nbformat Documentation

Release 4.3

Jupyter Development Team

Mar 29, 2019

Contents

Jupyter (né IPython) notebook files are simple JSON documents, containing text, source code, rich media output, and metadata. Each segment of the document is stored in a cell.

Contents:

The Notebook file format

Some general points about the notebook format:

Note: *All* metadata fields are optional. While the type and values of some metadata are defined, no metadata values are required to be defined.

1.1 Top-level structure

At the highest level, a Jupyter notebook is a dictionary with a few keys:

- metadata (dict)
- nbformat (int)
- nbformat_minor (int)
- cells (list)

```
{
  "metadata" : {
    "kernel_info": {
        # if kernel_info is defined, its name field is required.
        "name" : "the name of the kernel"
    },
    "language_info": {
        # if language_info is defined, its name field is required.
        "name" : "the programming language of the kernel",
        "version": "the version of the language",
        "codemirror_mode": "The name of the codemirror mode to use [optional]"
    }
  },
  "nbformat": 4,
  "nbformat_minor": 0,
```

```
    "cells" : [
        # list of cell dictionaries, see below
    ],
}
```

Some fields, such as code input and text output, are characteristically multi-line strings. When these fields are written to disk, they **may** be written as a list of strings, which should be joined with `''` when reading back into memory. In programmatic APIs for working with notebooks (Python, Javascript), these are always re-joined into the original multi-line string. If you intend to work with notebook files directly, you must allow multi-line string fields to be either a string or list of strings.

1.2 Cell Types

There are a few basic cell types for encapsulating code and text. All cells have the following basic structure:

```
{
  "cell_type" : "name",
  "metadata" : {},
  "source" : "single string or [list, of, strings]",
}
```

Note: On disk, multi-line strings **MAY** be split into lists of strings. When read with the nbformat Python API, these multi-line strings will always be a single string.

1.2.1 Markdown cells

Markdown cells are used for body-text, and contain markdown, as defined in GitHub-flavored markdown, and implemented in marked.

```
{
  "cell_type" : "markdown",
  "metadata" : {},
  "source" : "[multi-line *markdown*]",
}
```

Changed in version nbformat: 4.0

Heading cells have been removed, in favor of simple headings in markdown.

1.2.2 Code cells

Code cells are the primary content of Jupyter notebooks. They contain source code in the language of the document's associated kernel, and a list of outputs associated with executing that code. They also have an execution_count, which must be an integer or `null`.

```
{
  "cell_type" : "code",
  "execution_count": 1, # integer or null
  "metadata" : {
```

```
        "collapsed" : True, # whether the output of the cell is collapsed
        "autoscroll": False, # any of true, false or "auto"
    },
    "source" : "[some multi-line code]",
    "outputs": [{
        # list of output dicts (described below)
        "output_type": "stream",
        ...
    }],
}
```

Changed in version nbformat: 4.0

`input` was renamed to `source`, for consistency among cell types.

Changed in version nbformat: 4.0

`prompt_number` renamed to `execution_count`

1.2.3 Code cell outputs

A code cell can have a variety of outputs (stream data or rich mime-type output). These correspond to messages produced as a result of executing the cell.

All outputs have an `output_type` field, which is a string defining what type of output it is.

stream output

```
{
    "output_type" : "stream",
    "name" : "stdout", # or stderr
    "text" : "[multiline stream text]",
}
```

Changed in version nbformat: 4.0

The `stream` key was changed to `name` to match the stream message.

display_data

Rich display outputs, as created by `display_data` messages, contain data keyed by mime-type. This is often called a mime-bundle, and shows up in various locations in the notebook format and message spec. The metadata of these messages may be keyed by mime-type as well.

```
{
    "output_type" : "display_data",
    "data" : {
        "text/plain" : "[multiline text data]",
        "image/png": "[base64-encoded-multiline-png-data]",
        "application/json": {
            # JSON data is included as-is
            "json": "data",
        },
    },
```

```
    "metadata" : {
      "image/png": {
        "width": 640,
        "height": 480,
      },
    },
}
```

Changed in version nbformat: 4.0

`application/json` output is no longer double-serialized into a string.

Changed in version nbformat: 4.0

mime-types are used for keys, instead of a combination of short names (`text`) and mime-types, and are stored in a `data` key, rather than the top-level. i.e. `output.data['image/png']` instead of `output.png`.

execute_result

Results of executing a cell (as created by `displayhook` in Python) are stored in `execute_result` outputs. *execute_result* outputs are identical to `display_data`, adding only a `execution_count` field, which must be an integer.

```
{
  "output_type" : "execute_result",
  "execution_count": 42,
  "data" : {
    "text/plain" : "[multiline text data]",
    "image/png": "[base64-encoded-multiline-png-data]",
    "application/json": {
      # JSON data is included as-is
      "json": "data",
    },
  },
  "metadata" : {
    "image/png": {
      "width": 640,
      "height": 480,
    },
  },
}
```

Changed in version nbformat: 4.0

`pyout` renamed to `execute_result`

Changed in version nbformat: 4.0

`prompt_number` renamed to `execution_count`

error

Failed execution may show a traceback

```
{
      'output_type': 'error',
```

```
  'ename' : str,    # Exception name, as a string
  'evalue' : str,   # Exception value, as a string

  # The traceback will contain a list of frames,
  # represented each as a string.
  'traceback' : list,
}
```

Changed in version nbformat: 4.0

`pyerr` renamed to `error`

1.2.4 Raw NBConvert cells

A raw cell is defined as content that should be included *unmodified* in nbconvert output. For example, this cell could include raw LaTeX for nbconvert to pdf via latex, or restructured text for use in Sphinx documentation.

The notebook authoring environment does not render raw cells.

The only logic in a raw cell is the *format* metadata field. If defined, it specifies which nbconvert output format is the intended target for the raw cell. When outputting to any other format, the raw cell's contents will be excluded. In the default case when this value is undefined, a raw cell's contents will be included in any nbconvert output, regardless of format.

```
{
  "cell_type" : "raw",
  "metadata" : {
    # the mime-type of the target nbconvert format.
    # nbconvert to formats other than this will exclude this cell.
    "format" : "mime/type"
  },
  "source" : "[some nbformat output text]"
}
```

1.2.5 Cell attachments

New in version 4.1.

Markdown and raw cells can have a number of attachments, typically inline images that can be referenced in the markdown content of a cell. The `attachments` dictionary of a cell contains a set of mime-bundles (see display_data) keyed by filename that represents the files attached to the cell.

Note: The `attachments` dictionary is an optional field and can be undefined or empty if the cell does not have any attachments.

```
{
  "cell_type" : "markdown",
  "metadata" : {},
  "source" : ["Here is an *inline* image ![inline image](attachment:test.png)"],
  "attachments" : {
    "test.png": {
        "image/png" : "base64-encoded-png-data"
    }
```

```
    }
}
```

1.3 Backward-compatible changes

The notebook format is an evolving format. When backward-compatible changes are made, the notebook format minor version is incremented. When backward-incompatible changes are made, the major version is incremented.

As of nbformat 4.x, backward-compatible changes include:

- new fields in any dictionary (notebook, cell, output, metadata, etc.)
- new cell types
- new output types

New cell or output types will not be rendered in versions that do not recognize them, but they will be preserved.

1.4 Metadata

Metadata is a place that you can put arbitrary JSONable information about your notebook, cell, or output. Because it is a shared namespace, any custom metadata should use a sufficiently unique namespace, such as *metadata.kaylees_md.foo = "bar"*.

Metadata fields officially defined for Jupyter notebooks are listed here:

1.4.1 Notebook metadata

The following metadata keys are defined at the notebook level:

Key	Value	Interpretation
kernelspec	dict	A kernel specification
authors	list of dicts	A list of authors of the document

A notebook's authors is a list of dictionaries containing information about each author of the notebook. Currently, only the name is required. Additional fields may be added.

```
nb.metadata.authors = [
    {
        'name': 'Fernando Perez',
    },
    {
        'name': 'Brian Granger',
    },
]
```

1.4.2 Cell metadata

The following metadata keys are defined at the cell level:

Key	Value	Interpretation
collapsed	bool	Whether the cell's output container should be collapsed
autoscroll	bool or 'auto'	Whether the cell's output is scrolled, unscrolled, or autoscrolled
deletable	bool	If False, prevent deletion of the cell
format	'mime/type'	The mime-type of a *Raw NBConvert Cell*
name	str	A name for the cell. Should be unique
tags	list of str	A list of string tags on the cell. Commas are not allowed in a tag

1.4.3 Output metadata

The following metadata keys are defined for code cell outputs:

Key	Value	Interpretation
isolated	bool	Whether the output should be isolated into an IFrame

Python API for working with notebook files

2.1 Reading and writing

`nbformat.`**`read`**(*fp*, *as_version*, ***kwargs*)

Read a notebook from a file as a NotebookNode of the given version.

The string can contain a notebook of any version. The notebook will be returned *as_version*, converting, if necessary.

Notebook format errors will be logged.

> **Parameters**
>
> - **fp** (`file or str`) – A file-like object with a read method that returns unicode (use `io.open()` in Python 2), or a path to a file.
>
> - **as_version** (`int`) – The version of the notebook format to return. The notebook will be converted, if necessary. Pass nbformat.NO_CONVERT to prevent conversion.
>
> **Returns   nb** – The notebook that was read.
>
> **Return type** *NotebookNode*

`nbformat.`**`reads`**(*s*, *as_version*, ***kwargs*)

Read a notebook from a string and return the NotebookNode object as the given version.

The string can contain a notebook of any version. The notebook will be returned *as_version*, converting, if necessary.

Notebook format errors will be logged.

> **Parameters**
>
> - **s** (`unicode`) – The raw unicode string to read the notebook from.
>
> - **as_version** (`int`) – The version of the notebook format to return. The notebook will be converted, if necessary. Pass nbformat.NO_CONVERT to prevent conversion.
>
> **Returns   nb** – The notebook that was read.

> > **Return type** *NotebookNode*

The reading functions require you to pass the *as_version* parameter. Your code should specify the notebook format that it knows how to work with: for instance, if your code handles version 4 notebooks:

```
nb = nbformat.read('path/to/notebook.ipynb', as_version=4)
```

This will automatically upgrade or downgrade notebooks in other versions of the notebook format to the structure your code knows about.

nbformat.**write**(*nb*, *fp*, *version=nbformat.NO_CONVERT*, ***kwargs*)
> Write a notebook to a file in a given nbformat version.
>
> The file-like object must accept unicode input.
>
> > **Parameters**
> >
> > - **nb** (NotebookNode) – The notebook to write.
> >
> > - **fp** (*file or str*) – Any file-like object with a write method that accepts unicode, or a path to write a file.
> >
> > - **version** (*int, optional*) – The nbformat version to write. If nb is not this version, it will be converted. If unspecified, or specified as nbformat.NO_CONVERT, the notebook's own version will be used and no conversion performed.

nbformat.**writes**(*nb*, *version=nbformat.NO_CONVERT*, ***kwargs*)
> Write a notebook to a string in a given format in the given nbformat version.
>
> Any notebook format errors will be logged.
>
> > **Parameters**
> >
> > - **nb** (NotebookNode) – The notebook to write.
> >
> > - **version** (*int, optional*) – The nbformat version to write. If unspecified, or specified as nbformat.NO_CONVERT, the notebook's own version will be used and no conversion performed.
> >
> > **Returns** **s** – The notebook as a JSON string.
> >
> > **Return type** unicode

nbformat.**NO_CONVERT**
> This special value can be passed to the reading and writing functions, to indicate that the notebook should be loaded/saved in the format it's supplied.

nbformat.**current_nbformat**
nbformat.**current_nbformat_minor**
> These integers represent the current notebook format version that the nbformat module knows about.

2.2 NotebookNode objects

The functions in this module work with *NotebookNode* objects, which are like dictionaries, but allow attribute access (nb.cells). The structure of these objects matches the notebook format described in *The Notebook file format*.

class nbformat.**NotebookNode**(**args*, ***kw*)
> A dict-like node with attribute-access

---

nbformat.**from_dict**(*d*)
: Convert dict to dict-like NotebookNode

    Recursively converts any dict in the container to a NotebookNode. This does not check that the contents of the dictionary make a valid notebook or part of a notebook.

2.3 Other functions

nbformat.**convert**(*nb*, *to_version*)
: Convert a notebook node object to a specific version. Assumes that all the versions starting from 1 to the latest major X are implemented. In other words, there should never be a case where v1 v2 v3 v5 exist without a v4. Also assumes that all conversions can be made in one step increments between major versions and ignores minor revisions.

    Parameters

    - **nb** (`NotebookNode`) –

    - **to_version** (*int*) – Major revision to convert the notebook to. Can either be an upgrade or a downgrade.

nbformat.**validate**(*nbjson*, *ref=None*, *version=None*, *version_minor=None*)
: Checks whether the given notebook JSON conforms to the current notebook format schema.

    Raises ValidationError if not valid.

class nbformat.**ValidationError**(*message*, *validator=<unset>*, *path=()*, *cause=None*, *context=()*, *validator_value=<unset>*, *instance=<unset>*, *schema=<unset>*, *schema_path=()*, *parent=None*)

2.4 Constructing notebooks programmatically

These functions return *NotebookNode* objects with the necessary fields.

nbformat.v4.**new_notebook**(***kwargs*)
: Create a new notebook

nbformat.v4.**new_code_cell**(*source=''*, ***kwargs*)
: Create a new code cell

nbformat.v4.**new_markdown_cell**(*source=''*, ***kwargs*)
: Create a new markdown cell

nbformat.v4.**new_raw_cell**(*source=''*, ***kwargs*)
: Create a new raw cell

nbformat.v4.**new_output**(*output_type*, *data=None*, ***kwargs*)
: Create a new output, to go in the `cell.outputs` list of a code cell.

nbformat.v4.**output_from_msg**(*msg*)
: Create a NotebookNode for an output from a kernel's IOPub message.

    Returns NotebookNode

    Return type the output as a notebook node.

    Raises ValueError: if the message is not an output message.

2.5 Notebook signatures

This machinery is used by the notebook web application to record which notebooks are *trusted*, and may show dynamic output as soon as they're loaded. See Security in notebook documents for more information.

class `nbformat.sign.`**NotebookNotary**(***kwargs*)
> A class for computing and verifying notebook signatures.
>
> **sign**(*nb*)
> > Sign a notebook, indicating that its output is trusted on this machine
> >
> > Stores hash algorithm and hmac digest in a local database of trusted notebooks.
>
> **unsign**(*nb*)
> > Ensure that a notebook is untrusted
> >
> > by removing its signature from the trusted database, if present.
>
> **check_signature**(*nb*)
> > Check a notebook's stored signature
> >
> > If a signature is stored in the notebook's metadata, a new signature is computed and compared with the stored value.
> >
> > Returns True if the signature is found and matches, False otherwise.
> >
> > The following conditions must all be met for a notebook to be trusted: - a signature is stored in the form 'scheme:hexdigest' - the stored scheme matches the requested scheme - the requested scheme is available from hashlib - the computed hash from notebook_signature matches the stored hash
>
> **mark_cells**(*nb*, *trusted*)
> > Mark cells as trusted if the notebook's signature can be verified
> >
> > Sets `cell.metadata.trusted = True | False` on all code cells, depending on the *trusted* parameter. This will typically be the return value from `self.check_signature(nb)`.
> >
> > This function is the inverse of check_cells
>
> **check_cells**(*nb*)
> > Return whether all code cells are trusted.
> >
> > A cell is trusted if the 'trusted' field in its metadata is truthy, or if it has no potentially unsafe outputs. If there are no code cells, return True.
> >
> > This function is the inverse of mark_cells.

2.5.1 Signature storage

Signatures are stored using a pluggable *SignatureStore* subclass. To implement your own, override the methods below and configure `NotebookNotary.store_factory`.

class `nbformat.sign.`**SignatureStore**
> Base class for a signature store.
>
> **store_signature**(*digest*, *algorithm*)
> > Implement in subclass to store a signature.
> >
> > Should not raise if the signature is already stored.
>
> **remove_signature**(*digest*, *algorithm*)
> > Implement in subclass to delete a signature.

Should not raise if the signature is not stored.

check_signature(*digest*, *algorithm*)
Implement in subclass to check if a signature is known.

Return True for a known signature, False for unknown.

close()
Close any open connections this store may use.

If the store maintains any open connections (e.g. to a database), they should be closed.

By default, *NotebookNotary* will use an SQLite based store if SQLite bindings are available, and an in-memory store otherwise.

class nbformat.sign.**SQLiteSignatureStore**(*db_file*, ***kwargs*)
Store signatures in an SQLite database.

class nbformat.sign.**MemorySignatureStore**
Non-persistent storage of signatures in memory.

Changes in nbformat

3.1 4.3

- A new pluggable `SignatureStore` class allows specifying different ways to record the signatures of trusted notebooks. The default is still an SQLite database. See *Signature storage* for more information.

- `nbformat.read()` and `nbformat.write()` accept file paths as bytes as well as unicode.

- Fix for calling `nbformat.validate()` on an empty dictionary.

- Fix for running the tests where the locale makes ASCII the default encoding.

3.2 4.2

3.2.1 4.2.0

- Update nbformat spec version to 4.2, allowing JSON outputs to have any JSONable type, not just `object`, and mime-types of the form `application/anything+json`.

- Define basics of `authors` in notebook metadata. `nb.metadata.authors` shall be a list of objects with the property `name`, a string of each author's full name.

- Update use of traitlets API to require traitlets 4.1.

- Support trusting notebooks on stdin with `cat notebook | jupyter trust`

3.3 4.1

3.3.1 4.1.0

4.1 on GitHub

- Update nbformat spec version to 4.1, adding support for attachments on markdown and raw cells.
- Catch errors opening trust database, falling back on `:memory:` if the database cannot be opened.

3.4 4.0

4.0 on GitHub

The first release of nbformat as its own package.

Indices and tables

- genindex
- modindex
- search

Python Module Index

n