': 'value'})

# produce json messages
producer = KafkaProducer(value_serializer=lambda m: json.dumps(m).encode('ascii'))
producer.send('json-topic', {'key': 'value'})

# produce asynchronously
for _ in range(100):
    producer.send('my-topic', b'msg')

# block until all async messages are sent
producer.flush()

# configure multiple retries
producer = KafkaProducer(retries=5)
```



## 5.2 kafka-python API

### 5.2.1 KafkaConsumer

`class kafka.KafkaConsumer(*topics, **configs)`

Consume records from a Kafka cluster.

The consumer will transparently handle the failure of servers in the Kafka cluster, and adapt as topic-partitions are created or migrate between brokers. It also interacts with the assigned kafka Group Coordinator node to allow multiple consumers to load balance consumption of topics (requires kafka >= 0.9.0.0).

**Parameters** *\*topics* (*str*) – optional list of topics to subscribe to. If not set, call `subscribe()` or `assign()` before consuming records.

#### Keyword Arguments

- **bootstrap\_servers** – ‘host[:port]’ string (or list of ‘host[:port]’ strings) that the consumer should contact to bootstrap initial cluster metadata. This does not have to be the full node list. It just needs to have at least one broker that will respond to a Metadata API Request. Default port is 9092. If no servers are specified, will default to localhost:9092.
- **client\_id** (*str*) – a name for this client. This string is passed in each request to servers and can be used to identify specific server-side log entries that correspond to this client. Also submitted to GroupCoordinator for logging with respect to consumer group administration. Default: ‘kafka-python-{version}’
- **group\_id** (*str or None*) – name of the consumer group to join for dynamic partition assignment (if enabled), and to use for fetching and committing offsets. If None, auto-partition assignment (via group coordinator) and offset commits are disabled. Default: ‘kafka-python-default-group’
- **key\_deserializer** (*callable*) – Any callable that takes a raw message key and returns a deserialized key.
- **value\_deserializer** (*callable*) – Any callable that takes a raw message value and returns a deserialized value.
- **fetch\_min\_bytes** (*int*) – Minimum amount of data the server should return for a fetch request, otherwise wait up to `fetch_max_wait_ms` for more data to accumulate. Default: 1.
- **fetch\_max\_wait\_ms** (*int*) – The maximum amount of time in milliseconds the server will block before answering the fetch request if there isn’t sufficient data to immediately satisfy the requirement given by `fetch_min_bytes`. Default: 500.
- **max\_partition\_fetch\_bytes** (*int*) – The maximum amount of data per-partition the server will return. The maximum total memory used for a request = `#partitions * max_partition_fetch_bytes`. This size must be at least as large as the maximum message size the server allows or else it is possible for the producer to send messages larger than the consumer can fetch. If that happens, the consumer can get stuck trying to fetch a large message on a certain partition. Default: 1048576.
- **request\_timeout\_ms** (*int*) – Client request timeout in milliseconds. Default: 40000.
- **retry\_backoff\_ms** (*int*) – Milliseconds to backoff when retrying on errors. Default: 100.
- **reconnect\_backoff\_ms** (*int*) – The amount of time in milliseconds to wait before attempting to reconnect to a given host. Default: 50.
- **max\_in\_flight\_requests\_per\_connection** (*int*) – Requests are pipelined to kafka brokers up to this number of maximum requests per broker connection. Default: 5.

- **auto\_offset\_reset** (*str*) – A policy for resetting offsets on `OffsetOutOfRange` errors: ‘earliest’ will move to the oldest available message, ‘latest’ will move to the most recent. Any other value will raise the exception. Default: ‘latest’.
- **enable\_auto\_commit** (*bool*) – If true the consumer’s offset will be periodically committed in the background. Default: True.
- **auto\_commit\_interval\_ms** (*int*) – milliseconds between automatic offset commits, if `enable_auto_commit` is True. Default: 5000.
- **default\_offset\_commit\_callback** (*callable*) – called as `callback(offsets, response)` response will be either an Exception or a `OffsetCommitResponse` struct. This callback can be used to trigger custom actions when a commit request completes.
- **check\_crcs** (*bool*) – Automatically check the CRC32 of the records consumed. This ensures no on-the-wire or on-disk corruption to the messages occurred. This check adds some overhead, so it may be disabled in cases seeking extreme performance. Default: True
- **metadata\_max\_age\_ms** (*int*) – The period of time in milliseconds after which we force a refresh of metadata even if we haven’t seen any partition leadership changes to proactively discover any new brokers or partitions. Default: 300000
- **partition\_assignment\_strategy** (*list*) – List of objects to use to distribute partition ownership amongst consumer instances when group management is used. Default: [`RangePartitionAssignor`, `RoundRobinPartitionAssignor`]
- **heartbeat\_interval\_ms** (*int*) – The expected time in milliseconds between heartbeats to the consumer coordinator when using Kafka’s group management feature. Heartbeats are used to ensure that the consumer’s session stays active and to facilitate rebalancing when new consumers join or leave the group. The value must be set lower than `session_timeout_ms`, but typically should be set no higher than 1/3 of that value. It can be adjusted even lower to control the expected time for normal rebalances. Default: 3000
- **session\_timeout\_ms** (*int*) – The timeout used to detect failures when using Kafka’s group management facilities. Default: 30000
- **send\_buffer\_bytes** (*int*) – The size of the TCP send buffer (`SO_SNDBUF`) to use when sending data. Default: None (relies on system defaults). The java client defaults to 131072.
- **receive\_buffer\_bytes** (*int*) – The size of the TCP receive buffer (`SO_RCVBUF`) to use when reading data. Default: None (relies on system defaults). The java client defaults to 32768.
- **consumer\_timeout\_ms** (*int*) – number of milliseconds to block during message iteration before raising `StopIteration` (i.e., ending the iterator). Default -1 (block forever).
- **skip\_double\_compressed\_messages** (*bool*) – A bug in `KafkaProducer` <= 1.2.4 caused some messages to be corrupted via double-compression. By default, the fetcher will return these messages as a compressed blob of bytes with a single offset, i.e. how the message was actually published to the cluster. If you prefer to have the fetcher automatically detect corrupt messages and skip them, set this option to True. Default: False.
- **security\_protocol** (*str*) – Protocol used to communicate with brokers. Valid values are: `PLAINTEXT`, `SSL`. Default: `PLAINTEXT`.
- **ssl\_context** (*ssl.SSLContext*) – pre-configured `SSLContext` for wrapping socket connections. If provided, all other `ssl_*` configurations will be ignored. Default: None.
- **ssl\_check\_hostname** (*bool*) – flag to configure whether ssl handshake should verify that the certificate matches the brokers hostname. default: true.
- **ssl\_cafile** (*str*) – optional filename of ca file to use in certificate verification. default: none.

- **ssl\_certfile** (*str*) – optional filename of file in pem format containing the client certificate, as well as any ca certificates needed to establish the certificate’s authenticity. default: none.
- **ssl\_keyfile** (*str*) – optional filename containing the client private key. default: none.
- **ssl\_crlfile** (*str*) – optional filename containing the CRL to check for certificate expiration. By default, no CRL check is done. When providing a file, only the leaf certificate will be checked against this CRL. The CRL can only be checked with Python 3.4+ or 2.7.9+. default: none.
- **api\_version** (*str*) – specify which kafka API version to use. 0.9 enables full group coordination features; 0.8.2 enables kafka-storage offset commits; 0.8.1 enables zookeeper-storage offset commits; 0.8.0 is what is left. If set to ‘auto’, will attempt to infer the broker version by probing various APIs. Default: auto
- **api\_version\_auto\_timeout\_ms** (*int*) – number of milliseconds to throw a timeout exception from the constructor when checking the broker api version. Only applies if api\_version set to ‘auto’
- **metric\_reporters** (*list*) – A list of classes to use as metrics reporters. Implementing the AbstractMetricsReporter interface allows plugging in classes that will be notified of new metric creation. Default: []
- **metrics\_num\_samples** (*int*) – The number of samples maintained to compute metrics. Default: 2
- **metrics\_sample\_window\_ms** (*int*) – The number of samples maintained to compute metrics. Default: 30000

---

**Note:** Configuration parameters are described in more detail at <https://kafka.apache.org/090/configuration.html#newconsumerconfigs>

---

#### **assign** (*partitions*)

Manually assign a list of TopicPartitions to this consumer.

**Parameters** **partitions** (*list of TopicPartition*) – assignment for this instance.

**Raises** `IllegalStateException` – if consumer has already called `subscribe()`

**Warning:** It is not possible to use both manual partition assignment with `assign()` and group assignment with `subscribe()`.

---

**Note:** This interface does not support incremental assignment and will replace the previous assignment (if there was one).

---

---

**Note:** Manual topic assignment through this method does not use the consumer’s group management functionality. As such, there will be no rebalance operation triggered when group membership or cluster and topic metadata change.

---

#### **assignment** ()

Get the TopicPartitions currently assigned to this consumer.

If partitions were directly assigned using `assign()`, then this will simply return the same partitions that were previously assigned. If topics were subscribed using `subscribe()`, then this will give the set of topic

partitions currently assigned to the consumer (which may be none if the assignment hasn't happened yet, or if the partitions are in the process of being reassigned).

**Returns** {TopicPartition, ...}

**Return type** set

**close()**

Close the consumer, waiting indefinitely for any needed cleanup.

**commit** (*offsets=None*)

Commit offsets to kafka, blocking until success or error

This commits offsets only to Kafka. The offsets committed using this API will be used on the first fetch after every rebalance and also on startup. As such, if you need to store offsets in anything other than Kafka, this API should not be used. To avoid re-processing the last message read if a consumer is restarted, the committed offset should be the next message your application should consume, i.e.: `last_offset + 1`.

Blocks until either the commit succeeds or an unrecoverable error is encountered (in which case it is thrown to the caller).

Currently only supports kafka-topic offset storage (not zookeeper)

**Parameters** **offsets** (*dict, optional*) – {TopicPartition: OffsetAndMetadata} dict to commit with the configured `group_id`. Defaults to current consumed offsets for all subscribed partitions.

**commit\_async** (*offsets=None, callback=None*)

Commit offsets to kafka asynchronously, optionally firing callback

This commits offsets only to Kafka. The offsets committed using this API will be used on the first fetch after every rebalance and also on startup. As such, if you need to store offsets in anything other than Kafka, this API should not be used. To avoid re-processing the last message read if a consumer is restarted, the committed offset should be the next message your application should consume, i.e.: `last_offset + 1`.

This is an asynchronous call and will not block. Any errors encountered are either passed to the callback (if provided) or discarded.

**Parameters**

- **offsets** (*dict, optional*) – {TopicPartition: OffsetAndMetadata} dict to commit with the configured `group_id`. Defaults to current consumed offsets for all subscribed partitions.
- **callback** (*callable, optional*) – called as `callback(offsets, response)` with response as either an Exception or a `OffsetCommitResponse` struct. This callback can be used to trigger custom actions when a commit request completes.

**Returns** `kafka.future.Future`

**committed** (*partition*)

Get the last committed offset for the given partition

This offset will be used as the position for the consumer in the event of a failure.

This call may block to do a remote call if the partition in question isn't assigned to this consumer or if the consumer hasn't yet initialized its cache of committed offsets.

**Parameters** **partition** (*TopicPartition*) – the partition to check

**Returns** The last committed offset, or None if there was no prior commit.

**highwater** (*partition*)

Last known highwater offset for a partition

A highwater offset is the offset that will be assigned to the next message that is produced. It may be useful for calculating lag, by comparing with the reported position. Note that both position and highwater refer to the *next* offset – i.e., highwater offset is one greater than the newest available message.

Highwater offsets are returned in `FetchResponse` messages, so will not be available if not `FetchRequests` have been sent for this partition yet.

**Parameters** `partition` (`TopicPartition`) – partition to check

**Returns** offset if available

**Return type** int or None

**partitions\_for\_topic** (`topic`)

Get metadata about the partitions for a given topic.

**Parameters** `topic` (`str`) – topic to check

**Returns** partition ids

**Return type** set

**pause** (`*partitions`)

Suspend fetching from the requested partitions.

Future calls to `poll()` will not return any records from these partitions until they have been resumed using `resume()`. Note that this method does not affect partition subscription. In particular, it does not cause a group rebalance when automatic assignment is used.

**Parameters** `*partitions` (`TopicPartition`) – partitions to pause

**paused** ()

Get the partitions that were previously paused by a call to `pause()`.

**Returns** {partition (`TopicPartition`), ...}

**Return type** set

**poll** (`timeout_ms=0`)

Fetch data from assigned topics / partitions.

Records are fetched and returned in batches by topic-partition. On each poll, consumer will try to use the last consumed offset as the starting offset and fetch sequentially. The last consumed offset can be manually set through `seek(partition, offset)` or automatically set as the last committed offset for the subscribed list of partitions.

Incompatible with iterator interface – use one or the other, not both.

**Parameters** `timeout_ms` (`int`, *optional*) – milliseconds spent waiting in poll if data is not available in the buffer. If 0, returns immediately with any records that are available currently in the buffer, else returns empty. Must not be negative. Default: 0

**Returns** topic to list of records since the last fetch for the subscribed list of topics and partitions

**Return type** dict

**position** (`partition`)

Get the offset of the next record that will be fetched

**Parameters** `partition` (`TopicPartition`) – partition to check

**Returns** offset

**Return type** int

**resume** (\**partitions*)

Resume fetching from the specified (paused) partitions.

**Parameters** \***partitions** (*TopicPartition*) – partitions to resume

**seek** (*partition*, *offset*)

Manually specify the fetch offset for a *TopicPartition*.

Overrides the fetch offsets that the consumer will use on the next *poll()*. If this API is invoked for the same partition more than once, the latest offset will be used on the next *poll()*. Note that you may lose data if this API is arbitrarily used in the middle of consumption, to reset the fetch offsets.

**Parameters**

- **partition** (*TopicPartition*) – partition for seek operation
- **offset** (*int*) – message offset in partition

**Raises** *AssertionError* – if offset is not an *int*  $\geq 0$ ; or if partition is not currently assigned.

**seek\_to\_beginning** (\**partitions*)

Seek to the oldest available offset for partitions.

**Parameters** \***partitions** – optionally provide specific *TopicPartitions*, otherwise default to all assigned partitions

**Raises** *AssertionError* – if any partition is not currently assigned, or if no partitions are assigned

**seek\_to\_end** (\**partitions*)

Seek to the most recent available offset for partitions.

**Parameters** \***partitions** – optionally provide specific *TopicPartitions*, otherwise default to all assigned partitions

**Raises** *AssertionError* – if any partition is not currently assigned, or if no partitions are assigned

**subscribe** (*topics=()*, *pattern=None*, *listener=None*)

Subscribe to a list of topics, or a topic regex pattern

Partitions will be dynamically assigned via a group coordinator. Topic subscriptions are not incremental: this list will replace the current assignment (if there is one).

This method is incompatible with *assign()*

**Parameters**

- **topics** (*list*) – List of topics for subscription.
- **pattern** (*str*) – Pattern to match available topics. You must provide either topics or pattern, but not both.
- **listener** (*ConsumerRebalanceListener*) – Optionally include listener callback, which will be called before and after each rebalance operation.

As part of group management, the consumer will keep track of the list of consumers that belong to a particular group and will trigger a rebalance operation if one of the following events trigger:

- Number of partitions change for any of the subscribed topics
- Topic is created or deleted
- An existing member of the consumer group dies

- A new member is added to the consumer group

When any of these events are triggered, the provided listener will be invoked first to indicate that the consumer’s assignment has been revoked, and then again when the new assignment has been received. Note that this listener will immediately override any listener set in a previous call to subscribe. It is guaranteed, however, that the partitions revoked/assigned through this interface are from topics subscribed in this call.

#### Raises

- `IllegalStateException` – if called after previously calling `assign()`
- `AssertionError` – if neither topics or pattern is provided
- `TypeError` – if listener is not a `ConsumerRebalanceListener`

#### `subscription()`

Get the current topic subscription.

**Returns** {topic, ...}

**Return type** set

#### `topics()`

Get all topics the user is authorized to view.

**Returns** topics

**Return type** set

#### `unsubscribe()`

Unsubscribe from all topics and clear all assigned partitions.

## 5.2.2 KafkaProducer

**class** `kafka.KafkaProducer` (*\*\*configs*)

A Kafka client that publishes records to the Kafka cluster.

The producer is thread safe and sharing a single producer instance across threads will generally be faster than having multiple instances.

The producer consists of a pool of buffer space that holds records that haven’t yet been transmitted to the server as well as a background I/O thread that is responsible for turning these records into requests and transmitting them to the cluster.

The `send()` method is asynchronous. When called it adds the record to a buffer of pending record sends and immediately returns. This allows the producer to batch together individual records for efficiency.

The ‘acks’ config controls the criteria under which requests are considered complete. The “all” setting will result in blocking on the full commit of the record, the slowest but most durable setting.

If the request fails, the producer can automatically retry, unless ‘retries’ is configured to 0. Enabling retries also opens up the possibility of duplicates (see the documentation on message delivery semantics for details: <http://kafka.apache.org/documentation.html#semantics> ).

The producer maintains buffers of unsent records for each partition. These buffers are of a size specified by the ‘batch\_size’ config. Making this larger can result in more batching, but requires more memory (since we will generally have one of these buffers for each active partition).

By default a buffer is available to send immediately even if there is additional unused space in the buffer. However if you want to reduce the number of requests you can set ‘linger\_ms’ to something greater than 0. This will instruct the producer to wait up to that number of milliseconds before sending a request in hope that more records will arrive to fill up the same batch. This is analogous to Nagle’s algorithm in TCP. Note that

records that arrive close together in time will generally batch together even with `linger_ms=0` so under heavy load batching will occur regardless of the `linger` configuration; however setting this to something larger than 0 can lead to fewer, more efficient requests when not under maximal load at the cost of a small amount of latency.

The `buffer_memory` controls the total amount of memory available to the producer for buffering. If records are sent faster than they can be transmitted to the server then this buffer space will be exhausted. When the buffer space is exhausted additional send calls will block.

The `key_serializer` and `value_serializer` instruct how to turn the key and value objects the user provides into bytes.

### Keyword Arguments

- **`bootstrap_servers`** – ‘host[:port]’ string (or list of ‘host[:port]’ strings) that the producer should contact to bootstrap initial cluster metadata. This does not have to be the full node list. It just needs to have at least one broker that will respond to a Metadata API Request. Default port is 9092. If no servers are specified, will default to localhost:9092.
- **`client_id`** (*str*) – a name for this client. This string is passed in each request to servers and can be used to identify specific server-side log entries that correspond to this client. Default: ‘kafka-python-producer-#’ (appended with a unique number per instance)
- **`key_serializer`** (*callable*) – used to convert user-supplied keys to bytes If not None, called as `f(key)`, should return bytes. Default: None.
- **`value_serializer`** (*callable*) – used to convert user-supplied message values to bytes. If not None, called as `f(value)`, should return bytes. Default: None.

- **`acks`** (*0, 1, ‘all’*) – The number of acknowledgments the producer requires the leader to have received before considering a request complete. This controls the durability of records that are sent. The following settings are common:

**0: Producer will not wait for any acknowledgment from the server.** The message will immediately be added to the socket buffer and considered sent. No guarantee can be made that the server has received the record in this case, and the retries configuration will not take effect (as the client won’t generally know of any failures). The offset given back for each record will always be set to -1.

**1: Wait for leader to write the record to its local log only.** Broker will respond without awaiting full acknowledgement from all followers. In this case should the leader fail immediately after acknowledging the record but before the followers have replicated it then the record will be lost.

**all: Wait for the full set of in-sync replicas to write the record.** This guarantees that the record will not be lost as long as at least one in-sync replica remains alive. This is the strongest available guarantee.

If unset, defaults to `acks=1`.

- **`compression_type`** (*str*) – The compression type for all data generated by the producer. Valid values are ‘gzip’, ‘snappy’, ‘lz4’, or None. Compression is of full batches of data, so the efficacy of batching will also impact the compression ratio (more batching means better compression). Default: None.
- **`retries`** (*int*) – Setting a value greater than zero will cause the client to resend any record whose send fails with a potentially transient error. Note that this retry is no different than if the client resent the record upon receiving the error. Allowing retries will potentially change the ordering of records because if two records are sent to a single partition, and the first fails and is retried but the second succeeds, then the second record may appear first. Default: 0.



- **batch\_size** (*int*) – Requests sent to brokers will contain multiple batches, one for each partition with data available to be sent. A small batch size will make batching less common and may reduce throughput (a batch size of zero will disable batching entirely). Default: 16384
- **linger\_ms** (*int*) – The producer groups together any records that arrive in between request transmissions into a single batched request. Normally this occurs only under load when records arrive faster than they can be sent out. However in some circumstances the client may want to reduce the number of requests even under moderate load. This setting accomplishes this by adding a small amount of artificial delay; that is, rather than immediately sending out a record the producer will wait for up to the given delay to allow other records to be sent so that the sends can be batched together. This can be thought of as analogous to Nagle’s algorithm in TCP. This setting gives the upper bound on the delay for batching: once we get batch\_size worth of records for a partition it will be sent immediately regardless of this setting, however if we have fewer than this many bytes accumulated for this partition we will ‘linger’ for the specified time waiting for more records to show up. This setting defaults to 0 (i.e. no delay). Setting linger\_ms=5 would have the effect of reducing the number of requests sent but would add up to 5ms of latency to records sent in the absence of load. Default: 0.
- **partitioner** (*callable*) – Callable used to determine which partition each message is assigned to. Called (after key serialization): partitioner(key\_bytes, all\_partitions, available\_partitions). The default partitioner implementation hashes each non-None key using the same murmur2 algorithm as the java client so that messages with the same key are assigned to the same partition. When a key is None, the message is delivered to a random partition (filtered to partitions with available leaders only, if possible).
- **buffer\_memory** (*int*) – The total bytes of memory the producer should use to buffer records waiting to be sent to the server. If records are sent faster than they can be delivered to the server the producer will block up to max\_block\_ms, raising an exception on timeout. In the current implementation, this setting is an approximation. Default: 33554432 (32MB)
- **max\_block\_ms** (*int*) – Number of milliseconds to block during send() and partitions\_for(). These methods can be blocked either because the buffer is full or metadata unavailable. Blocking in the user-supplied serializers or partitioner will not be counted against this timeout. Default: 60000.
- **max\_request\_size** (*int*) – The maximum size of a request. This is also effectively a cap on the maximum record size. Note that the server has its own cap on record size which may be different from this. This setting will limit the number of record batches the producer will send in a single request to avoid sending huge requests. Default: 1048576.
- **metadata\_max\_age\_ms** (*int*) – The period of time in milliseconds after which we force a refresh of metadata even if we haven’t seen any partition leadership changes to proactively discover any new brokers or partitions. Default: 300000
- **retry\_backoff\_ms** (*int*) – Milliseconds to backoff when retrying on errors. Default: 100.
- **request\_timeout\_ms** (*int*) – Client request timeout in milliseconds. Default: 30000.
- **receive\_buffer\_bytes** (*int*) – The size of the TCP receive buffer (SO\_RCVBUF) to use when reading data. Default: None (relies on system defaults). Java client defaults to 32768.
- **send\_buffer\_bytes** (*int*) – The size of the TCP send buffer (SO\_SNDBUF) to use when sending data. Default: None (relies on system defaults). Java client defaults to 131072.
- **reconnect\_backoff\_ms** (*int*) – The amount of time in milliseconds to wait before attempting to reconnect to a given host. Default: 50.
- **max\_in\_flight\_requests\_per\_connection** (*int*) – Requests are pipelined to kafka brokers up to this number of maximum requests per broker connection. Default: 5.

- **security\_protocol** (*str*) – Protocol used to communicate with brokers. Valid values are: PLAINTEXT, SSL. Default: PLAINTEXT.
- **ssl\_context** (*ssl.SSLContext*) – pre-configured SSLContext for wrapping socket connections. If provided, all other ssl\_\* configurations will be ignored. Default: None.
- **ssl\_check\_hostname** (*bool*) – flag to configure whether ssl handshake should verify that the certificate matches the brokers hostname. default: true.
- **ssl\_cafile** (*str*) – optional filename of ca file to use in certificate verification. default: none.
- **ssl\_certfile** (*str*) – optional filename of file in pem format containing the client certificate, as well as any ca certificates needed to establish the certificate’s authenticity. default: none.
- **ssl\_keyfile** (*str*) – optional filename containing the client private key. default: none.
- **ssl\_crlfile** (*str*) – optional filename containing the CRL to check for certificate expiration. By default, no CRL check is done. When providing a file, only the leaf certificate will be checked against this CRL. The CRL can only be checked with Python 3.4+ or 2.7.9+. default: none.
- **api\_version** (*str*) – specify which kafka API version to use. If set to ‘auto’, will attempt to infer the broker version by probing various APIs. Default: auto
- **api\_version\_auto\_timeout\_ms** (*int*) – number of milliseconds to throw a timeout exception from the constructor when checking the broker api version. Only applies if api\_version set to ‘auto’

---

**Note:** Configuration parameters are described in more detail at <https://kafka.apache.org/090/configuration.html#producerconfigs>

---

**close** (*timeout=None*)  
Close this producer.

**Parameters** *timeout* (*float*, *optional*) – timeout in seconds to wait for completion.

**flush** (*timeout=None*)

Invoking this method makes all buffered records immediately available to send (even if *linger\_ms* is greater than 0) and blocks on the completion of the requests associated with these records. The post-condition of `flush()` is that any previously sent record will have completed (e.g. `Future.is_done() == True`). A request is considered completed when either it is successfully acknowledged according to the ‘acks’ configuration for the producer, or it results in an error.

Other threads can continue sending messages while one thread is blocked waiting for a flush call to complete; however, no guarantee is made about the completion of messages sent after the flush call begins.

**Parameters** *timeout* (*float*, *optional*) – timeout in seconds to wait for completion.

**partitions\_for** (*topic*)

Returns set of all known partitions for the topic.

**send** (*topic*, *value=None*, *key=None*, *partition=None*, *timestamp\_ms=None*)

Publish a message to a topic.

**Parameters**

- **topic** (*str*) – topic where the message will be published
- **value** (*optional*) – message value. Must be type bytes, or be serializable to bytes via configured *value\_serializer*. If value is None, key is required and message acts as a ‘delete’. See kafka compaction documentation for more details:

<http://kafka.apache.org/documentation.html#compaction> (compaction requires kafka >= 0.8.1)

- **partition** (*int*, *optional*) – optionally specify a partition. If not set, the partition will be selected using the configured ‘partitioner’.
- **key** (*optional*) – a key to associate with the message. Can be used to determine which partition to send the message to. If partition is None (and producer’s partitioner config is left as default), then messages with the same key will be delivered to the same partition (but if key is None, partition is chosen randomly). Must be type bytes, or be serializable to bytes via configured key\_serializer.
- **timestamp\_ms** (*int*, *optional*) – epoch milliseconds (from Jan 1 1970 UTC) to use as the message timestamp. Defaults to current time.

**Returns** resolves to RecordMetadata

**Return type** FutureRecordMetadata

**Raises** KafkaTimeoutError – if unable to fetch topic metadata, or unable to obtain memory buffer prior to configured max\_block\_ms

### 5.2.3 KafkaClient

**class** kafka.client.**KafkaClient** (\*\**configs*)

A network client for asynchronous request/response network i/o. This is an internal class used to implement the user-facing producer and consumer clients.

This class is not thread-safe!

**add\_topic** (*topic*)

Add a topic to the list of topics tracked via metadata.

**Parameters** **topic** (*str*) – topic to track

**Returns** resolves after metadata request/response

**Return type** Future

**check\_version** (*node\_id=None*, *timeout=2*, *strict=False*)

Attempt to guess a broker version

**Note: it is possible that this method blocks longer than the** specified timeout. This can happen if the entire cluster is down and the client enters a bootstrap backoff sleep. This is only possible if node\_id is None.

Returns: version str, i.e. ‘0.10’, ‘0.9’, ‘0.8.2’, ‘0.8.1’, ‘0.8.0’

**Raises**

- NodeNotReadyError (if node\_id is provided)
- NoBrokersAvailable (if node\_id is None)
- UnrecognizedBrokerVersion – please file bug if seen!
- AssertionError (if strict=True) – please file bug if seen!

**close** (*node\_id=None*)

Closes one or all broker connections.

**Parameters** **node\_id** (*int*, *optional*) – the id of the node to close

**connected** (*node\_id*)

Return True iff the *node\_id* is connected.

**connection\_delay** (*node\_id*)

Returns the number of milliseconds to wait, based on the connection state, before attempting to send data. When disconnected, this respects the reconnect backoff time. When connecting, returns 0 to allow non-blocking connect to finish. When connected, returns a very large number to handle slow/stalled connections.

**Parameters** *node\_id* (*int*) – The id of the node to check

**Returns** The number of milliseconds to wait.

**Return type** *int*

**in\_flight\_request\_count** (*node\_id=None*)

Get the number of in-flight requests for a node or all nodes.

**Parameters** *node\_id* (*int*, *optional*) – a specific node to check. If unspecified, return the total for all nodes

**Returns** pending in-flight requests for the node, or all nodes if None

**Return type** *int*

**is\_disconnected** (*node\_id*)

Check whether the node connection has been disconnected failed.

A disconnected node has either been closed or has failed. Connection failures are usually transient and can be resumed in the next `ready()` call, but there are cases where transient failures need to be caught and re-acted upon.

**Parameters** *node\_id* (*int*) – the id of the node to check

**Returns** True iff the node exists and is disconnected

**Return type** *bool*

**is\_ready** (*node\_id*)

Check whether a node is ready to send more requests.

In addition to connection-level checks, this method also is used to block additional requests from being sent during a metadata refresh.

**Parameters** *node\_id* (*int*) – id of the node to check

**Returns** True if the node is ready and metadata is not refreshing

**Return type** *bool*

**least\_loaded\_node** ()

Choose the node with fewest outstanding requests, with fallbacks.

This method will prefer a node with an existing connection, but will potentially choose a node for which we don't yet have a connection if all existing connections are in use. This method will never choose a node that was disconnected within the reconnect backoff period. If all else fails, the method will attempt to bootstrap again using the `bootstrap_servers` list.

**Returns** *node\_id* or None if no suitable node was found

**poll** (*timeout\_ms=None*, *future=None*, *sleep=True*)

Try to read and write to sockets.

This method will also attempt to complete node connections, refresh stale metadata, and run previously-scheduled tasks.

**Parameters**

- **timeout\_ms** (*int, optional*) – maximum amount of time to wait (in ms) for at least one response. Must be non-negative. The actual timeout will be the minimum of timeout, request timeout and metadata timeout. Default: request\_timeout\_ms
- **future** (*Future, optional*) – if provided, blocks until future.is\_done
- **sleep** (*bool*) – if True and there is nothing to do (no connections or requests in flight), will sleep for duration timeout before returning empty results. Default: False.

**Returns** responses received (can be empty)

**Return type** list

**ready** (*node\_id*)

Check whether a node is connected and ok to send more requests.

**Parameters** **node\_id** (*int*) – the id of the node to check

**Returns** True if we are ready to send to the given node

**Return type** bool

**schedule** (*task, at*)

Schedule a new task to be executed at the given time.

This is “best-effort” scheduling and should only be used for coarse synchronization. A task cannot be scheduled for multiple times simultaneously; any previously scheduled instance of the same task will be cancelled.

**Parameters**

- **task** (*callable*) – task to be scheduled
- **at** (*float or int*) – epoch seconds when task should run

**Returns** resolves to result of task call, or exception if raised

**Return type** Future

**send** (*node\_id, request*)

Send a request to a specific node.

**Parameters**

- **node\_id** (*int*) – destination node
- **request** (*Struct*) – request object (not-encoded)

**Raises** `AssertionError` – if node\_id is not in current cluster metadata

**Returns** resolves to Response struct or Error

**Return type** Future

**set\_topics** (*topics*)

Set specific topics to track for metadata.

**Parameters** **topics** (*list of str*) – topics to check for metadata

**Returns** resolves after metadata request/response

**Return type** Future

**unsubscribe** (*task*)

Unsubscribe a task.

This will remove all instances of the task from the task queue. This is a no-op if the task is not scheduled.

**Parameters** **task** (*callable*) – task to be unscheduled

## 5.2.4 BrokerConnection

`class kafka.BrokerConnection (host, port, afi, **configs)`

**blacked\_out** ()

Return true if we are disconnected from the given node and can't re-establish a connection yet

**can\_send\_more** ()

Return True unless there are max\_in\_flight\_requests.

**check\_version** (timeout=2, strict=False)

Attempt to guess the broker version. This is a blocking call.

**close** (error=None)

Close socket and fail all in-flight-requests.

**Parameters** **error** (*Exception, optional*) – pending in-flight-requests will be failed with this exception. Default: kafka.errors.ConnectionError.

**connect** ()

Attempt to connect and return ConnectionState

**connected** ()

Return True iff socket is connected.

**connecting** ()

Returns True if still connecting (this may encompass several different states, such as SSL handshake, authorization, etc).

**disconnected** ()

Return True iff socket is closed

**recv** ()

Non-blocking network receive.

Return response if available

**send** (request, expect\_response=True)

send request, return Future()

Can block on network if request is larger than send\_buffer\_bytes

## 5.3 Simple APIs (DEPRECATED)

### 5.3.1 SimpleConsumer (DEPRECATED)

```
from kafka import SimpleProducer, SimpleClient

# To consume messages
client = SimpleClient('localhost:9092')
consumer = SimpleConsumer(client, "my-group", "my-topic")
for message in consumer:
    # message is raw byte string -- decode if necessary!
```

```

# e.g., for unicode: `message.decode('utf-8')`
print(message)

# Use multiprocessing for parallel consumers
from kafka import MultiProcessConsumer

# This will split the number of partitions among two processes
consumer = MultiProcessConsumer(client, "my-group", "my-topic", num_procs=2)

# This will spawn processes such that each handles 2 partitions max
consumer = MultiProcessConsumer(client, "my-group", "my-topic",
                                partitions_per_proc=2)

for message in consumer:
    print(message)

for message in consumer.get_messages(count=5, block=True, timeout=4):
    print(message)

client.close()

```

## 5.3.2 SimpleProducer (DEPRECATED)

### Asynchronous Mode

```

from kafka import SimpleProducer, SimpleClient

# To send messages asynchronously
client = SimpleClient('localhost:9092')
producer = SimpleProducer(client, async=True)
producer.send_messages('my-topic', b'async message')

# To send messages in batch. You can use any of the available
# producers for doing this. The following producer will collect
# messages in batch and send them to Kafka after 20 messages are
# collected or every 60 seconds
# Notes:
# * If the producer dies before the messages are sent, there will be losses
# * Call producer.stop() to send the messages and cleanup
producer = SimpleProducer(client,
                          async=True,
                          batch_send_every_n=20,
                          batch_send_every_t=60)

```

### Synchronous Mode

```

from kafka import SimpleProducer, SimpleClient

# To send messages synchronously
client = SimpleClient('localhost:9092')
producer = SimpleProducer(client, async=False)

# Note that the application is responsible for encoding messages to type bytes
producer.send_messages('my-topic', b'some message')

```

```
producer.send_messages('my-topic', b'this method', b'is variadic')

# Send unicode message
producer.send_messages('my-topic', u'.'.encode('utf-8'))

# To wait for acknowledgements
# ACK_AFTER_LOCAL_WRITE : server will wait till the data is written to
#                          a local log before sending response
# ACK_AFTER_CLUSTER_COMMIT : server will block until the message is committed
#                          by all in sync replicas before sending a response
producer = SimpleProducer(client,
                           async=False,
                           req_acks=SimpleProducer.ACK_AFTER_LOCAL_WRITE,
                           ack_timeout=2000,
                           sync_fail_on_error=False)

responses = producer.send_messages('my-topic', b'another message')
for r in responses:
    logging.info(r.offset)
```

### 5.3.3 KeyedProducer (DEPRECATED)

```
from kafka import (
    SimpleClient, KeyedProducer,
    Murmur2Partitioner, RoundRobinPartitioner)

kafka = SimpleClient('localhost:9092')

# HashedPartitioner is default (currently uses python hash())
producer = KeyedProducer(kafka)
producer.send_messages(b'my-topic', b'key1', b'some message')
producer.send_messages(b'my-topic', b'key2', b'this methode')

# Murmur2Partitioner attempts to mirror the java client hashing
producer = KeyedProducer(kafka, partitioner=Murmur2Partitioner)

# Or just produce round-robin (or just use SimpleProducer)
producer = KeyedProducer(kafka, partitioner=RoundRobinPartitioner)
```

### 5.3.4 SimpleClient (DEPRECATED)

```
import time
from kafka import SimpleClient
from kafka.errors import LeaderNotAvailableError, NotLeaderForPartitionError
from kafka.protocol import create_message
from kafka.structs import ProduceRequestPayload

kafka = SimpleClient('localhost:9092')
payload = ProduceRequestPayload(topic='my-topic', partition=0,
                                messages=[create_message("some message")])

retries = 5
resps = []
while retries and not resps:
    retries -= 1
```



```

try:
    resps = kafka.send_produce_request(
        payloads=[payload], fail_on_error=True)
except LeaderNotAvailableError, NotLeaderForPartitionError:
    kafka.load_metadata_for_topics()
    time.sleep(1)

# Other exceptions you might consider handling:
# UnknownTopicOrPartitionError, TopicAuthorizationFailedError,
# RequestTimedOutError, MessageSizeTooLargeError, InvalidTopicError,
# RecordListTooLargeError, InvalidRequiredAcksError,
# NotEnoughReplicasError, NotEnoughReplicasAfterAppendError

kafka.close()

resps[0].topic      # 'my-topic'
resps[0].partition  # 0
resps[0].error      # 0
resps[0].offset     # offset of the first message sent in this request

```

## 5.4 Install

Install with your favorite package manager

### 5.4.1 Latest Release

Pip:

```
pip install kafka-python
```

Releases are also listed at <https://github.com/dpkp/kafka-python/releases>

### 5.4.2 Bleeding-Edge

```
git clone https://github.com/dpkp/kafka-python
pip install ./kafka-python
```

Setuptools:

```
git clone https://github.com/dpkp/kafka-python
easy_install ./kafka-python
```

Using *setup.py* directly:

```
git clone https://github.com/dpkp/kafka-python
cd kafka-python
python setup.py install
```

### 5.4.3 Optional LZ4 install

To enable LZ4 compression/decompression, install lz4tools and xxhash:

```
>>> pip install lz4tools
>>> pip install xxhash
```

*Note:* these modules do not support python2.6

## 5.4.4 Optional Snappy install

### Install Development Libraries

Download and build Snappy from <http://code.google.com/p/snappy/downloads/list>

Ubuntu:

```
apt-get install libsnappy-dev
```

OSX:

```
brew install snappy
```

From Source:

```
wget http://snappy.googlecode.com/files/snappy-1.0.5.tar.gz
tar xzvf snappy-1.0.5.tar.gz
cd snappy-1.0.5
./configure
make
sudo make install
```

### Install Python Module

Install the *python-snappy* module

```
pip install python-snappy
```

## 5.5 Tests

Test environments are managed via tox. The test suite is run via pytest. Individual tests are written using unittest, pytest, and in some cases, doctest.

Linting is run via pylint, but is generally skipped on python2.6 and pypy due to pylint compatibility / performance issues.

For test coverage details, see <https://coveralls.io/github/dpkp/kafka-python>

The test suite includes unit tests that mock network interfaces, as well as integration tests that setup and teardown kafka broker (and zookeeper) fixtures for client / consumer / producer testing.

### 5.5.1 Unit tests

To run the tests locally, install tox – *pip install tox* See <http://tox.readthedocs.org/en/latest/install.html>

Then simply run tox, optionally setting the python environment. If unset, tox will loop through all environments.

```
tox -e py27
tox -e py35

# run protocol tests only
tox -- -v test.test_protocol

# re-run the last failing test, dropping into pdb
tox -e py27 -- --lf --pdb

# see available (pytest) options
tox -e py27 -- --help
```

## 5.5.2 Integration tests

```
KAFKA_VERSION=0.10.0.0 tox -e py27
KAFKA_VERSION=0.8.2.2 tox -e py35
```

Integration tests start Kafka and Zookeeper fixtures. This requires downloading kafka server binaries:

```
./build_integration.sh
```

By default, this will install 0.8.1.1, 0.8.2.2, 0.9.0.1, and 0.10.0.0 brokers into the `servers/` directory. To install a specific version, set `KAFKA_VERSION=1.2.3`:

```
KAFKA_VERSION=0.8.0 ./build_integration.sh
```

Then run the tests against supported Kafka versions, simply set the `KAFKA_VERSION` env variable to the server build you want to use for testing:

```
KAFKA_VERSION=0.9.0.1 tox -e py27
```

To test against the kafka source tree, set `KAFKA_VERSION=trunk` [optionally set `SCALA_VERSION` (defaults to 2.10)]

```
SCALA_VERSION=2.11 KAFKA_VERSION=trunk ./build_integration.sh
KAFKA_VERSION=trunk tox -e py35
```

## 5.6 Compatibility

kafka-python is compatible with (and tested against) broker versions 0.10 through 0.8.0 . kafka-python is not compatible with the 0.8.2-beta release.

kafka-python is tested on python 2.6, 2.7, 3.3, 3.4, 3.5, and pypy.

Builds and tests via Travis-CI. See <https://travis-ci.org/dpkp/kafka-python>

## 5.7 Support

For support, see github issues at <https://github.com/dpkp/kafka-python>

Limited IRC chat at #kafka-python on freenode (general chat is #apache-kafka).

For information about Apache Kafka generally, see <https://kafka.apache.org/>

For general discussion of kafka-client design and implementation (not python specific), see <https://groups.google.com/forum/m/#!forum/kafka-clients>

## 5.8 License

Apache License, v2.0. See [LICENSE](#).

Copyright 2016, Dana Powers, David Arthur, and Contributors (See [AUTHORS](#)).

## 5.9 Changelog

### 5.9.1 1.2.5 (July 15, 2016)

#### Bugfixes

- Fix bug causing KafkaProducer to double-compress message batches on retry
- Check for double-compressed messages in KafkaConsumer, log warning and optionally skip
- Drop recursion in `_unpack_message_set`; only decompress once

### 5.9.2 1.2.4 (July 8, 2016)

#### Bugfixes

- Update `consumer_timeout_ms` docstring - KafkaConsumer raises `StopIteration`, no longer `ConsumerTimeout`
- Use explicit subscription state flag to handle `seek()` during message iteration
- Fix consumer iteration on compacted topics (dppk PR 752)
- Support `ssl_password` config when loading cert chains (amckemie PR 750)

### 5.9.3 1.2.3 (July 2, 2016)

#### Patch Improvements

- Fix gc error log: avoid `AttributeError` in `_unregister_cleanup` (dppk PR 747)
- Wakeup socket optimizations (dppk PR 740)
- Assert will be disabled by “python -O” (tyronecai PR 736)
- Randomize order of topics/partitions processed by fetcher to improve balance (dppk PR 732)
- Allow `client.check_version` timeout to be set in Producer and Consumer constructors (eastlondoner PR 647)

### 5.9.4 1.2.2 (June 21, 2016)

#### Bugfixes

- Clarify timeout unit in KafkaProducer close and flush (ms7s PR 734)
- Avoid busy poll during metadata refresh failure with retry\_backoff\_ms (dpgk PR 733)
- Check\_version should scan nodes until version found or timeout (dpgk PR 731)
- Fix bug which could cause least\_loaded\_node to always return the same unavailable node (dpgk PR 730)
- Fix producer garbage collection with weakref in atexit handler (dpgk PR 728)
- Close client selector to fix fd leak (msmith PR 729)
- Tweak spelling mistake in error const (steve8918 PR 719)
- Rearrange connection tests to separate legacy KafkaConnection

### 5.9.5 1.2.1 (June 1, 2016)

#### Bugfixes

- Fix regression in MessageSet decoding wrt PartialMessages (#716)
- Catch response decode errors and log details (#715)
- Fix Legacy support url (#712 - JonasGroeger)
- Update sphinx docs re 0.10 broker support

### 5.9.6 1.2.0 (May 24, 2016)

#### Support Kafka 0.10 Features

- Add protocol support for ApiVersionRequest (dpgk PR 678)
- KAFKA-3025: Message v1 – add timestamp and relative offsets (dpgk PR 693)
- Use Fetch/Produce API v2 for brokers  $\geq 0.10$  (uses message format v1) (dpgk PR 694)
- Use standard LZ4 framing for v1 messages / kafka 0.10 (dpgk PR 695)

#### Consumers

- Update SimpleConsumer / legacy protocol to handle compressed messages (paulcavallaro PR 684)

#### Producers

- KAFKA-3388: Fix expiration of batches sitting in the accumulator (dpgk PR 699)
- KAFKA-3197: when max.in.flight.request.per.connection = 1, attempt to guarantee ordering (dpgk PR 698)
- Dont use soon-to-be-reserved keyword await as function name (FutureProduceResult) (dpgk PR 697)

## Clients

- Fix socket leaks in KafkaClient (dpgk PR 696)

## Documentation

<none>

## Internals

- Support SSL CRL [requires python 2.7.9+ / 3.4+] (vincentbernat PR 683)
- Use original hostname for SSL checks (vincentbernat PR 682)
- Always pass encoded message bytes to MessageSet.encode()
- Raise ValueError on protocol encode/decode errors
- Supplement socket.gaierror exception in BrokerConnection.connect() (erikbeebe PR 687)
- BrokerConnection check\_version: expect 0.9 to fail with CorrelationIdError
- Fix small bug in Sensor (zackdever PR 679)

### 5.9.7 1.1.1 (Apr 26, 2016)

#### Bugfixes

- Fix throttle\_time\_ms sensor handling (zackdever PR 667)
- Improve handling of disconnected sockets (EasyPost PR 666 / dpgk)
- Disable standard metadata refresh triggers during bootstrap (dpgk)
- More predictable Future callback/errback exceptions (zackdever PR 670)
- Avoid some exceptions in Coordinator.\_\_del\_\_ (dpgk PR 668)

### 5.9.8 1.1.0 (Apr 25, 2016)

#### Consumers

- Avoid resending FetchRequests that are pending on internal queue
- Log debug messages when skipping fetched messages due to offset checks
- KAFKA-3013: Include topic-partition in exception for expired batches
- KAFKA-3318: clean up consumer logging and error messages
- Improve unknown coordinator error handling
- Improve auto-commit error handling when group\_id is None
- Add paused() API (zackdever PR 602)
- Add default\_offset\_commit\_callback to KafkaConsumer DEFAULT\_CONFIGS

## Producers

<none>

## Clients

- Support SSL connections
- Use selectors module for non-blocking IO
- Refactor KafkaClient connection management
- Fix AttributeError in \_\_del\_\_
- SimpleClient: catch errors thrown by \_get\_leader\_for\_partition (zackdever PR 606)

## Documentation

- Fix serializer/deserializer examples in README
- Update max.block.ms docstring
- Remove errant next(consumer) from consumer documentation
- Add producer.flush() to usage docs

## Internals

- Add initial metrics implementation (zackdever PR 637)
- KAFKA-2136: support Fetch and Produce v1 (throttle\_time\_ms)
- Use version-indexed lists for request/response protocol structs (dpkp PR 630)
- Split kafka.common into kafka.structs and kafka.errors
- Handle partial socket send() (dpkp PR 611)
- Fix windows support (dpkp PR 603)
- IPv6 support (TimEvens PR 615; Roguelazer PR 642)

### 5.9.9 1.0.2 (Mar 14, 2016)

## Consumers

- Improve KafkaConsumer Heartbeat handling (dpkp PR 583)
- Fix KafkaConsumer.position bug (stefanth PR 578)
- Raise TypeError when partition is not a TopicPartition (dpkp PR 587)
- KafkaConsumer.poll should sleep to prevent tight-loops (dpkp PR 597)

## Producers

- Fix producer threading bug that can crash sender (dppk PR 590)
- Fix bug in producer buffer pool reallocation (dppk PR 585)
- Remove spurious warnings when closing sync SimpleProducer (twm PR 567)
- Fix FutureProduceResult.await() on python2.6 (dppk)
- Add optional timeout parameter to KafkaProducer.flush() (dppk)
- KafkaProducer optimizations (zackdever PR 598)

## Clients

- Improve error handling in SimpleClient.load\_metadata\_for\_topics (dppk)
- Improve handling of KafkaClient.least\_loaded\_node failure (dppk PR 588)

## Documentation

- Fix KafkaError import error in docs (shichao-an PR 564)
- Fix serializer / deserializer examples (scribu PR 573)

## Internals

- Update to Kafka 0.9.0.1 for integration testing
- Fix ifr.future.failure in conn.py (mortenlj PR 566)
- Improve Zookeeper / Kafka Fixture management (dppk)

## 5.9.10 1.0.1 (Feb 19, 2016)

### Consumers

- Add RangePartitionAssignor (and use as default); add assignor tests (dppk PR 550)
- Make sure all consumers are in same generation before stopping group test
- Verify node ready before sending offset fetch request from coordinator
- Improve warning when offset fetch request returns unknown topic / partition

### Producers

- Warn if pending batches failed during flush
- Fix concurrency bug in RecordAccumulator.ready()
- Fix bug in SimpleBufferPool memory condition waiting / timeout
- Support batch\_size = 0 in producer buffers (dppk PR 558)
- Catch duplicate batch.done() calls [e.g., maybe\_expire then a response errback]



## Clients

### Documentation

- Improve kafka.cluster docstrings
- Migrate load\_example.py to KafkaProducer / KafkaConsumer

### Internals

- Dont override system rcvbuf or sndbuf unless configured explicitly (dppk PR 557)
- Some attributes may not exist in \_\_del\_\_ if we failed assertions
- Break up some circular references and close client wake pipes on \_\_del\_\_ (aisch PR 554)

## 5.9.11 1.0.0 (Feb 15, 2016)

This release includes significant code changes. Users of older kafka-python versions are encouraged to test upgrades before deploying to production as some interfaces and configuration options have changed.

Users of SimpleConsumer / SimpleProducer / SimpleClient (formerly KafkaClient) from prior releases should migrate to KafkaConsumer / KafkaProducer. Low-level APIs (Simple\*) are no longer being actively maintained and will be removed in a future release.

For comprehensive API documentation, please see `python help()` / docstrings, [kafka-python.readthedocs.org](http://kafka-python.readthedocs.org), or run `'tox -e docs'` from source to build documentation locally.

### Consumers

- KafkaConsumer re-written to emulate the new 0.9 kafka consumer (java client) and support coordinated consumer groups (feature requires `>= 0.9.0.0` brokers)
  - Methods no longer available:
    - \* `configure` [initialize a new consumer instead]
    - \* `set_topic_partitions` [use `subscribe()` or `assign()`]
    - \* `fetch_messages` [use `poll()` or iterator interface]
    - \* `get_partition_offsets`
    - \* `offsets` [use `committed(partition)`]
    - \* `task_done` [handled internally by auto-commit; or commit offsets manually]
  - Configuration changes (consistent with updated java client):
    - \* lots of new configuration parameters – see docs for details
    - \* `auto_offset_reset`: previously values were 'smallest' or 'largest', now values are 'earliest' or 'latest'
    - \* `fetch_wait_max_ms` is now `fetch_max_wait_ms`
    - \* `max_partition_fetch_bytes` is now `max_partition_fetch_bytes`
    - \* `deserializer_class` is now `value_deserializer` and `key_deserializer`
    - \* `auto_commit_enable` is now `enable_auto_commit`
    - \* `auto_commit_interval_messages` was removed

- \* `socket_timeout_ms` was removed
- \* `refresh_leader_backoff_ms` was removed
- `SimpleConsumer` and `MultiProcessConsumer` are now deprecated and will be removed in a future release. Users are encouraged to migrate to `KafkaConsumer`.

### Producers

- new producer class: `KafkaProducer`. Exposes the same interface as official java client. Async by default; `returned future.get()` can be called for synchronous blocking
- `SimpleProducer` is now deprecated and will be removed in a future release. Users are encouraged to migrate to `KafkaProducer`.

### Clients

- synchronous `KafkaClient` renamed to `SimpleClient`. For backwards compatibility, you will get a `SimpleClient` via `'from kafka import KafkaClient'`. This will change in a future release.
- All client calls use non-blocking IO under the hood.
- Add probe method `check_version()` to infer broker versions.

### Documentation

- Updated README and sphinx documentation to address new classes.
- Docstring improvements to make python `help()` easier to use.

### Internals

- Old protocol stack is deprecated. It has been moved to `kafka.protocol.legacy` and may be removed in a future release.
- Protocol layer re-written using Type classes, Schemas and Structs (modeled on the java client).
- Add support for LZ4 compression (including broken framing header checksum).

## 5.9.12 0.9.5 (Dec 6, 2015)

### Consumers

- Initial support for consumer coordinator: offsets only (toddpalino PR 420)
- Allow blocking until some messages are received in `SimpleConsumer` (saaros PR 457)
- Support subclass config changes in `KafkaConsumer` (zackdever PR 446)
- Support retry semantics in `MultiProcessConsumer` (barricadeio PR 456)
- Support `partition_info` in `MultiProcessConsumer` (scrapinghub PR 418)
- Enable `seek()` to an absolute offset in `SimpleConsumer` (haosdent PR 412)
- Add `KafkaConsumer.close()` (ucarion PR 426)

## Producers

- Catch client.reinit() exceptions in async producer (dpgk)
- Producer.stop() now blocks until async thread completes (dpgk PR 485)
- Catch errors during load\_metadata\_for\_topics in async producer (bschopman PR 467)
- Add compression-level support for codecs that support it (trbs PR 454)
- Fix translation of Java murmur2 code, fix byte encoding for Python 3 (chrischamberlin PR 439)
- Only call stop() on not-stopped producer objects (docker-hub PR 435)
- Allow null payload for deletion feature (scrapinghub PR 409)

## Clients

- Use non-blocking io for broker aware requests (ecanzonieri PR 473)
- Use debug logging level for metadata request (ecanzonieri PR 415)
- Catch KafkaUnavailableError in \_send\_broker\_aware\_request (mutability PR 436)
- Lower logging level on replica not available and commit (ecanzonieri PR 415)

## Documentation

- Update docs and links wrt maintainer change (mumrah -> dpgk)

## Internals

- Add py35 to tox testing
- Update travis config to use container infrastructure
- Add 0.8.2.2 and 0.9.0.0 resources for integration tests; update default official releases
- new pylint disables for pylint 1.5.1 (zackdever PR 481)
- Fix python3 / python2 comments re queue/Queue (dpgk)
- Add Murmur2Partitioner to kafka \_\_all\_\_ imports (dpgk Issue 471)
- Include LICENSE in PyPI sdist (koobs PR 441)

### 5.9.13 0.9.4 (June 11, 2015)

## Consumers

- Refactor SimpleConsumer internal fetch handling (dpgk PR 399)
- Handle exceptions in SimpleConsumer commit() and reset\_partition\_offset() (dpgk PR 404)
- Improve FailedPayloadsError handling in KafkaConsumer (dpgk PR 398)
- KafkaConsumer: avoid raising KeyError in task\_done (dpgk PR 389)
- MultiProcessConsumer – support configured partitions list (dpgk PR 380)
- Fix SimpleConsumer leadership change handling (dpgk PR 393)

- Fix SimpleConsumer connection error handling (reAsOn2010 PR 392)
- Improve Consumer handling of ‘falsy’ partition values (wting PR 342)
- Fix \_offsets call error in KafkaConsumer (hellais PR 376)
- Fix str/bytes bug in KafkaConsumer (dppk PR 365)
- Register atexit handlers for consumer and producer thread/multiprocess cleanup (dppk PR 360)
- Always fetch commit offsets in base consumer unless group is None (dppk PR 356)
- Stop consumer threads on delete (dppk PR 357)
- Deprecate metadata\_broker\_list in favor of bootstrap\_servers in KafkaConsumer (dppk PR 340)
- Support pass-through parameters in multiprocess consumer (scrapinghub PR 336)
- Enable offset commit on SimpleConsumer.seek (ecanzonieri PR 350)
- Improve multiprocess consumer partition distribution (scrapinghub PR 335)
- Ignore messages with offset less than requested (wkiser PR 328)
- Handle OffsetOutOfRange in SimpleConsumer (ecanzonieri PR 296)

## Producers

- Add Murmur2Partitioner (dppk PR 378)
- Log error types in SimpleProducer and SimpleConsumer (dppk PR 405)
- SimpleProducer support configuration of fail\_on\_error (dppk PR 396)
- Deprecate KeyedProducer.send() (dppk PR 379)
- Further improvements to async producer code (dppk PR 388)
- Add more configuration parameters for async producer (dppk)
- Deprecate SimpleProducer batch\_send=True in favor of async (dppk)
- Improve async producer error handling and retry logic (vshlapakov PR 331)
- Support message keys in async producer (vshlapakov PR 329)
- Use threading instead of multiprocessing for Async Producer (vshlapakov PR 330)
- Stop threads on \_\_del\_\_ (chmdquesne PR 324)
- Fix leadership failover handling in KeyedProducer (dppk PR 314)

## KafkaClient

- Add .topics property for list of known topics (dppk)
- Fix request / response order guarantee bug in KafkaClient (dppk PR 403)
- Improve KafkaClient handling of connection failures in \_get\_conn (dppk)
- Client clears local metadata cache before updating from server (dppk PR 367)
- KafkaClient should return a response or error for each request - enable better retry handling (dppk PR 366)
- Improve str/bytes conversion in KafkaClient and KafkaConsumer (dppk PR 332)
- Always return sorted partition ids in client.get\_partition\_ids\_for\_topic() (dppk PR 315)

## Documentation

- Cleanup Usage Documentation
- Improve KafkaConsumer documentation (dppk PR 341)
- Update consumer documentation (sontek PR 317)
- Add doc configuration for tox (sontek PR 316)
- Switch to .rst doc format (sontek PR 321)
- Fixup google groups link in README (sontek PR 320)
- Automate documentation at kafka-python.readthedocs.org

## Internals

- Switch integration testing from 0.8.2.0 to 0.8.2.1 (dppk PR 402)
- Fix most flaky tests, improve debug logging, improve fixture handling (dppk)
- General style cleanups (dppk PR 394)
- Raise error on duplicate topic-partition payloads in protocol grouping (dppk)
- Use module-level loggers instead of simply 'kafka' (dppk)
- Remove pkg\_resources check for \_\_version\_\_ at runtime (dppk PR 387)
- Make external API consistently support python3 strings for topic (kecaps PR 361)
- Fix correlation id overflow (dppk PR 355)
- Cleanup kafka/common structs (dppk PR 338)
- Use context managers in gzip\_encode / gzip\_decode (dppk PR 337)
- Save failed request as FailedPayloadsError attribute (jobevers PR 302)
- Remove unused kafka.queue (mumrah)

### 5.9.14 0.9.3 (Feb 3, 2015)

- Add coveralls.io support (sontek PR 307)
- Fix python2.6 threading.Event bug in ReentrantTimer (dppk PR 312)
- Add kafka 0.8.2.0 to travis integration tests (dppk PR 310)
- Auto-convert topics to utf-8 bytes in Producer (sontek PR 306)
- Fix reference cycle between SimpleConsumer and ReentrantTimer (zhaopengzp PR 309)
- Add Sphinx API docs (wedaly PR 282)
- Handle additional error cases exposed by 0.8.2.0 kafka server (dppk PR 295)
- Refactor error class management (alexcb PR 289)
- Expose KafkaConsumer in \_\_all\_\_ for easy imports (Dinoshauer PR 286)
- SimpleProducer starts on random partition by default (alexcb PR 288)
- Add keys to compressed messages (meandthewallaby PR 281)

- Add new high-level `KafkaConsumer` class based on java client api (dcpk PR 234)
- Add `KeyedProducer.send_messages` api (pubnub PR 277)
- Fix consumer `pending()` method (jettify PR 276)
- Update low-level demo in README (sunisdown PR 274)
- Include key in `KeyedProducer` messages (se7entyse7en PR 268)
- Fix `SimpleConsumer` timeout behavior in `get_messages` (dcpk PR 238)
- Fix error in `consumer.py` test against `max_buffer_size` (rthille/wizzat PR 225/242)
- Improve string concat performance on pypy / py3 (dcpk PR 233)
- Reorg directory layout for consumer/producer/partitioners (dcpk/wizzat PR 232/243)
- Add `OffsetCommitContext` (locationlabs PR 217)
- Metadata Refactor (dcpk PR 223)
- Add Python 3 support (brutasse/wizzat - PR 227)
- Minor cleanups - imports / README / PyPI classifiers (dcpk - PR 221)
- Fix socket test (dcpk - PR 222)
- Fix exception catching bug in `test_failover_integration` (zever - PR 216)

### **5.9.15 0.9.2 (Aug 26, 2014)**

- Warn users that async producer does not reliably handle failures (dcpk - PR 213)
- Fix spurious `ConsumerFetchSizeTooSmall` error in consumer (DataDog - PR 136)
- Use PyLint for static error checking (dcpk - PR 208)
- Strictly enforce str message type in `producer.send_messages` (dcpk - PR 211)
- Add test timers via nose-timer plugin; list 10 slowest timings by default (dcpk)
- Move fetching last known offset logic to a stand alone function (zever - PR 177)
- Improve `KafkaConnection` and add more tests (dcpk - PR 196)
- Raise `TypeError` if necessary when encoding strings (mdaniel - PR 204)
- Use Travis-CI to publish tagged releases to pypi (tkuhman / mumrah)
- Use official binary tarballs for integration tests and parallelize travis tests (dcpk - PR 193)
- Improve new-topic creation handling (wizzat - PR 174)

### **5.9.16 0.9.1 (Aug 10, 2014)**

- Add codec parameter to Producers to enable compression (patricklucas - PR 166)
- Support IPv6 hosts and network (snaury - PR 169)
- Remove dependency on distribute (patricklucas - PR 163)
- Fix connection error timeout and improve tests (wizzat - PR 158)
- `SimpleProducer` randomization of initial round robin ordering (alexcb - PR 139)
- Fix connection timeout in `KafkaClient` and `KafkaConnection` (maciejkula - PR 161)

- Fix seek + commit behavior (wizzat - PR 148)

### 5.9.17 0.9.0 (Mar 21, 2014)

- Connection refactor and test fixes (wizzat - PR 134)
- Fix when partition has no leader (mrtheb - PR 109)
- Change Producer API to take topic as send argument, not as instance variable (rdiomar - PR 111)
- Substantial refactor and Test Fixing (rdiomar - PR 88)
- Fix Multiprocess Consumer on windows (mahendra - PR 62)
- Improve fault tolerance; add integration tests (jimjh)
- PEP8 / Flakes / Style cleanups (Vetoshkin Nikita; mrtheb - PR 59)
- Setup Travis CI (jimjh - PR 53/54)
- Fix import of BufferUnderflowError (jimjh - PR 49)
- Fix code examples in README (StevenLeRoux - PR 47/48)

### 5.9.18 0.8.0

- Changing auto\_commit to False in [SimpleConsumer](kafka/consumer.py), until 0.8.1 is release offset commits are unsupported
- Adding fetch\_size\_bytes to SimpleConsumer constructor to allow for user-configurable fetch sizes
- Allow SimpleConsumer to automatically increase the fetch size if a partial message is read and no other messages were read during that fetch request. The increase factor is 1.5
- Exception classes moved to kafka.common





## A

`add_topic()` (kafka.client.KafkaClient method), 23  
`assign()` (kafka.KafkaConsumer method), 15  
`assignment()` (kafka.KafkaConsumer method), 15

## B

`blacked_out()` (kafka.BrokerConnection method), 26  
`BrokerConnection` (class in kafka), 26

## C

`can_send_more()` (kafka.BrokerConnection method), 26  
`check_version()` (kafka.BrokerConnection method), 26  
`check_version()` (kafka.client.KafkaClient method), 23  
`close()` (kafka.BrokerConnection method), 26  
`close()` (kafka.client.KafkaClient method), 23  
`close()` (kafka.KafkaConsumer method), 16  
`close()` (kafka.KafkaProducer method), 22  
`commit()` (kafka.KafkaConsumer method), 16  
`commit_async()` (kafka.KafkaConsumer method), 16  
`committed()` (kafka.KafkaConsumer method), 16  
`connect()` (kafka.BrokerConnection method), 26  
`connected()` (kafka.BrokerConnection method), 26  
`connected()` (kafka.client.KafkaClient method), 23  
`connecting()` (kafka.BrokerConnection method), 26  
`connection_delay()` (kafka.client.KafkaClient method), 24

## D

`disconnected()` (kafka.BrokerConnection method), 26

## F

`flush()` (kafka.KafkaProducer method), 22

## H

`highwater()` (kafka.KafkaConsumer method), 16

## I

`in_flight_request_count()` (kafka.client.KafkaClient method), 24  
`is_disconnected()` (kafka.client.KafkaClient method), 24

`is_ready()` (kafka.client.KafkaClient method), 24

## K

`KafkaClient` (class in kafka.client), 23  
`KafkaConsumer` (class in kafka), 13  
`KafkaProducer` (class in kafka), 19

## L

`least_loaded_node()` (kafka.client.KafkaClient method), 24

## P

`partitions_for()` (kafka.KafkaProducer method), 22  
`partitions_for_topic()` (kafka.KafkaConsumer method), 17  
`pause()` (kafka.KafkaConsumer method), 17  
`paused()` (kafka.KafkaConsumer method), 17  
`poll()` (kafka.client.KafkaClient method), 24  
`poll()` (kafka.KafkaConsumer method), 17  
`position()` (kafka.KafkaConsumer method), 17

## R

`ready()` (kafka.client.KafkaClient method), 25  
`recv()` (kafka.BrokerConnection method), 26  
`resume()` (kafka.KafkaConsumer method), 17

## S

`schedule()` (kafka.client.KafkaClient method), 25  
`seek()` (kafka.KafkaConsumer method), 18  
`seek_to_beginning()` (kafka.KafkaConsumer method), 18  
`seek_to_end()` (kafka.KafkaConsumer method), 18  
`send()` (kafka.BrokerConnection method), 26  
`send()` (kafka.client.KafkaClient method), 25  
`send()` (kafka.KafkaProducer method), 22  
`set_topics()` (kafka.client.KafkaClient method), 25  
`subscribe()` (kafka.KafkaConsumer method), 18  
`subscription()` (kafka.KafkaConsumer method), 19

## T

`topics()` (kafka.KafkaConsumer method), 19

## U

`unschedule()` (`kafka.client.KafkaClient` method), [25](#)

`unsubscribe()` (`kafka.KafkaConsumer` method), [19](#)