
Python

Release

Jan 19, 2018

Contents

1	Getting Started	3
1.1	Prerequisites	3
1.2	Quick Install	3
1.3	Source Install	3
2	Command List	5
2.1	kafka-assigner	5
3	Contribution guide	7
3.1	Bug fixes	7
3.2	Code guidelines	7
4	API Documentation	9
4.1	kafka.tools	9
4.2	kafka.tools.assigner	10
4.3	kafka.tools	13
5	Indices and tables	15
	Python Module Index	17

kafka-tools is a collection of various tools for managing Apache Kafka.

This document describes how to install and configure kafka-tools.

1.1 Prerequisites

Software that must be installed:

- Python (<https://www.python.org/>)
- Kazoo (<https://kazoo.readthedocs.io/en/latest/>)
- Kafka (<https://kafka.apache.org/>)

In addition, you will need to run it on a host that has:

- Access to the Zookeeper ensemble for the Kafka cluster.
- SSH access to the Kafka brokers (with credentials preferably loaded into ssh-agent).

1.2 Quick Install

1. Download the Kafka binaries from <https://kafka.apache.org/downloads.html>.
2. Use pip to install the kafka-tools package from pypi.

1.3 Source Install

1. Download the Kafka binaries from <https://kafka.apache.org/downloads.html>.
2. Clone the kafka-tools repository from <https://github.com/linkedin/kafka-tools>
3. Run the tests using tox.

4. Install kafka-tools using setup.py.

2.1 kafka-assigner

kafka-assigner is used for performing partition reassignments and preferred replica elections. It uses the admin CLI utilities provided with Kafka and layers on additional logic to perform tasks like removing a broker, rebalancing partitions, fixing partition replication factors, and performing preferred replica elections.

We're always open to fixes and new features! Please open a PR for any changes that you have and someone will review and merge it. If you're not up for writing the code, open an issue for any problems or requests.

3.1 Bug fixes

If you find a bug in “kafka-tools” we welcome patches to help fix the problem.

A few things to keep in mind when submitting bug fixes:

- Please include details on reproducing the bug
- Explain how your patch resolves the issue
- Test cases for the fix are extremely helpful
- Would your patch affect existing expected behaviors? For example, catching an exception may fix an issue for your code, but another caller may expect that exception in order to handle the error case differently.

3.2 Code guidelines

Please follow these Python best practices:

- Provide documentation for new or modified APIs
- Provide test cases for new or modified APIs

4.1 kafka.tools

4.1.1 kafka.tools.exceptions

4.1.2 kafka.tools.models.broker

4.1.3 kafka.tools.models.cluster

4.1.4 kafka.tools.models.partition

4.1.5 kafka.tools.models.topic

4.1.6 kafka.tools.modules

4.1.7 kafka.tools.utilities

`kafka.tools.utilities.check_java_home()`

Make sure that JAVA_HOME in the current environment is specified and is valid.

Raises ConfigurationException if JAVA_HOME is not set or does not contain java

`kafka.tools.utilities.find_path_containing(fname)`

Search the PATH for the given executable filename

Parameters `fname` – the filename to check

Returns the path that contains the filename

Raises ConfigurationException if the filename cannot be found, or if it is not executable

`kafka.tools.utilities.get_tools_path(tools_path=None)`

Find the Kafka admin utilities, either from the provided arg or the PATH.

Parameters `tools_path` – the path to use for locating the Kafka admin utilities.

Returns the path that contains Kafka admin utilities

Raises `ConfigurationException` if the path cannot be determined

`kafka.tools.utilities.is_exec_file` (*fname*)

Check if the given filename is a regular file and is executable.

Parameters `fname` – the filename to check.

Returns True if the filename given exists and is executable, False otherwise

`kafka.tools.utilities.json_loads` (*json_str*)

Load the provided string as JSON data. Make sure to try the python2 way and the python3 way

Parameters `json_str` – The JSON encoded string

Returns The decoded JSON object

`kafka.tools.utilities.synchronized` (*item*)

Decorator that synchronizes access to the instance method it decorates using a preexisting lock in the `_lock` attribute of the instance

4.2 kafka.tools.assigner

4.2.1 kafka.tools.assigner.actions.balance

4.2.2 kafka.tools.assigner.actions.balancemodules.count

4.2.3 kafka.tools.assigner.actions.balancemodules.even

4.2.4 kafka.tools.assigner.actions.balancemodules.leader

4.2.5 kafka.tools.assigner.actions.balancemodules.rackaware

`kafka.tools.assigner.actions.balancemodules.rackaware.check_partition_swappable` (*replicas_a*,
replicas_b,
pos)

Check if the broker at position `pos` in the first replica list can be swapped with the replica at position `pos` in the second list 1. `replicas_a[pos]` must not be in the `replicas_b` list 2. `replicas_b[pos]` must not be in the `replicas_a` list 3. `replicas_a[pos]` must have a different rack than the replicas in `replicas_b` (except for `replicas_b[pos]`) 4. `replicas_b[pos]` must have a different rack than the replicas in `replicas_a` (except for `replicas_a[pos]`)

Params `replicas_a` the first replica list

Params `replicas_b` the second replica list

Params `pos` the position in the replica list to be replaced

Returns True if the broker can be swapped into this replica list, False otherwise

`kafka.tools.assigner.actions.balancemodules.rackaware.difference_in_size_to_last_partition`

Return the difference in size between the specified Partition and the last Partition in the provided list. If the list is empty, return infinity.

Params partition a Partition object to use for calculating the difference

Params partitions a list of Partition objects

Returns The difference in size between partition and the last Partition in the partitions list, or infinity

`kafka.tools.assigner.actions.balancemodules.rackaware.racks_for_replica_list` (*replicas*,
pos=None)

Returns a set of racks for each of the given replicas in the list Skip the replica at position pos, if specified

Params replicas a list of Broker objects

Params pos a replica position to skip, or None to not skip a replica

Returns a list of racks

4.2.6 kafka.tools.assigner.actions.balancemodules.size

4.2.7 kafka.tools.assigner.actions.clone

4.2.8 kafka.tools.assigner.actions.elect

4.2.9 kafka.tools.assigner.actions.remove

4.2.10 kafka.tools.assigner.actions.reorder

4.2.11 kafka.tools.assigner.actions.setrf

4.2.12 kafka.tools.assigner.actions.trim

4.2.13 kafka.tools.assigner.arguments

4.2.14 kafka.tools.assigner.batcher

4.2.15 kafka.tools.assigner.models.reassignment

4.2.16 kafka.tools.assigner.models.replica_election

4.2.17 kafka.tools.assigner.sizers.ssh

4.3 kafka.tools

4.3.1 kafka.tools.protocol.arguments

4.3.2 kafka.tools.protocol.errors

4.3.3 kafka.tools.protocol.help

4.3.4 kafka.tools.protocol.types.bytebuffer

4.3.5 kafka.tools.protocol.requests.api_versions_v0

4.3.6 kafka.tools.protocol.requests.controlled_shutdown_v1

4.3.7 kafka.tools.protocol.requests.create_topics_v0

4.3.8 kafka.tools.protocol.requests.delete_topics_v0

4.3.9 kafka.tools.protocol.requests.describe_groups_v0

4.3.10 kafka.tools.protocol.requests.group_coordinator_v0

4.3.11 kafka.tools.protocol.requests.heartbeat_v0

4.3.12 kafka.tools.protocol.requests.join_group_v0

4.3.13 kafka.tools.protocol.requests.leader_and_isr_v0

4.3.14 kafka.tools.protocol.requests.leave_group_v0

CHAPTER 5

Indices and tables

- `genindex`
- `modindex`
- `search`

Python Module Index

k

kafka.tools.assigner.actions.balance, 10
kafka.tools.assigner.actions.balancemodules.count, 10
kafka.tools.assigner.actions.balancemodules.even, 10
kafka.tools.assigner.actions.balancemodules.leader, 10
kafka.tools.assigner.actions.balancemodules.rackaware, 10
kafka.tools.assigner.actions.balancemodules.size, 13
kafka.tools.assigner.actions.clone, 13
kafka.tools.assigner.actions elect, 13
kafka.tools.assigner.actions.remove, 13
kafka.tools.assigner.actions.reorder, 13
kafka.tools.assigner.actions.setrf, 13
kafka.tools.assigner.actions.trim, 13
kafka.tools.assigner.arguments, 13
kafka.tools.assigner.batcher, 13
kafka.tools.assigner.models.reassignment, 13
kafka.tools.assigner.models.replica_election, 13
kafka.tools.assigner.sizers.ssh, 13
kafka.tools.exceptions, 9
kafka.tools.models.broker, 9
kafka.tools.models.partition, 9
kafka.tools.models.topic, 9
kafka.tools.modules, 9
kafka.tools.protocol.arguments, 13
kafka.tools.protocol.errors, 13
kafka.tools.protocol.help, 13
kafka.tools.protocol.requests.api_versions_v0, 13
kafka.tools.protocol.requests.controlled_shutdown_v1, 13
kafka.tools.protocol.requests.create_topics_v0, 13
kafka.tools.protocol.requests.delete_topics_v0, 13
kafka.tools.protocol.requests.describe_groups_v0, 13
kafka.tools.protocol.requests.group_coordinator_v0, 13
kafka.tools.protocol.requests.heartbeat_v0, 13
kafka.tools.protocol.requests.join_group_v0, 13
kafka.tools.protocol.requests.leader_and_isr_v0, 13
kafka.tools.protocol.requests.leave_group_v0, 13
kafka.tools.protocol.requests.list_groups_v0, 13
kafka.tools.protocol.requests.list_offset_v0, 13
kafka.tools.protocol.requests.offset_commit_v0, 13
kafka.tools.protocol.requests.offset_commit_v1, 13
kafka.tools.protocol.requests.offset_commit_v2, 13
kafka.tools.protocol.requests.offset_fetch_v0, 13
kafka.tools.protocol.requests.offset_fetch_v1, 13
kafka.tools.protocol.requests.sasl_handshake_v0, 13
kafka.tools.protocol.requests.stop_replica_v0, 13
kafka.tools.protocol.requests.sync_group_v0, 13
kafka.tools.protocol.requests.topic_metadata_v0, 13
kafka.tools.protocol.requests.topic_metadata_v1, 13

kafka.tools.protocol.requests.update_metadata_v0,
13
kafka.tools.protocol.responses.api_versions_v0,
13
kafka.tools.protocol.responses.controlled_shutdown_v1,
13
kafka.tools.protocol.responses.create_topics_v0,
13
kafka.tools.protocol.responses.delete_topics_v0,
13
kafka.tools.protocol.responses.describe_groups_v0,
13
kafka.tools.protocol.responses.group_coordinator_v0,
13
kafka.tools.protocol.responses.heartbeat_v0,
13
kafka.tools.protocol.responses.join_group_v0,
13
kafka.tools.protocol.responses.leader_and_isr_v0,
13
kafka.tools.protocol.responses.leave_group_v0,
13
kafka.tools.protocol.responses.list_groups_v0,
13
kafka.tools.protocol.responses.list_offset_v0,
13
kafka.tools.protocol.responses.metadata_v0,
13
kafka.tools.protocol.responses.metadata_v1,
13
kafka.tools.protocol.responses.offset_commit_v0,
13
kafka.tools.protocol.responses.offset_commit_v1,
13
kafka.tools.protocol.responses.offset_commit_v2,
13
kafka.tools.protocol.responses.offset_fetch_v0,
13
kafka.tools.protocol.responses.offset_fetch_v1,
13
kafka.tools.protocol.responses.sasl_handshake_v0,
13
kafka.tools.protocol.responses.stop_replica_v0,
13
kafka.tools.protocol.responses.sync_group_v0,
13
kafka.tools.protocol.responses.update_metadata_v0,
13
kafka.tools.protocol.types.bytebuffer,
13
kafka.tools.utilities, 9

C

check_java_home() (in module kafka.tools.utilities), 9

check_partition_swappable() (in module kafka.tools.assigner.actions.balance
modules.rackaware),
10

D

difference_in_size_to_last_partition() (in module kafka.tools.assigner.actions.balance
modules.rackaware),
10

F

find_path_containing() (in module kafka.tools.utilities), 9

G

get_tools_path() (in module kafka.tools.utilities), 9

I

is_exec_file() (in module kafka.tools.utilities), 10

J

json_loads() (in module kafka.tools.utilities), 10

K

kafka.tools.assigner.actions.balance (module), 10

kafka.tools.assigner.actions.balance.modules.count (mod-
ule), 10

kafka.tools.assigner.actions.balance.modules.even (mod-
ule), 10

kafka.tools.assigner.actions.balance.modules.leader
(module), 10

kafka.tools.assigner.actions.balance.modules.rackaware
(module), 10

kafka.tools.assigner.actions.balance.modules.size (mod-
ule), 13

kafka.tools.assigner.actions.clone (module), 13

kafka.tools.assigner.actions.elect (module), 13

kafka.tools.assigner.actions.remove (module), 13

kafka.tools.assigner.actions.reorder (module), 13

kafka.tools.assigner.actions.setrf (module), 13

kafka.tools.assigner.actions.trim (module), 13

kafka.tools.assigner.arguments (module), 13

kafka.tools.assigner.batcher (module), 13

kafka.tools.assigner.models.reassignment (module), 13

kafka.tools.assigner.models.replica_election (module), 13

kafka.tools.assigner.sizers.ssh (module), 13

kafka.tools.exceptions (module), 9

kafka.tools.models.broker (module), 9

kafka.tools.models.partition (module), 9

kafka.tools.models.topic (module), 9

kafka.tools.modules (module), 9

kafka.tools.protocol.arguments (module), 13

kafka.tools.protocol.errors (module), 13

kafka.tools.protocol.help (module), 13

kafka.tools.protocol.requests.api_versions_v0 (module),
13

kafka.tools.protocol.requests.controlled_shutdown_v1
(module), 13

kafka.tools.protocol.requests.create_topics_v0 (module),
13

kafka.tools.protocol.requests.delete_topics_v0 (module),
13

kafka.tools.protocol.requests.describe_groups_v0 (mod-
ule), 13

kafka.tools.protocol.requests.group_coordinator_v0
(module), 13

kafka.tools.protocol.requests.heartbeat_v0 (module), 13

kafka.tools.protocol.requests.join_group_v0 (module), 13

kafka.tools.protocol.requests.leader_and_isr_v0 (mod-
ule), 13

kafka.tools.protocol.requests.leave_group_v0 (module),
13

kafka.tools.protocol.requests.list_groups_v0 (module),
13

kafka.tools.protocol.requests.list_offset_v0 (module), 13

kafka.tools.protocol.requests.offset_commit_v0 (mod-
ule), 13

kafka.tools.protocol.requests.offset_commit_v1 (mod-
ule), 13

kafka.tools.protocol.requests.offset_commit_v2 (module), 13
kafka.tools.protocol.requests.offset_fetch_v0 (module), 13
kafka.tools.protocol.requests.offset_fetch_v1 (module), 13
kafka.tools.protocol.requests.sasl_handshake_v0 (module), 13
kafka.tools.protocol.requests.stop_replica_v0 (module), 13
kafka.tools.protocol.requests.sync_group_v0 (module), 13
kafka.tools.protocol.requests.topic_metadata_v0 (module), 13
kafka.tools.protocol.requests.topic_metadata_v1 (module), 13
kafka.tools.protocol.requests.update_metadata_v0 (module), 13
kafka.tools.protocol.responses.api_versions_v0 (module), 13
kafka.tools.protocol.responses.controlled_shutdown_v1 (module), 13
kafka.tools.protocol.responses.create_topics_v0 (module), 13
kafka.tools.protocol.responses.delete_topics_v0 (module), 13
kafka.tools.protocol.responses.describe_groups_v0 (module), 13
kafka.tools.protocol.responses.group_coordinator_v0 (module), 13
kafka.tools.protocol.responses.heartbeat_v0 (module), 13
kafka.tools.protocol.responses.join_group_v0 (module), 13
kafka.tools.protocol.responses.leader_and_isr_v0 (module), 13
kafka.tools.protocol.responses.leave_group_v0 (module), 13
kafka.tools.protocol.responses.list_groups_v0 (module), 13
kafka.tools.protocol.responses.list_offset_v0 (module), 13
kafka.tools.protocol.responses.metadata_v0 (module), 13
kafka.tools.protocol.responses.metadata_v1 (module), 13
kafka.tools.protocol.responses.offset_commit_v0 (module), 13
kafka.tools.protocol.responses.offset_commit_v1 (module), 13
kafka.tools.protocol.responses.offset_commit_v2 (module), 13
kafka.tools.protocol.responses.offset_fetch_v0 (module), 13
kafka.tools.protocol.responses.offset_fetch_v1 (module), 13
kafka.tools.protocol.responses.sasl_handshake_v0 (module), 13
kafka.tools.protocol.responses.stop_replica_v0 (module), 13
kafka.tools.protocol.responses.sync_group_v0 (module), 13
kafka.tools.protocol.responses.update_metadata_v0 (module), 13
kafka.tools.protocol.types.bytebuffer (module), 13
kafka.tools.utilities (module), 9

R
racks_for_replica_list() (in module kafka.tools.assigner.actions.balancemodules.rackaware), 11

S
synchronized() (in module kafka.tools.utilities), 10