
jupyter_client Documentation

Release 5.2.3

Jupyter Development Team

Mar 12, 2018

1	Messaging in Jupyter	3
1.1	Versioning	3
1.2	Introduction	4
1.3	General Message Format	5
1.4	Compatibility	6
1.5	The Wire Protocol	6
1.6	Python API	7
1.7	Messages on the shell (ROUTER/DEALER) channel	7
1.8	Messages on the IOPub (PUB/SUB) channel	18
1.9	Messages on the stdin (ROUTER/DEALER) channel	22
1.10	Heartbeat for kernels	23
1.11	Custom Messages	23
1.12	Notes	24
2	Making kernels for Jupyter	27
2.1	Connection files	27
2.2	Handling messages	28
2.3	Kernel specs	28
3	Making simple Python wrapper kernels	31
3.1	Required steps	31
3.2	Example	32
3.3	Optional steps	33
4	jupyter_client API	35
4.1	kernelspec - discovering kernels	35
4.2	manager - starting, stopping, signalling	36
4.3	client - communicating with kernels	38
5	Changes in Jupyter Client	43
5.1	5.2.3	43
5.2	5.2.2	43
5.3	5.2.1	43
5.4	5.2	44
5.5	5.1	44
5.6	5.0	45
5.7	4.4	45

5.8	4.3	46
5.9	4.2	46
5.10	4.1	46
5.11	4.0	47
6	Indices and tables		49
	Python Module Index		51

This package provides the Python API for starting, managing and communicating with Jupyter kernels.

Important: This document contains the authoritative description of the Jupyter messaging protocol. All developers are strongly encouraged to keep it updated as the implementation evolves, so that we have a single common reference for all protocol details.

CHAPTER 1

Messaging in Jupyter

This document explains the basic communications design and messaging specification for how Jupyter frontends and kernels communicate. The [ZeroMQ](#) library provides the low-level transport layer over which these messages are sent.

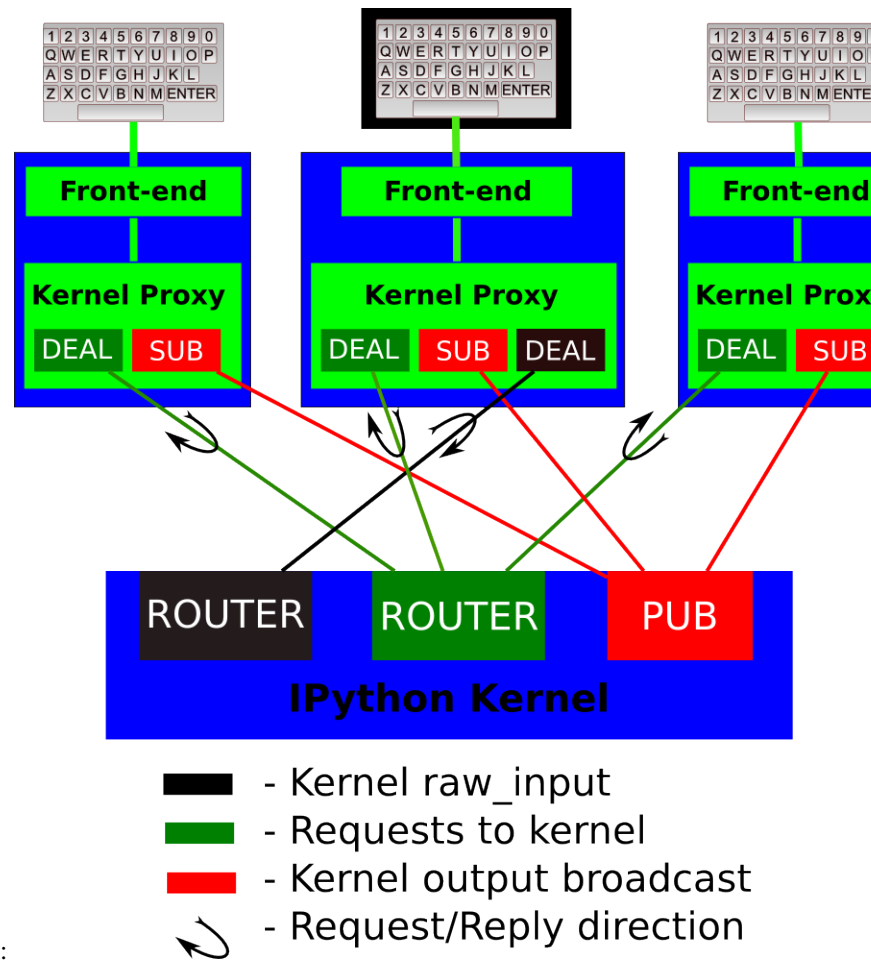
Important: This document contains the authoritative description of the IPython messaging protocol. All developers are strongly encouraged to keep it updated as the implementation evolves, so that we have a single common reference for all protocol details.

1.1 Versioning

The Jupyter message specification is versioned independently of the packages that use it. The current version of the specification is 5.3.

Note: *New in* and *Changed in* messages in this document refer to versions of the **Jupyter message specification**, not versions of [jupyter_client](#).

1.2 Introduction



The basic design is explained in the following diagram:

A single kernel can be simultaneously connected to one or more frontends. The kernel has dedicated sockets for the following functions:

1. **Shell:** this single ROUTER socket allows multiple incoming connections from frontends, and this is the socket where requests for code execution, object information, prompts, etc. are made to the kernel by any frontend. The communication on this socket is a sequence of request/reply actions from each frontend and the kernel.
2. **IOPub:** this socket is the ‘broadcast channel’ where the kernel publishes all side effects (stdout, stderr, etc.) as well as the requests coming from any client over the shell socket and its own requests on the stdin socket. There are a number of actions in Python which generate side effects: `print()` writes to `sys.stdout`, errors generate tracebacks, etc. Additionally, in a multi-client scenario, we want all frontends to be able to know what each other has sent to the kernel (this can be useful in collaborative scenarios, for example). This socket allows both side effects and the information about communications taking place with one client over the shell channel to be made available to all clients in a uniform manner.
3. **stdin:** this ROUTER socket is connected to all frontends, and it allows the kernel to request input from the active frontend when `raw_input()` is called. The frontend that executed the code has a DEALER socket that acts as a ‘virtual keyboard’ for the kernel while this communication is happening (illustrated in the figure by the black outline around the central keyboard). In practice, frontends may display such kernel requests using a special input widget or otherwise indicating that the user is to type input for the kernel instead of normal commands in the frontend.

All messages are tagged with enough information (details below) for clients to know which messages come

from their own interaction with the kernel and which ones are from other clients, so they can display each type appropriately.

4. **Control:** This channel is identical to Shell, but operates on a separate socket, to allow important messages to avoid queueing behind execution requests (e.g. shutdown or abort).
5. **Heartbeat:** This socket allows for simple bytestring messages to be sent between the frontend and the kernel to ensure that they are still connected.

The actual format of the messages allowed on each of these channels is specified below. Messages are dicts of dicts with string keys and values that are reasonably representable in JSON.

1.3 General Message Format

A message is defined by the following four-dictionary structure:

```
{
  # The message header contains a pair of unique identifiers for the
  # originating session and the actual message id, in addition to the
  # username for the process that generated the message. This is useful in
  # collaborative settings where multiple users may be interacting with the
  # same kernel simultaneously, so that frontends can label the various
  # messages in a meaningful way.
  'header' : {
    'msg_id' : str, # typically UUID, must be unique per message
    'username' : str,
    'session' : str, # typically UUID, should be unique per session
    # ISO 8601 timestamp for when the message is created
    'date': str,
    # All recognized message type strings are listed below.
    'msg_type' : str,
    # the message protocol version
    'version' : '5.0',

  },

  # In a chain of messages, the header from the parent is copied so that
  # clients can track where messages come from.
  'parent_header' : dict,

  # Any metadata associated with the message.
  'metadata' : dict,

  # The actual content of the message must be a dict, whose structure
  # depends on the message type.
  'content' : dict,

  # optional: buffers is a list of binary data buffers for implementations
  # that support binary extensions to the protocol.
  'buffers': list,
}
```

Changed in version 5.0: version key added to the header.

Changed in version 5.1: date in the header was accidentally omitted from the spec prior to 5.1, but it has always been in the canonical implementation, so implementers are strongly encouraged to include it. It will be mandatory in 5.1.

1.4 Compatibility

Kernels must implement the *execute* and *kernel info* messages in order to be usable. All other message types are optional, although we recommend implementing *completion* if possible. Kernels do not need to send any reply for messages they don't handle, and frontends should provide sensible behaviour if no reply arrives (except for the required execution and kernel info messages).

stdin messages are unique in that the request comes from the kernel, and the reply from the frontend. The frontend is not required to support this, but if it does not, it must set 'allow_stdin' : False in its *execute requests*. In this case, the kernel may not send stdin requests. If that field is true, the kernel may send stdin requests and block waiting for a reply, so the frontend must answer.

Both sides should allow unexpected message types, and extra fields in known message types, so that additions to the protocol do not break existing code.

1.5 The Wire Protocol

This message format exists at a high level, but does not describe the actual *implementation* at the wire level in zeromq. The canonical implementation of the message spec is our `Session` class.

Note: This section should only be relevant to non-Python consumers of the protocol. Python consumers should simply import and use implementation of the wire protocol in `jupyter_client.session.Session`.

Every message is serialized to a sequence of at least six blobs of bytes:

```
[
    b'u-u-i-d',          # zmq identity(ies)
    b'<IDS|MSG>',        # delimiter
    b'baddad42',         # HMAC signature
    b'{header}',          # serialized header dict
    b'{parent_header}',  # serialized parent header dict
    b'{metadata}',       # serialized metadata dict
    b'{content}',        # serialized content dict
    b'\xf0\xf9\x00\x01' # extra raw data buffer(s)
    ...
]
```

The front of the message is the ZeroMQ routing prefix, which can be zero or more socket identities. This is every piece of the message prior to the delimiter key `<IDS|MSG>`. In the case of `IOPub`, there should be just one prefix component, which is the topic for `IOPub` subscribers, e.g. `execute_result`, `display_data`.

Note: In most cases, the `IOPub` topics are irrelevant and completely ignored, because frontends just subscribe to all topics. The convention used in the IPython kernel is to use the `msg_type` as the topic, and possibly extra information about the message, e.g. `kernel.{u-u-i-d}.execute_result` or `stream.stdout`

After the delimiter is the **HMAC** signature of the message, used for authentication. If authentication is disabled, this should be an empty string. By default, the hashing function used for computing these signatures is `sha256`.

Note: To disable authentication and signature checking, set the `key` field of a connection file to an empty string.

The signature is the HMAC hex digest of the concatenation of:

- A shared key (typically the `key` field of a connection file)
- The serialized header dict
- The serialized parent header dict
- The serialized metadata dict
- The serialized content dict

In Python, this is implemented via:

```
# once:
digester = HMAC(key, digestmod=hashlib.sha256)

# for each message
d = digester.copy()
for serialized_dict in (header, parent, metadata, content):
    d.update(serialized_dict)
signature = d.hexdigest()
```

After the signature is the actual message, always in four frames of bytes. The four dictionaries that compose a message are serialized separately, in the order of header, parent header, metadata, and content. These can be serialized by any function that turns a dict into bytes. The default and most common serialization is JSON, but msgpack and pickle are common alternatives.

After the serialized dicts are zero to many raw data buffers, which can be used by message types that support binary data, which can be used in custom messages, such as comms and extensions to the protocol.

1.6 Python API

As messages are dicts, they map naturally to a `func(**kw)` call form. We should develop, at a few key points, functional forms of all the requests that take arguments in this manner and automatically construct the necessary dict for sending.

In addition, the Python implementation of the message specification extends messages upon deserialization to the following form for convenience:

```
{
  'header' : dict,
  # The msg's unique identifier and type are always stored in the header,
  # but the Python implementation copies them to the top level.
  'msg_id' : str,
  'msg_type' : str,
  'parent_header' : dict,
  'content' : dict,
  'metadata' : dict,
}
```

All messages sent to or received by any IPython process should have this extended structure.

1.7 Messages on the shell (ROUTER/DEALER) channel

1.7.1 Request-Reply

In general, the ROUTER/DEALER sockets follow a request-reply pattern:

The client sends an `<action>_request` message (such as `execute_request`) on its shell (DEALER) socket. The kernel receives that request and immediately publishes a `status: busy` message on IOPub. The kernel then processes the request and sends the appropriate `<action>_reply` message, such as `execute_reply`. After processing the request and publishing associated IOPub messages, if any, the kernel publishes a `status: idle` message. This idle status message indicates that IOPub messages associated with a given request have all been received.

All reply messages have a `'status'` field, which will have one of the following values:

- `status='ok'`: The request was processed successfully, and the remaining content of the reply is specified in the appropriate section below.
- **`status='error'`: The request failed due to an error.** When status is `'error'`, the usual content of a successful reply should be omitted, instead the following fields should be present:

```
{
  'status' : 'error',
  'ename' : str,    # Exception name, as a string
  'evalue' : str,   # Exception value, as a string
  'traceback' : list(str), # traceback frames as strings
}
```

- `status='abort'`: This is the same as `status='error'` but with no information about the error. No fields should be present other than `status`.

Changed in version 5.1: `status='abort'` has not proved useful, and is considered deprecated. Kernels should send `status='error'` instead.

1.7.2 Execute

This message type is used by frontends to ask the kernel to execute code on behalf of the user, in a namespace reserved to the user's variables (and thus separate from the kernel's own internal code and variables).

Message type: `execute_request`:

```
content = {
    # Source code to be executed by the kernel, one or more lines.
    'code' : str,

    # A boolean flag which, if True, signals the kernel to execute
    # this code as quietly as possible.
    # silent=True forces store_history to be False,
    # and will *not*:
    #   - broadcast output on the IOPUB channel
    #   - have an execute_result
    # The default is False.
    'silent' : bool,

    # A boolean flag which, if True, signals the kernel to populate history
    # The default is True if silent is False. If silent is True, store_history
    # is forced to be False.
    'store_history' : bool,

    # A dict mapping names to expressions to be evaluated in the
    # user's dict. The rich display-data representation of each will be evaluated after_
    ↪ execution.
    # See the display_data content for the structure of the representation data.
    'user_expressions' : dict,
```

(continues on next page)

(continued from previous page)

```
# Some frontends do not support stdin requests.
# If this is true, code running in the kernel can prompt the user for input
# with an input_request message (see below). If it is false, the kernel
# should not send these messages.
'allow_stdin' : True,

# A boolean flag, which, if True, does not abort the execution queue, if an exception
↳is encountered.
# This allows the queued execution of multiple execute_requests, even if they
↳generate exceptions.
'stop_on_error' : False,
}
```

Changed in version 5.0: `user_variables` removed, because it is redundant with `user_expressions`.

The `code` field contains a single string (possibly multiline) to be executed.

The `user_expressions` field deserves a detailed explanation. In the past, IPython had the notion of a prompt string that allowed arbitrary code to be evaluated, and this was put to good use by many in creating prompts that displayed system status, path information, and even more esoteric uses like remote instrument status acquired over the network. But now that IPython has a clean separation between the kernel and the clients, the kernel has no prompt knowledge; prompts are a frontend feature, and it should be even possible for different frontends to display different prompts while interacting with the same kernel. `user_expressions` can be used to retrieve this information.

Any error in evaluating any expression in `user_expressions` will result in only that key containing a standard error message, of the form:

```
{
  'status' : 'error',
  'ename' : 'NameError',
  'evalue' : 'foo',
  'traceback' : ...
}
```

Note: In order to obtain the current execution counter for the purposes of displaying input prompts, frontends may make an execution request with an empty code string and `silent=True`.

Upon completion of the execution request, the kernel *always* sends a reply, with a status code indicating what happened and additional data depending on the outcome. See [below](#) for the possible return codes and associated data.

See also:

[Execution semantics in the IPython kernel](#)

Execution counter (prompt number)

The kernel should have a single, monotonically increasing counter of all execution requests that are made with `store_history=True`. This counter is used to populate the `In[n]` and `Out[n]` prompts. The value of this counter will be returned as the `execution_count` field of all `execute_reply` and `execute_input` messages.

Execution results

Message type: `execute_reply`:

```
content = {
    # One of: 'ok' OR 'error' OR 'abort'
    'status' : str,

    # The global kernel counter that increases by one with each request that
    # stores history. This will typically be used by clients to display
    # prompt numbers to the user. If the request did not store history, this will
    # be the current value of the counter in the kernel.
    'execution_count' : int,
}
```

When status is `'ok'`, the following extra fields are present:

```
{
    # 'payload' will be a list of payload dicts, and is optional.
    # payloads are considered deprecated.
    # The only requirement of each payload dict is that it have a 'source' key,
    # which is a string classifying the payload (e.g. 'page').

    'payload' : list(dict),

    # Results for the user_expressions.
    'user_expressions' : dict,
}
```

Changed in version 5.0: `user_variables` is removed, use `user_expressions` instead.

Payloads (DEPRECATED)

Execution payloads

Payloads are considered **deprecated**, though their replacement is not yet implemented.

Payloads are a way to trigger frontend actions from the kernel. Current payloads:

page: display data in a pager.

Pager output is used for introspection, or other displayed information that's not considered output. Pager payloads are generally displayed in a separate pane, that can be viewed alongside code, and are not included in notebook documents.

```
{
    "source": "page",
    # mime-bundle of data to display in the pager.
    # Must include text/plain.
    "data": mimebundle,
    # line offset to start from
    "start": int,
}
```

set_next_input: create a new output

used to create new cells in the notebook, or set the next input in a console interface. The main example being `%load`.

```
{
  "source": "set_next_input",
  # the text contents of the cell to create
  "text": "some cell content",
  # If true, replace the current cell in document UIs instead of inserting
  # a cell. Ignored in console UIs.
  "replace": bool,
}
```

edit: open a file for editing.

Triggered by `%edit`. Only the QtConsole currently supports edit payloads.

```
{
  "source": "edit",
  "filename": "/path/to/file.py", # the file to edit
  "line_number": int, # the line number to start with
}
```

ask_exit: instruct the frontend to prompt the user for exit

Allows the kernel to request exit, e.g. via `%exit` in IPython. Only for console frontends.

```
{
  "source": "ask_exit",
  # whether the kernel should be left running, only closing the client
  "keepkernel": bool,
}
```

1.7.3 Introspection

Code can be inspected to show useful information to the user. It is up to the Kernel to decide what information should be displayed, and its formatting.

Message type: `inspect_request`:

```
content = {
  # The code context in which introspection is requested
  # this may be up to an entire multiline cell.
  'code' : str,

  # The cursor position within 'code' (in unicode characters) where inspection is_
  ↪requested
  'cursor_pos' : int,

  # The level of detail desired. In IPython, the default (0) is equivalent to_
  ↪typing
  # 'x?' at the prompt, 1 is equivalent to 'x??'.
  # The difference is up to kernels, but in IPython level 1 includes the source code
  # if available.
  'detail_level' : 0 or 1,
}
```

Changed in version 5.0: `object_info_request` renamed to `inspect_request`.

Changed in version 5.0: `name` key replaced with `code` and `cursor_pos`, moving the lexing responsibility to the kernel.

Changed in version 5.2: Due to a widespread bug in many frontends, `cursor_pos` in versions prior to 5.2 is ambiguous in the presence of “astral-plane” characters. In 5.2, `cursor_pos` **must be** the actual encoding-independent offset in unicode codepoints. See [cursor_pos and unicode offsets](#) for more.

The reply is a mime-bundle, like a [display_data](#) message, which should be a formatted representation of information about the context. In the notebook, this is used to show tooltips over function calls, etc.

Message type: `inspect_reply`:

```
content = {
    # 'ok' if the request succeeded or 'error', with error information as in all_
    ↪ other replies.
    'status' : 'ok',

    # found should be true if an object was found, false otherwise
    'found' : bool,

    # data can be empty if nothing is found
    'data' : dict,
    'metadata' : dict,
}
```

Changed in version 5.0: `object_info_reply` renamed to `inspect_reply`.

Changed in version 5.0: Reply is changed from structured data to a mime bundle, allowing formatting decisions to be made by the kernel.

1.7.4 Completion

Message type: `complete_request`:

```
content = {
    # The code context in which completion is requested
    # this may be up to an entire multiline cell, such as
    # 'foo = a.isal'
    'code' : str,

    # The cursor position within 'code' (in unicode characters) where completion is_
    ↪ requested
    'cursor_pos' : int,
}
```

Changed in version 5.0: `line`, `block`, and `text` keys are removed in favor of a single `code` for context. Lexing is up to the kernel.

Changed in version 5.2: Due to a widespread bug in many frontends, `cursor_pos` in versions prior to 5.2 is ambiguous in the presence of “astral-plane” characters. In 5.2, `cursor_pos` **must be** the actual encoding-independent offset in unicode codepoints. See [cursor_pos and unicode offsets](#) for more.

Message type: `complete_reply`:

```
content = {
    # The list of all matches to the completion request, such as
    # ['a.isalnum', 'a.isalpha'] for the above example.
    'matches' : list,

    # The range of text that should be replaced by the above matches when a completion is_
    ↪ accepted.
}
```

(continues on next page)

(continued from previous page)

```
# typically cursor_end is the same as cursor_pos in the request.
'cursor_start' : int,
'cursor_end' : int,

# Information that frontend plugins might use for extra display information about
# completions.
'metadata' : dict,

# status should be 'ok' unless an exception was raised during the request,
# in which case it should be 'error', along with the usual error message content
# in other messages.
'status' : 'ok'
}
```

Changed in version 5.0: `matched_text` is removed in favor of `cursor_start` and `cursor_end`. `metadata` is added for extended information.

1.7.5 History

For clients to explicitly request history from a kernel. The kernel has all the actual execution history stored in a single location, so clients can request it from the kernel when needed.

Message type: `history_request`:

```
• content = {

    # If True, also return output history in the resulting dict.
    'output' : bool,

    # If True, return the raw input history, else the transformed input.
    'raw' : bool,

    # So far, this can be 'range', 'tail' or 'search'.
    'hist_access_type' : str,

    # If hist_access_type is 'range', get a range of input cells. session
    # is a number counting up each time the kernel starts; you can give
    # a positive session number, or a negative number to count back from
    # the current session.
    'session' : int,
    # start and stop are line (cell) numbers within that session.
    'start' : int,
    'stop' : int,

    # If hist_access_type is 'tail' or 'search', get the last n cells.
    'n' : int,

    # If hist_access_type is 'search', get cells matching the specified glob
    # pattern (with * and ? as wildcards).
    'pattern' : str,

    # If hist_access_type is 'search' and unique is true, do not
    # include duplicated history. Default is false.
    'unique' : bool,

}
```

New in version 4.0: The key `unique` for `history_request`.

Message type: `history_reply`:

```
content = {
    # A list of 3 tuples, either:
    # (session, line_number, input) or
    # (session, line_number, (input, output)),
    # depending on whether output was False or True, respectively.
    'history' : list,
}
```

Note: Most of the history messaging options are not used by Jupyter frontends, and many kernels do not implement them. If you're implementing these messages in a kernel, the 'tail' request is the most useful; this is used by the Qt console, for example. The notebook interface does not use history messages at all.

This interface was designed by exposing all the main options of IPython's history interface. We may remove some options in a future version of the message spec.

1.7.6 Code completeness

New in version 5.0.

When the user enters a line in a console style interface, the console must decide whether to immediately execute the current code, or whether to show a continuation prompt for further input. For instance, in Python `a = 5` would be executed immediately, while `for i in range(5):` would expect further input.

There are four possible replies:

- *complete* code is ready to be executed
- *incomplete* code should prompt for another line
- *invalid* code will typically be sent for execution, so that the user sees the error soonest.
- *unknown* - if the kernel is not able to determine this. The frontend should also handle the kernel not replying promptly. It may default to sending the code for execution, or it may implement simple fallback heuristics for whether to execute the code (e.g. execute after a blank line).

Frontends may have ways to override this, forcing the code to be sent for execution or forcing a continuation prompt.

Message type: `is_complete_request`:

```
content = {
    # The code entered so far as a multiline string
    'code' : str,
}
```

Message type: `is_complete_reply`:

```
content = {
    # One of 'complete', 'incomplete', 'invalid', 'unknown'
    'status' : str,

    # If status is 'incomplete', indent should contain the characters to use
    # to indent the next line. This is only a hint: frontends may ignore it
    # and use their own autoindentation rules. For other statuses, this
```

(continues on next page)

(continued from previous page)

```
# field does not exist.
'indent': str,
}
```

1.7.7 Connect

Deprecated since version 5.1: `connect_request/reply` have not proved useful, and are considered deprecated. Kernels are not expected to implement handlers for this message.

When a client connects to the request/reply socket of the kernel, it can issue a connect request to get basic information about the kernel, such as the ports the other ZeroMQ sockets are listening on. This allows clients to only have to know about a single port (the shell channel) to connect to a kernel. The ports for any additional channels the kernel is listening on should be included in the reply. If any ports are omitted from the reply, this indicates that the channels are not running.

Message type: `connect_request`:

```
content = {}
```

For example, a kernel with all channels running:

Message type: `connect_reply`:

```
content = {
    'shell_port' : int,      # The port the shell ROUTER socket is listening on.
    'iopub_port' : int,     # The port the PUB socket is listening on.
    'stdin_port' : int,     # The port the stdin ROUTER socket is listening on.
    'hb_port' : int,        # The port the heartbeat socket is listening on.
    'control_port' : int,    # The port the control ROUTER socket is listening on.
}
```

1.7.8 Comm info

When a client needs the currently open comms in the kernel, it can issue a request for the currently open comms. When the optional `target_name` is specified, the reply only contains the currently open comms for the target.

Message type: `comm_info_request`:

```
content = {
    # Optional, the target name
    'target_name': str,
}
```

Message type: `comm_info_reply`:

```
content = {
    # A dictionary of the comms, indexed by uuids.
    'comms': {
        comm_id: {
            'target_name': str,
        },
    },
}
```

New in version 5.1: `comm_info` is a proposed addition for msgspec v5.1.

1.7.9 Kernel info

If a client needs to know information about the kernel, it can make a request of the kernel's information. This message can be used to fetch core information of the kernel, including language (e.g., Python), language version number and IPython version number, and the IPython message spec version number.

Message type: `kernel_info_request`:

```
content = {  
}
```

Message type: `kernel_info_reply`:

```
content = {  
    # Version of messaging protocol.  
    # The first integer indicates major version. It is incremented when  
    # there is any backward incompatible change.  
    # The second integer indicates minor version. It is incremented when  
    # there is any backward compatible change.  
    'protocol_version': 'X.Y.Z',  
  
    # The kernel implementation name  
    # (e.g. 'ipython' for the IPython kernel)  
    'implementation': str,  
  
    # Implementation version number.  
    # The version number of the kernel's implementation  
    # (e.g. IPython.__version__ for the IPython kernel)  
    'implementation_version': 'X.Y.Z',  
  
    # Information about the language of code for the kernel  
    'language_info': {  
        # Name of the programming language that the kernel implements.  
        # Kernel included in IPython returns 'python'.  
        'name': str,  
  
        # Language version number.  
        # It is Python version number (e.g., '2.7.3') for the kernel  
        # included in IPython.  
        'version': 'X.Y.Z',  
  
        # mimetype for script files in this language  
        'mimetype': str,  
  
        # Extension including the dot, e.g. '.py'  
        'file_extension': str,  
  
        # Pygments lexer, for highlighting  
        # Only needed if it differs from the 'name' field.  
        'pygments_lexer': str,  
  
        # Codemirror mode, for for highlighting in the notebook.  
        # Only needed if it differs from the 'name' field.  
        'codemirror_mode': str or dict,  
  
        # Nbconvert exporter, if notebooks written with this kernel should  
        # be exported with something other than the general 'script'  
        # exporter.
```

(continues on next page)

(continued from previous page)

```

    'nbconvert_exporter': str,
},

# A banner of information about the kernel,
# which may be displayed in console environments.
'banner' : str,

# Optional: A list of dictionaries, each with keys 'text' and 'url'.
# These will be displayed in the help menu in the notebook UI.
'help_links': [
    {'text': str, 'url': str}
],
}

```

Refer to the lists of available [Pygments lexers](#) and [codemirror modes](#) for those fields.

Changed in version 5.0: Versions changed from lists of integers to strings.

Changed in version 5.0: `ipython_version` is removed.

Changed in version 5.0: `language_info`, `implementation`, `implementation_version`, `banner` and `help_links` keys are added.

Changed in version 5.0: `language_version` moved to `language_info.version`

Changed in version 5.0: `language` moved to `language_info.name`

1.7.10 Kernel shutdown

The clients can request the kernel to shut itself down; this is used in multiple cases:

- when the user chooses to close the client application via a menu or window control.
- when the user types ‘exit’ or ‘quit’ (or their uppercase magic equivalents).
- when the user chooses a GUI method (like the ‘Ctrl-C’ shortcut in the IPythonQt client) to force a kernel restart to get a clean kernel without losing client-side state like history or inlined figures.

The client sends a shutdown request to the kernel, and once it receives the reply message (which is otherwise empty), it can assume that the kernel has completed shutdown safely. The request can be sent on either the *control* or *shell* channels.

Upon their own shutdown, client applications will typically execute a last minute sanity check and forcefully terminate any kernel that is still alive, to avoid leaving stray processes in the user’s machine.

Message type: `shutdown_request`:

```

content = {
    'restart' : bool # False if final shutdown, or True if shutdown precedes a restart
}

```

Message type: `shutdown_reply`:

```

content = {
    'restart' : bool # False if final shutdown, or True if shutdown precedes a restart
}

```

Note: When the clients detect a dead kernel thanks to inactivity on the heartbeat socket, they simply send a forceful process termination signal, since a dead process is unlikely to respond in any useful way to messages.

1.7.11 Kernel interrupt

In case a kernel can not catch operating system interrupt signals (e.g. the used runtime handles signals and does not allow a user program to define a callback), a kernel can choose to be notified using a message instead. For this to work, the kernels `kernelspec` must set `interrupt_mode` to `message`. An interruption will then result in the following message on the `control` channel:

Message type: `interrupt_request`:

```
content = {}
```

Message type: `interrupt_reply`:

```
content = {}
```

New in version 5.3.

1.8 Messages on the IOPub (PUB/SUB) channel

1.8.1 Streams (stdout, stderr, etc)

Message type: `stream`:

```
content = {
    # The name of the stream is one of 'stdout', 'stderr'
    'name' : str,

    # The text is an arbitrary string to be written to that stream
    'text' : str,
}
```

Changed in version 5.0: ‘data’ key renamed to ‘text’ for consistency with the notebook format.

1.8.2 Display Data

This type of message is used to bring back data that should be displayed (text, html, svg, etc.) in the frontends. This data is published to all frontends. Each message can have multiple representations of the data; it is up to the frontend to decide which to use and how. A single message should contain all possible representations of the same information. Each representation should be a JSON’able data structure, and should be a valid MIME type.

Some questions remain about this design:

- Do we use this message type for `execute_result/displayhook`? Probably not, because the `displayhook` also has to handle the Out prompt display. On the other hand we could put that information into the metadata section.

Message type: `display_data`:

```

content = {

    # Who create the data
    # Used in V4. Removed in V5.
    # 'source' : str,

    # The data dict contains key/value pairs, where the keys are MIME
    # types and the values are the raw data of the representation in that
    # format.
    'data' : dict,

    # Any metadata that describes the data
    'metadata' : dict,

    # Optional transient data introduced in 5.1. Information not to be
    # persisted to a notebook or other documents. Intended to live only
    # during a live kernel session.
    'transient': dict,
}

```

The metadata contains any metadata that describes the output. Global keys are assumed to apply to the output as a whole. The metadata dict can also contain mime-type keys, which will be sub-dictionaries, which are interpreted as applying only to output of that type. Third parties should put any data they write into a single dict with a reasonably unique name to avoid conflicts.

The only metadata keys currently defined in IPython are the width and height of images:

```

metadata = {
    'image/png' : {
        'width': 640,
        'height': 480
    }
}

```

and expanded for JSON data:

```

metadata = {
    'application/json' : {
        'expanded': True
    }
}

```

The transient dict contains runtime metadata that should not be persisted to document formats and is fully optional. The only transient key currently defined in Jupyter is `display_id`:

```

transient = {
    'display_id': 'abcd'
}

```

Changed in version 5.0: *application/json* data should be unpacked JSON data, not double-serialized as a JSON string.

Changed in version 5.1: *transient* is a new field.

1.8.3 Update Display Data

New in version 5.1.

Displays can now be named with a `display_id` within the `transient` field of `display_data` or `execute_result`.

When a `display_id` is specified for a display, it can be updated later with an `update_display_data` message. This message has the same format as [display_data](#) messages and must contain a `transient` field with a `display_id`.

Message type: `update_display_data`:

```
content = {  
  
    # The data dict contains key/value pairs, where the keys are MIME  
    # types and the values are the raw data of the representation in that  
    # format.  
    'data' : dict,  
  
    # Any metadata that describes the data  
    'metadata' : dict,  
  
    # Any information not to be persisted to a notebook or other environment  
    # Intended to live only during a kernel session  
    'transient': dict,  
}
```

Frontends can choose how they update prior outputs (or if they regard this as a regular `display_data` message). Within the jupyter and `nteract` notebooks, all displays that match the `display_id` are updated (even if there are multiple).

1.8.4 Code inputs

To let all frontends know what code is being executed at any given time, these messages contain a re-broadcast of the code portion of an [execute_request](#), along with the [execution_count](#).

Message type: `execute_input`:

```
content = {  
    'code' : str, # Source code to be executed, one or more lines  
  
    # The counter for this execution is also provided so that clients can  
    # display it, since IPython automatically creates variables called _iN  
    # (for input prompt In[N]).  
    'execution_count' : int  
}
```

Changed in version 5.0: `pyin` is renamed to `execute_input`.

1.8.5 Execution results

Results of an execution are published as an `execute_result`. These are identical to [display_data](#) messages, with the addition of an `execution_count` key.

Results can have multiple simultaneous formats depending on its configuration. A plain text representation should always be provided in the `text/plain` mime-type. Frontends are free to display any or all of these according to its capabilities. Frontends should ignore mime-types they do not understand. The data itself is any JSON object and depends on the format. It is often, but not always a string.

Message type: `execute_result`:


```

content = {

    # The counter for this execution is also provided so that clients can
    # display it, since IPython automatically creates variables called _N
    # (for prompt N).
    'execution_count' : int,

    # data and metadata are identical to a display_data message.
    # the object being displayed is that passed to the display hook,
    # i.e. the *result* of the execution.
    'data' : dict,
    'metadata' : dict,
}

```

1.8.6 Execution errors

When an error occurs during code execution

Message type: error:

```

content = {
    # Similar content to the execute_reply messages for the 'error' case,
    # except the 'status' field is omitted.
}

```

Changed in version 5.0: pyerr renamed to error

1.8.7 Kernel status

This message type is used by frontends to monitor the status of the kernel.

Message type: status:

```

content = {
    # When the kernel starts to handle a message, it will enter the 'busy'
    # state and when it finishes, it will enter the 'idle' state.
    # The kernel will publish state 'starting' exactly once at process startup.
    execution_state : ('busy', 'idle', 'starting')
}

```

When a kernel receives a request and begins processing it, the kernel shall immediately publish a status message with `execution_state: 'busy'`. When that kernel has completed processing the request and has finished publishing associated IOPub messages, if any, it shall publish a status message with `execution_state: 'idle'`. Thus, the outputs associated with a given execution shall generally arrive between the busy and idle status messages associated with a given request.

Note: A caveat for asynchronous output

Asynchronous output (e.g. from background threads) may be produced after the kernel has sent the idle status message that signals the completion of the request. The handling of these out-of-order output messages is currently undefined in this specification, but the Jupyter Notebook continues to handle IOPub messages associated with a given request after the idle message has arrived, as long as the output area corresponding to that request is still active.

Changed in version 5.0: Busy and idle messages should be sent before/after handling every request, not just execution.

Note: Extra status messages are added between the notebook webserver and websocket clients that are not sent by the kernel. These are:

- restarting (kernel has died, but will be automatically restarted)
 - dead (kernel has died, restarting has failed)
-

1.8.8 Clear output

This message type is used to clear the output that is visible on the frontend.

Message type: `clear_output`:

```
content = {  
  
    # Wait to clear the output until new output is available. Clears the  
    # existing output immediately before the new output is displayed.  
    # Useful for creating simple animations with minimal flickering.  
    'wait' : bool,  
}
```

Changed in version 4.1: `stdout`, `stderr`, and `display` boolean keys for selective clearing are removed, and `wait` is added. The selective clearing keys are ignored in v4 and the default behavior remains the same, so v4 `clear_output` messages will be safely handled by a v4.1 frontend.

1.9 Messages on the stdin (ROUTER/DEALER) channel

With the stdin ROUTER/DEALER socket, the request/reply pattern goes in the opposite direction of most kernel communication. With the stdin socket, the kernel makes the request, and the single frontend provides the response. This pattern allows code to prompt the user for a line of input, which would normally be read from stdin in a terminal.

Many programming languages provide a function which displays a prompt, blocks until the user presses return, and returns the text they typed before pressing return. In Python 3, this is the `input()` function; in R it is called `readline()`. If the *execute_request* message has `allow_stdin==True`, kernels may implement these functions so that they send an `input_request` message and wait for a corresponding `input_reply`. The frontend is responsible for displaying the prompt and getting the user's input.

If `allow_stdin` is `False`, the kernel must not send `stdin_request`. The kernel may decide what to do instead, but it's most likely that calls to the 'prompt for input' function should fail immediately in this case.

Message type: `input_request`:

```
content = {  
    # the text to show at the prompt  
    'prompt' : str,  
    # Is the request for a password?  
    # If so, the frontend shouldn't echo input.  
    'password' : bool  
}
```

Message type: `input_reply`:

```
content = { 'value' : str }
```

When `password` is `True`, the frontend should not show the input as it is entered. Different frontends may obscure it in different ways; e.g. showing each character entered as the same neutral symbol, or not showing anything at all as the user types.

Changed in version 5.0: `password` key added.

Note: The `stdin` socket of the client is required to have the same `zmq IDENTITY` as the client's shell socket. Because of this, the `input_request` must be sent with the same `IDENTITY` routing prefix as the `execute_reply` in order for the frontend to receive the message.

Note: This pattern of requesting user input is quite different from how `stdin` works at a lower level. The Jupyter protocol does not support everything code running in a terminal can do with `stdin`, but we believe that this enables the most common use cases.

1.10 Heartbeat for kernels

Clients send ping messages on a `REQ` socket, which are echoed right back from the Kernel's `REP` socket. These are simple bytestrings, not full JSON messages described above.

1.11 Custom Messages

New in version 4.1.

Message spec 4.1 (IPython 2.0) added a messaging system for developers to add their own objects with Frontend and Kernel-side components, and allow them to communicate with each other. To do this, IPython adds a notion of a `Comm`, which exists on both sides, and can communicate in either direction.

These messages are fully symmetrical - both the Kernel and the Frontend can send each message, and no messages expect a reply. The Kernel listens for these messages on the Shell channel, and the Frontend listens for them on the `IOPub` channel.

1.11.1 Opening a Comm

Opening a `Comm` produces a `comm_open` message, to be sent to the other side:

```
{
  'comm_id' : 'u-u-i-d',
  'target_name' : 'my_comm',
  'data' : {}
}
```

Every `Comm` has an ID and a target name. The code handling the message on the receiving side is responsible for maintaining a mapping of `target_name` keys to constructors. After a `comm_open` message has been sent, there should be a corresponding `Comm` instance on both sides. The `data` key is always a dict and can be any extra JSON information used in initialization of the `comm`.

If the `target_name` key is not found on the receiving side, then it should immediately reply with a `comm_close` message to avoid an inconsistent state.

1.11.2 Comm Messages

Comm messages are one-way communications to update comm state, used for synchronizing widget state, or simply requesting actions of a comm’s counterpart.

Essentially, each comm pair defines their own message specification implemented inside the `data` dict.

There are no expected replies (of course, one side can send another `comm_msg` in reply).

Message type: `comm_msg`:

```
{
  'comm_id' : 'u-u-i-d',
  'data' : {}
}
```

1.11.3 Tearing Down Comms

Since comms live on both sides, when a comm is destroyed the other side must be notified. This is done with a `comm_close` message.

Message type: `comm_close`:

```
{
  'comm_id' : 'u-u-i-d',
  'data' : {}
}
```

1.11.4 Output Side Effects

Since comm messages can execute arbitrary user code, handlers should set the parent header and publish status busy / idle, just like an execute request.

1.12 Notes

1.12.1 `cursor_pos` and unicode offsets

Many frontends, especially those implemented in javascript, reported `cursor_pos` as the interpreter’s string index, which is not the same as the unicode character offset if the interpreter uses UTF-16 (e.g. javascript or Python 2 on macOS), which stores “astral-plane” characters such as (U+1D41A) as surrogate pairs, taking up two indices instead of one, causing a unicode offset drift of one per astral-plane character. Not all frontends have this behavior, however, and after JSON serialization information about which encoding was used when calculating the offset is lost, so assuming `cursor_pos` is calculated in UTF-16 could result in a similarly incorrect offset for frontends that did the right thing.

For this reason, in protocol versions prior to 5.2, `cursor_pos` is officially ambiguous in the presence of astral plane unicode characters. Frontends claiming to implement protocol 5.2 **MUST** identify `cursor_pos` as the encoding-independent unicode character offset. Kernels may choose to expect the UTF-16 offset from requests implementing protocol 5.1 and earlier, in order to behave correctly with the most popular frontends. But they should know that doing so *introduces* the inverse bug for the frontends that do not have this bug.

Known affected frontends (as of 2017-06):

- Jupyter Notebook < 5.1

- JupyterLab < 0.24
- nteract < 0.2.0
- CoCalc
- Jupyter Console and QtConsole with Python 2 on macOS and Windows

Known *not* affected frontends:

- QtConsole, Jupyter Console with Python 3 or Python 2 on Linux

See also:

[Discussion on GitHub](#)

Making kernels for Jupyter

A ‘kernel’ is a program that runs and introspects the user’s code. IPython includes a kernel for Python code, and people have written kernels for [several other languages](#).

When Jupyter starts a kernel, it passes it a connection file. This specifies how to set up communications with the frontend.

There are two options for writing a kernel:

1. You can reuse the IPython kernel machinery to handle the communications, and just describe how to execute your code. This is much simpler if the target language can be driven from Python. See [Making simple Python wrapper kernels](#) for details.
2. You can implement the kernel machinery in your target language. This is more work initially, but the people using your kernel might be more likely to contribute to it if it’s in the language they know.

2.1 Connection files

Your kernel will be given the path to a connection file when it starts (see [Kernel specs](#) for how to specify the command line arguments for your kernel). This file, which is accessible only to the current user, will contain a JSON dictionary looking something like this:

```
{
  "control_port": 50160,
  "shell_port": 57503,
  "transport": "tcp",
  "signature_scheme": "hmac-sha256",
  "stdin_port": 52597,
  "hb_port": 42540,
  "ip": "127.0.0.1",
  "iopub_port": 40885,
  "key": "a0436f6c-1916-498b-8eb9-e81ab9368e84"
}
```

The `transport`, `ip` and five `_port` fields specify five ports which the kernel should bind to using [ZeroMQ](#). For instance, the address of the shell socket in the example above would be:

```
tcp://127.0.0.1:57503
```

New ports are chosen at random for each kernel started.

`signature_scheme` and `key` are used to cryptographically sign messages, so that other users on the system can't send code to run in this kernel. See [The Wire Protocol](#) for the details of how this signature is calculated.

2.2 Handling messages

After reading the connection file and binding to the necessary sockets, the kernel should go into an event loop, listening on the `hb` (heartbeat), `control` and `shell` sockets.

[Heartbeat](#) messages should be echoed back immediately on the same socket - the frontend uses this to check that the kernel is still alive.

Messages on the `control` and `shell` sockets should be parsed, and their signature validated. See [The Wire Protocol](#) for how to do this.

The kernel will send messages on the `iopub` socket to display output, and on the `stdin` socket to prompt the user for textual input.

See also:

[Messaging in Jupyter](#) Details of the different sockets and the messages that come over them

[Creating Language Kernels for IPython](#) A blog post by the author of [IHaskell](#), a Haskell kernel

[simple_kernel](#) A simple example implementation of the kernel machinery in Python

2.3 Kernel specs

A kernel identifies itself to IPython by creating a directory, the name of which is used as an identifier for the kernel. These may be created in a number of locations:

	Unix	Windows
System	<code>/usr/share/jupyter/kernels</code> <code>/usr/local/share/jupyter/kernels</code>	<code>%PROGRAMDATA%\jupyter\kernels</code>
Env	<code>{sys.prefix}/share/jupyter/kernels</code>	
User	<code>~/.local/share/jupyter/kernels</code> (Linux) <code>~/Library/Jupyter/kernels</code> (Mac)	<code>%APPDATA%\jupyter\kernels</code>

The user location takes priority over the system locations, and the case of the names is ignored, so selecting kernels works the same way whether or not the filesystem is case sensitive. Since kernelspecs show up in URLs and other places, a kernelspec is required to have a simple name, only containing ASCII letters, ASCII numbers, and the simple separators: `-` hyphen, `.` period, `_` underscore.

Other locations may also be searched if the `JUPYTER_PATH` environment variable is set.

Inside the kernel directory, three types of files are presently used: `kernel.json`, `kernel.js`, and logo image files. Currently, no other files are used, but this may change in the future.

Inside the directory, the most important file is `kernel.json`. This should be a JSON serialised dictionary containing the following keys and values:

- **argv**: A list of command line arguments used to start the kernel. The text `{connection_file}` in any argument will be replaced with the path to the connection file.
- **display_name**: The kernel's name as it should be displayed in the UI. Unlike the kernel name used in the API, this can contain arbitrary unicode characters.
- **language**: The name of the language of the kernel. When loading notebooks, if no matching kernelspec key (may differ across machines) is found, a kernel with a matching *language* will be used. This allows a notebook written on any Python or Julia kernel to be properly associated with the user's Python or Julia kernel, even if they aren't listed under the same name as the author's.
- **interrupt_mode** (optional): May be either `signal` or `message` and specifies how a client is supposed to interrupt cell execution on this kernel, either by sending an interrupt signal via the operating system's signalling facilities (e.g. *SIGINT* on POSIX systems), or by sending an `interrupt_request` message on the control channel (see *Kernel interrupt*). If this is not specified the client will default to `signal` mode.
- **env** (optional): A dictionary of environment variables to set for the kernel. These will be added to the current environment variables before the kernel is started.
- **metadata** (optional): A dictionary of additional attributes about this kernel; used by clients to aid clients in kernel selection. Metadata added here should be namespaced for the tool reading and writing that metadata.

For example, the kernel.json file for IPython looks like this:

```
{
  "argv": ["python3", "-m", "IPython.kernel",
           "-f", "{connection_file}"],
  "display_name": "Python 3",
  "language": "python"
}
```

To see the available kernel specs, run:

```
jupyter kernelspec list
```

To start the terminal console or the Qt console with a specific kernel:

```
jupyter console --kernel bash
jupyter qtconsole --kernel bash
```

The notebook offers you the available kernels in a dropdown menu from the 'New' button.

Making simple Python wrapper kernels

You can re-use IPython's kernel machinery to easily make new kernels. This is useful for languages that have Python bindings, such as [Octave](#) (via [Oct2Py](#)), or languages where the REPL can be controlled in a tty using [pexpect](#), such as `bash`.

See also:

[bash_kernel](#) A simple kernel for `bash`, written using this machinery

3.1 Required steps

Subclass `ipykernel.kernelbase.Kernel`, and implement the following methods and attributes:

class MyKernel

implementation

implementation_version

banner

Information for [Kernel info](#) replies. 'Implementation' refers to the kernel (e.g. IPython), rather than the language (e.g. Python). The 'banner' is displayed to the user in console UIs before the first prompt. All of these values are strings.

language_info

Language information for [Kernel info](#) replies, in a dictionary. This should contain the key `mimetype` with the mimetype of code in the target language (e.g. `'text/x-python'`), the name of the language being implemented (e.g. `'python'`), and `file_extension` (e.g. `'.py'`). It may also contain keys `codemirror_mode` and `pygments_lexer` if they need to differ from `language`.

Other keys may be added to this later.

do_execute (*code, silent, store_history=True, user_expressions=None, allow_stdin=False*)

Execute user code.

Parameters

- **code** (*str*) – The code to be executed.
- **silent** (*bool*) – Whether to display output.
- **store_history** (*bool*) – Whether to record this code in history and increase the execution count. If silent is True, this is implicitly False.
- **user_expressions** (*dict*) – Mapping of names to expressions to evaluate after the code has run. You can ignore this if you need to.
- **allow_stdin** (*bool*) – Whether the frontend can provide input on request (e.g. for Python's `raw_input()`).

Your method should return a dict containing the fields described in *Execution results*. To display output, it can send messages using `send_response()`. See *Messaging in Jupyter* for details of the different message types.

To launch your kernel, add this at the end of your module:

```
if __name__ == '__main__':
    from ipykernel.kernelapp import IPKernelApp
    IPKernelApp.launch_instance(kernel_class=MyKernel)
```

Now create a JSON *kernel spec* file and install it using `jupyter kernelspec install </path/to/kernel>`. Place your kernel module anywhere Python can import it (try current directory for testing). Finally, you can run your kernel using `jupyter console --kernel <mykernelname>`. Note that `<mykernelname>` in the below example is `echo`.

3.2 Example

See also:

echo_kernel A packaged, installable version of the condensed example below.

`echokernel.py` will simply echo any input it's given to stdout:

```
from ipykernel.kernelbase import Kernel

class EchoKernel(Kernel):
    implementation = 'Echo'
    implementation_version = '1.0'
    language = 'no-op'
    language_version = '0.1'
    language_info = {
        'name': 'Any text',
        'mimetype': 'text/plain',
        'file_extension': '.txt',
    }
    banner = "Echo kernel - as useful as a parrot"

    def do_execute(self, code, silent, store_history=True, user_expressions=None,
                  allow_stdin=False):
        if not silent:
            stream_content = {'name': 'stdout', 'text': code}
            self.send_response(self.iopub_socket, 'stream', stream_content)

        return {'status': 'ok',
                # The base class increments the execution count
```

(continues on next page)

(continued from previous page)

```

        'execution_count': self.execution_count,
        'payload': [],
        'user_expressions': {},
    }

if __name__ == '__main__':
    from ipykernel.kernelapp import IPKernelApp
    IPKernelApp.launch_instance(kernel_class=EchoKernel)

```

Here's the Kernel spec `kernel.json` file for this:

```

{
  "argv": ["python", "-m", "echokernel", "-f", "{connection_file}"],
  "display_name": "Echo"
}

```

3.3 Optional steps

You can override a number of other methods to improve the functionality of your kernel. All of these methods should return a dictionary as described in the relevant section of the [messaging spec](#).

class MyKernel

do_complete (*code*, *cursor_pos*)

Code completion

Parameters

- **code** (*str*) – The code already present
- **cursor_pos** (*int*) – The position in the code where completion is requested

See also:

[Completion](#) messages

do_inspect (*code*, *cursor_pos*, *detail_level=0*)

Object introspection

Parameters

- **code** (*str*) – The code
- **cursor_pos** (*int*) – The position in the code where introspection is requested
- **detail_level** (*int*) – 0 or 1 for more or less detail. In IPython, 1 gets the source code.

See also:

[Introspection](#) messages

do_history (*hist_access_type*, *output*, *raw*, *session=None*, *start=None*, *stop=None*, *n=None*, *pattern=None*, *unique=False*)

History access. Only the relevant parameters for the type of history request concerned will be passed, so your method definition must have defaults for all the arguments shown with defaults here.

See also:

[History](#) messages

do_is_complete (*code*)

Is code entered in a console-like interface complete and ready to execute, or should a continuation prompt be shown?

Parameters **code** (*str*) – The code entered so far - possibly multiple lines

See also:

Code completeness messages

do_shutdown (*restart*)

Shutdown the kernel. You only need to handle your own clean up - the kernel machinery will take care of cleaning up its own things before stopping.

Parameters **restart** (*bool*) – Whether the kernel will be started again afterwards

See also:

Kernel shutdown messages

4.1 kernelspec - discovering kernels

See also:

Kernel specs

class `jupyter_client.kernelspec.KernelSpec`

argv

The list of arguments to start this kernel.

env

A dictionary of extra environment variables to declare, in addition to the current environment variables, when launching this kernel.

display_name

The name to display for this kernel in UI.

language

The name of the language the kernel implements, to help with picking appropriate kernels when loading notebooks.

metadata

Additional kernel-specific metadata; clients can use this as needed, for instance to aid in kernel selection and filtering.

Metadata added here should be namespaced for the tool reading and writing that metadata. Concretely, if you're adding a new field called `supported_versions` which your tool recognizes, then you should add it as `"mytool": {"supported_versions": [1, 2]}`, **not** as a top-level field called `supported_versions`.

resource_dir

The path to the directory with this kernel's resources, such as icons.

to_json()
Serialise this kernelspec to a JSON object.
Returns a string.

class jupyter_client.kernelspec.**KernelSpecManager**

find_kernel_specs()
Returns a dict mapping kernel names to resource directories.

get_all_specs()
Returns a dict mapping kernel names to kernelspecs.
Returns a dict of the form:

```
{
  'kernel_name': {
    'resource_dir': '/path/to/kernel_name',
    'spec': {"the spec itself": ...}
  },
  ...
}
```

get_kernel_spec(kernel_name)
Returns a *KernelSpec* instance for the given kernel_name.
Raises *NoSuchKernel* if the given kernel name is not found.

install_kernel_spec(source_dir, kernel_name=None, user=False, replace=None, prefix=None)
Install a kernel spec by copying its directory.

If kernel_name is not given, the basename of source_dir will be used.

If user is False, it will attempt to install into the systemwide kernel registry. If the process does not have appropriate permissions, an *OSError* will be raised.

If prefix is given, the kernelspec will be installed to PREFIX/share/jupyter/kernels/KERNEL_NAME.
This can be sys.prefix for installation inside virtual or conda envs.

exception jupyter_client.kernelspec.**NoSuchKernel**

name
The name of the kernel which was requested.

```
jupyter_client.kernelspec.find_kernel_specs()
jupyter_client.kernelspec.get_kernel_spec(kernel_name)
jupyter_client.kernelspec.install_kernel_spec(source_dir, kernel_name=None,
                                              user=False, replace=False)
```

These methods from *KernelSpecManager* are exposed as functions on the module as well; they will use all the default settings.

4.2 manager - starting, stopping, signalling

class jupyter_client.**KernelManager** (**kwargs)
Manages a single kernel in a subprocess on this host.

This version starts kernels with Popen.

kernel_name

The name of the kernel to launch (see *Kernel specs*).

start_kernel (***kw*)

Starts a kernel on this host in a separate process.

If random ports (port=0) are being used, this method must be called before the channels are created.

Parameters ****kw** (*optional*) – keyword arguments that are passed down to build the kernel_cmd and launching the kernel (e.g. Popen kwargs).

kernel

Once the kernel has been started, this is the `subprocess.Popen` class for the kernel process.

is_alive ()

Is the kernel process still running?

interrupt_kernel ()

Interrupts the kernel by sending it a signal.

Unlike `signal_kernel`, this operation is well supported on all platforms.

signal_kernel (*signum*)

Sends a signal to the process group of the kernel (this usually includes the kernel and any subprocesses spawned by the kernel).

Note that since only SIGTERM is supported on Windows, this function is only useful on Unix systems.

client (***kwargs*)

Create a client configured to connect to our kernel

For the client API, see `jupyter_client.client`.

blocking_client ()

Make a blocking client connected to my kernel

shutdown_kernel (*now=False, restart=False*)

Attempts to stop the kernel process cleanly.

This attempts to shutdown the kernels cleanly by:

1. Sending it a shutdown message over the shell channel.
2. If that fails, the kernel is shutdown forcibly by sending it a signal.

Parameters

- **now** (*bool*) – Should the kernel be forcible killed *now*. This skips the first, nice shutdown attempt.
- **restart** (*bool*) – Will this kernel be restarted after it is shutdown. When this is True, connection files will not be cleaned up.

restart_kernel (*now=False, newports=False, **kw*)

Restarts a kernel with the arguments that were used to launch it.

Parameters

- **now** (*bool, optional*) – If True, the kernel is forcefully restarted *immediately*, without having a chance to do any cleanup action. Otherwise the kernel is given 1s to clean up before a forceful restart is issued.

In all cases the kernel is restarted, the only difference is whether it is given a chance to perform a clean shutdown or not.

- **newports** (*bool, optional*) – If the old kernel was launched with random ports, this flag decides whether the same ports and connection file will be used again. If False, the same ports and connection file are used. This is the default. If True, new random port numbers are chosen and a new connection file is written. It is still possible that the newly chosen random port numbers happen to be the same as the old ones.
- ****kw** (*optional*) – Any options specified here will overwrite those used to launch the kernel.

4.2.1 multikernelmanager - controlling multiple kernels

class `jupyter_client.MultiKernelManager` (***kwargs*)

A class for managing multiple kernels.

This exposes the same methods as `KernelManager`, but their first parameter is a kernel ID, a string identifying the kernel instance. Typically these are UUIDs picked by `start_kernel()`

start_kernel (*kernel_name=None, **kwargs*)

Start a new kernel.

The caller can pick a `kernel_id` by passing one in as a keyword arg, otherwise one will be picked using a uuid.

The kernel ID for the newly started kernel is returned.

list_kernel_ids ()

Return a list of the kernel ids of the active kernels.

get_kernel (*kernel_id*)

Get the single `KernelManager` object for a kernel by its uuid.

Parameters `kernel_id` (*uuid*) – The id of the kernel.

remove_kernel (*kernel_id*)

remove a kernel from our mapping.

Mainly so that a kernel can be removed if it is already dead, without having to call `shutdown_kernel`.

The kernel object is returned.

shutdown_all (*now=False*)

Shutdown all kernels.

4.2.2 Utility functions

`jupyter_client.run_kernel` (***kwargs*)

Context manager to create a kernel in a subprocess.

The kernel is shut down when the context exits.

Returns `kernel_client`

Return type connected `KernelClient` instance

4.3 client - communicating with kernels

See also:

Messaging in Jupyter The Jupyter messaging specification

class `jupyter_client.KernelClient` (***kwargs*)
 Communicates with a single kernel on any host via zmq channels.

There are four channels associated with each kernel:

- `shell`: for request/reply calls to the kernel.
- `iopub`: for the kernel to publish results to frontends.
- `hb`: for monitoring the kernel's heartbeat.
- `stdin`: for frontends to reply to `raw_input` calls in the kernel.

The messages that can be sent on these channels are exposed as methods of the client (`KernelClient.execute`, `complete`, `history`, etc.). These methods only send the message, they don't wait for a reply. To get results, use e.g. `get_shell_msg()` to fetch messages from the shell channel.

load_connection_file (*connection_file=None*)
 Load connection info from JSON dict in `self.connection_file`.

Parameters `connection_file` (*unicode, optional*) – Path to connection file to load.
 If unspecified, use `self.connection_file`

load_connection_info (*info*)
 Load connection info from a dict containing connection info.

Typically this data comes from a connection file and is called by `load_connection_file`.

Parameters `info` (*dict*) – Dictionary containing `connection_info`. See the `connection_file` spec for details.

start_channels (*shell=True, iopub=True, stdin=True, hb=True*)
 Starts the channels for this kernel.

This will create the channels if they do not exist and then start them (their activity runs in a thread). If port numbers of 0 are being used (random ports) then you must first call `start_kernel()`. If the channels have been stopped and you call this, `RuntimeError` will be raised.

execute (*code, silent=False, store_history=True, user_expressions=None, allow_stdin=None, stop_on_error=True*)
 Execute code in the kernel.

Parameters

- **code** (*str*) – A string of code in the kernel's language.
- **silent** (*bool, optional (default False)*) – If set, the kernel will execute the code as quietly possible, and will force `store_history` to be `False`.
- **store_history** (*bool, optional (default True)*) – If set, the kernel will store command history. This is forced to be `False` if `silent` is `True`.
- **user_expressions** (*dict, optional*) – A dict mapping names to expressions to be evaluated in the user's dict. The expression values are returned as strings formatted using `repr()`.
- **allow_stdin** (*bool, optional (default self.allow_stdin)*) – Flag for whether the kernel can send `stdin` requests to frontends.

Some frontends (e.g. the Notebook) do not support `stdin` requests. If `raw_input` is called from code executed from such a frontend, a `StdinNotImplementedError` will be raised.

- **stop_on_error** (*bool, optional (default True)*) – Flag whether to abort the execution queue, if an exception is encountered.

Returns

Return type The msg_id of the message sent.

complete (*code*, *cursor_pos=None*)

Tab complete text in the kernel's namespace.

Parameters

- **code** (*str*) – The context in which completion is requested. Can be anything between a variable name and an entire cell.
- **cursor_pos** (*int*, *optional*) – The position of the cursor in the block of code where the completion was requested. Default: `len (code)`

Returns

Return type The msg_id of the message sent.

inspect (*code*, *cursor_pos=None*, *detail_level=0*)

Get metadata information about an object in the kernel's namespace.

It is up to the kernel to determine the appropriate object to inspect.

Parameters

- **code** (*str*) – The context in which info is requested. Can be anything between a variable name and an entire cell.
- **cursor_pos** (*int*, *optional*) – The position of the cursor in the block of code where the info was requested. Default: `len (code)`
- **detail_level** (*int*, *optional*) – The level of detail for the introspection (0-2)

Returns

Return type The msg_id of the message sent.

history (*raw=True*, *output=False*, *hist_access_type='range'*, ***kwargs*)

Get entries from the kernel's history list.

Parameters

- **raw** (*bool*) – If True, return the raw input.
- **output** (*bool*) – If True, then return the output as well.
- **hist_access_type** (*str*) –
 'range' (fill in session, start and stop params), 'tail' (fill in n) or 'search' (fill in pattern param).
- **session** (*int*) – For a range request, the session from which to get lines. Session numbers are positive integers; negative ones count back from the current session.
- **start** (*int*) – The first line number of a history range.
- **stop** (*int*) – The final (excluded) line number of a history range.
- **n** (*int*) – The number of lines of history to get for a tail request.
- **pattern** (*str*) – The glob-syntax pattern for a search request.

Returns

Return type The ID of the message sent.

comm_info (*target_name=None*)

Request comm info

Returns

Return type The `msg_id` of the message sent

is_complete (*code*)

Ask the kernel whether some code is complete and ready to execute.

input (*string*)

Send a string of raw input to the kernel.

This should only be called in response to the kernel sending an `input_request` message on the `stdin` channel.

shutdown (*restart=False*)

Request an immediate kernel shutdown.

Upon receipt of the (empty) reply, client code can safely assume that the kernel has shut down and it's safe to forcefully terminate it if it's still alive.

The kernel will send the reply via a function registered with Python's `atexit` module, ensuring it's truly done as the kernel is done with all normal operation.

Returns

Return type The `msg_id` of the message sent

get_shell_msg (**args, **kwargs*)

Get a message from the shell channel

get_iopub_msg (**args, **kwargs*)

Get a message from the iopub channel

get_stdin_msg (**args, **kwargs*)

Get a message from the stdin channel

class `jupyter_client.BlockingKernelClient` (***kwargs*)

A `KernelClient` with blocking APIs

`get_[channel]_msg()` methods wait for and return messages on channels, raising `queue.Empty` if no message arrives within `timeout` seconds.

execute_interactive (*code, silent=False, store_history=True, user_expressions=None, allow_stdin=None, stop_on_error=True, timeout=None, output_hook=None, stdin_hook=None*)

Execute code in the kernel interactively

Output will be redisplayed, and stdin prompts will be relayed as well. If an IPython kernel is detected, rich output will be displayed.

You can pass a custom `output_hook` callable that will be called with every IOPub message that is produced instead of the default redisplay.

New in version 5.0.

Parameters

- **code** (*str*) – A string of code in the kernel's language.
- **silent** (*bool, optional (default False)*) – If set, the kernel will execute the code as quietly possible, and will force `store_history` to be `False`.
- **store_history** (*bool, optional (default True)*) – If set, the kernel will store command history. This is forced to be `False` if `silent` is `True`.

- **user_expressions** (*dict, optional*) – A dict mapping names to expressions to be evaluated in the user’s dict. The expression values are returned as strings formatted using `repr()`.
- **allow_stdin** (*bool, optional (default self.allow_stdin)*) – Flag for whether the kernel can send stdin requests to frontends.

Some frontends (e.g. the Notebook) do not support stdin requests. If `raw_input` is called from code executed from such a frontend, a `StdinNotImplementedError` will be raised.

- **stop_on_error** (*bool, optional (default True)*) – Flag whether to abort the execution queue, if an exception is encountered.
- **timeout** (*float or None (default: None)*) – Timeout to use when waiting for a reply
- **output_hook** (*callable(msg)*) – Function to be called with output messages. If not specified, output will be redisplayed.
- **stdin_hook** (*callable(msg)*) – Function to be called with `stdin_request` messages. If not specified, `input/getpass` will be called.

Returns `reply` – The reply message for this request

Return type `dict`

Changes in Jupyter Client

5.1 5.2.3

5.2.3 on GitHub

- Fix hang on close in `ThreadedKernelClient` (used in `QtConsole`) when using tornado with asyncio (default behavior of tornado 5, see [PR #352](#)).
- Fix errors when using deprecated `KernelManager.kernel_cmd` ([PR #343](#), [PR #344](#)).

5.2 5.2.2

5.2.2 on GitHub

- Fix `KernelSpecManager.get_all_specs()` method in subclasses that only override `KernelSpecManager.find_kernel_specs()` and `KernelSpecManager.get_kernel_spec()`. See [#338](#) and [PR #339](#).
- Eliminate occasional error messages during process exit ([PR #336](#)).
- Improve error message when attempting to bind on invalid address ([PR #330](#)).
- Add missing direct dependency on tornado ([PR #323](#)).

5.3 5.2.1

5.2.1 on GitHub

- Add parenthesis to conditional pytest requirement to work around a bug in the `wheel` package, that generate a `.whl` which otherwise always depends on `pytest` see [#324](#) and [PR #325](#).

5.4 5.2

5.2 on GitHub

- Define Jupyter protocol version 5.3:
 - Kernels can now opt to be interrupted by a message sent on the control channel instead of a system signal. See *Kernel specs* and *Kernel interrupt* (PR #294).
- New `jupyter kernel` command to launch an installed kernel by name (PR #240).
- Kernelspecs where the command starts with e.g. `python3` or `python3.6`—matching the version `jupyter_client` is running on—are now launched with the same Python executable as the launching process (PR #306). This extends the special handling of `python` added in 5.0.
- Command line arguments specified by a kernelspec can now include `{resource_dir}`, which will be substituted with the kernelspec resource directory path when the kernel is launched (PR #289).
- Kernelspecs now have an optional `metadata` field to hold arbitrary metadata about kernels—see *Kernel specs* (PR #274).
- Make the `KernelRestarter` class used by a `KernelManager` configurable (PR #290).
- When killing a kernel on Unix, kill its process group (PR #314).
- If a kernel dies soon after starting, reassign random ports before restarting it, in case one of the previously chosen ports has been bound by another process (PR #279).
- Avoid unnecessary filesystem operations when finding a kernelspec with `KernelSpecManager.get_kernel_spec()` (PR #311).
- `KernelSpecManager.get_all_specs()` will no longer raise an exception on encountering an invalid `kernel.json` file. It will raise a warning and continue (PR #310).
- Check for non-contiguous buffers before trying to send them through ZMQ (PR #258).
- Compatibility with upcoming Tornado version 5.0 (PR #304).
- Simplify setup code by always using `setuptools` (PR #284).
- Soften warnings when setting the sticky bit on runtime files fails (PR #286).
- Various corrections and improvements to documentation.

5.5 5.1

5.1 on GitHub

- Define Jupyter protocol version 5.2, resolving ambiguity of `cursor_pos` field in the presence of unicode surrogate pairs.
See also:
cursor_pos and unicode offsets
- Add `Session.clone()` for making a copy of a `Session` object without sharing the digest history. Reusing a single `Session` object to connect multiple sockets to the same `IOPub` peer can cause digest collisions.
- Avoid global references preventing garbage collection of background threads.

5.6 5.0

5.6.1 5.0.1

[5.0.1 on GitHub](#)

- Update internal protocol version number to 5.1, which should have been done in 5.0.0.

5.6.2 5.0.0

[5.0.0 on GitHub](#)

New features:

- Implement Jupyter protocol version 5.1.
- Introduce **jupyter run** command for running scripts with a kernel, for instance:

```
jupyter run --kernel python3 myscript.py
```

- New method `BlockingKernelClient.execute_interactive()` for running code and capturing or redisplaying its output.
- New `KernelManager.shutdown_wait_time` configurable for adjusting the time for a kernel manager to wait after politely requesting shutdown before it resorts to forceful termination.

Fixes:

- Set sticky bit on connection-file directory to avoid getting cleaned up.
- `jupyter_client.launcher.launch_kernel()` passes through additional options to the underlying Popen, matching `KernelManager.start_kernel()`.
- Check types of `buffers` argument in `Session.send()`, so that `TypeError`s are raised immediately, rather than in the eventloop.

Changes:

- In kernelspecs, if the executable is the string `python` (as opposed to an absolute path), `sys.executable` will be used rather than resolving `python` on `PATH`. This should enable Python-based kernels to install kernelspecs as part of wheels.
- kernelspec names are now validated. They should only include ascii letters and numbers, plus period, hyphen, and underscore.

Backward-incompatible changes:

- `datetime` objects returned in parsed messages are now always timezone-aware. Timestamps in messages without timezone info are interpreted as the local timezone, as this was the behavior in earlier versions.

5.7 4.4

5.7.1 4.4.0

[4.4 on GitHub](#)

- Add `KernelClient.load_connection_info()` on `KernelClient`, etc. for loading connection info directly from a dict, not just from files.

- Include parent headers when adapting messages from older protocol implementations (treats parent headers the same as headers).
- Compatibility fixes in tests for recent changes in ipykernel.

5.8 4.3

5.8.1 4.3.0

[4.3 on GitHub](#)

- Adds `--sys-prefix` argument to `jupyter kernelspec install`, for better symmetry with `jupyter nbextension install`, etc.

5.9 4.2

5.9.1 4.2.2

[4.2.2 on GitHub](#)

- Another fix for the `start_new_kernel()` issue in 4.2.1 affecting slow-starting kernels.

5.9.2 4.2.1

[4.2.1 on GitHub](#)

- Fix regression in 4.2 causing `start_new_kernel()` to fail while waiting for kernels to become available.

5.9.3 4.2.0

[4.2.0 on GitHub](#)

- added `jupyter kernelspec remove` for removing kernelspecs
- allow specifying the environment for kernel processes via the `env` argument
- added `name` field to connection files identifying the kernelspec name, so that consumers of connection files (alternate frontends) can identify the kernelspec in use
- added `KernelSpecManager.get_all_specs()` for getting all kernelspecs more efficiently
- various improvements to error messages and documentation

5.10 4.1

5.10.1 4.1.0

[4.1.0 on GitHub](#)

Highlights:

- Setuptools fixes for `jupyter kernelspec`

- `jupyter kernelspec list` includes paths
- add `KernelManager.blocking_client()`
- provisional implementation of `comm_info` requests from upcoming 5.1 release of the protocol

5.11 4.0

The first release of Jupyter Client as its own package.

CHAPTER 6

Indices and tables

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

j

`jupyter_client`, [35](#)

`jupyter_client.kernelspec`, [35](#)

A

argv (jupyter_client.kernelspec.KernelSpec attribute), 35

B

banner (MyKernel attribute), 31

blocking_client() (jupyter_client.KernelManager method), 37

BlockingKernelClient (class in jupyter_client), 41

C

client() (jupyter_client.KernelManager method), 37

comm_info() (jupyter_client.KernelClient method), 40

complete() (jupyter_client.KernelClient method), 40

D

display_name (jupyter_client.kernelspec.KernelSpec attribute), 35

do_complete() (MyKernel method), 33

do_execute() (MyKernel method), 31

do_history() (MyKernel method), 33

do_inspect() (MyKernel method), 33

do_is_complete() (MyKernel method), 33

do_shutdown() (MyKernel method), 34

E

env (jupyter_client.kernelspec.KernelSpec attribute), 35

environment variable
JUPYTER_PATH, 28

execute() (jupyter_client.KernelClient method), 39

execute_interactive() (jupyter_client.BlockingKernelClient method), 41

F

find_kernel_specs() (in module
jupyter_client.kernelspec), 36

find_kernel_specs() (jupyter_client.kernelspec.KernelSpecManager method), 36

G

get_all_specs() (jupyter_client.kernelspec.KernelSpecManager method), 36

get_iopub_msg() (jupyter_client.KernelClient method), 41

get_kernel() (jupyter_client.MultiKernelManager method), 38

get_kernel_spec() (in module jupyter_client.kernelspec), 36

get_kernel_spec() (jupyter_client.kernelspec.KernelSpecManager method), 36

get_shell_msg() (jupyter_client.KernelClient method), 41

get_stdin_msg() (jupyter_client.KernelClient method), 41

H

history() (jupyter_client.KernelClient method), 40

I

implementation (MyKernel attribute), 31

implementation_version (MyKernel attribute), 31

input() (jupyter_client.KernelClient method), 41

inspect() (jupyter_client.KernelClient method), 40

install_kernel_spec() (in module
jupyter_client.kernelspec), 36

install_kernel_spec() (jupyter_client.kernelspec.KernelSpecManager method), 36

interrupt_kernel() (jupyter_client.KernelManager method), 37

is_alive() (jupyter_client.KernelManager method), 37

is_complete() (jupyter_client.KernelClient method), 41

J

jupyter_client (module), 35

jupyter_client.kernelspec (module), 35

JUPYTER_PATH, 28

K

kernel (jupyter_client.KernelManager attribute), 37

kernel_name (jupyter_client.KernelManager attribute),
36
 KernelClient (class in jupyter_client), 39
 KernelManager (class in jupyter_client), 36
 KernelSpec (class in jupyter_client.kernelspec), 35
 KernelSpecManager (class in jupyter_client.kernelspec),
36

L

language (jupyter_client.kernelspec.KernelSpec attribute), 35
 language_info (MyKernel attribute), 31
 list_kernel_ids() (jupyter_client.MultiKernelManager method), 38
 load_connection_file() (jupyter_client.KernelClient method), 39
 load_connection_info() (jupyter_client.KernelClient method), 39

M

metadata (jupyter_client.kernelspec.KernelSpec attribute), 35
 MultiKernelManager (class in jupyter_client), 38
 MyKernel (built-in class), 31, 33

N

name (jupyter_client.kernelspec.NoSuchKernel attribute), 36
 NoSuchKernel, 36

R

remove_kernel() (jupyter_client.MultiKernelManager method), 38
 resource_dir (jupyter_client.kernelspec.KernelSpec attribute), 35
 restart_kernel() (jupyter_client.KernelManager method),
37
 run_kernel() (in module jupyter_client), 38

S

shutdown() (jupyter_client.KernelClient method), 41
 shutdown_all() (jupyter_client.MultiKernelManager method), 38
 shutdown_kernel() (jupyter_client.KernelManager method), 37
 signal_kernel() (jupyter_client.KernelManager method),
37
 start_channels() (jupyter_client.KernelClient method), 39
 start_kernel() (jupyter_client.KernelManager method), 37
 start_kernel() (jupyter_client.MultiKernelManager method), 38

T

to_json() (jupyter_client.kernelspec.KernelSpec method),
35