
Tastypie Documentation

Release 0.9.11

Daniel Lindsley, Cody Soyland & Matt Croydon

December 22, 2012

CONTENTS

Tastypie is an webservice API framework for Django. It provides a convenient, yet powerful and highly customizable, abstraction for creating REST-style interfaces.

GETTING STARTED WITH TASTYPIE

Tastypie is a reusable app (that is, it relies only on its own code and focuses on providing just a REST-style API) and is suitable for providing an API to any application without having to modify the sources of that app.

Not everyone's needs are the same, so Tastypie goes out of its way to provide plenty of hooks for overriding or extending how it works.

Note: If you hit a stumbling block, you can join #tastypie on irc.freenode.net to get help.

This tutorial assumes that you have a basic understanding of Django as well as how proper REST-style APIs ought to work. We will only explain the portions of the code that are Tastypie-specific in any kind of depth.

For example purposes, we'll be adding an API to a simple blog application. Here is `myapp/models.py`:

```
import datetime
from django.contrib.auth.models import User
from django.db import models
from django.template.defaultfilters import slugify

class Entry(models.Model):
    user = models.ForeignKey(User)
    pub_date = models.DateTimeField(default=datetime.datetime.now)
    title = models.CharField(max_length=200)
    slug = models.SlugField()
    body = models.TextField()

    def __unicode__(self):
        return self.title

    def save(self, *args, **kwargs):
        # For automatic slug generation.
        if not self.slug:
            self.slug = slugify(self.title)[:50]

        return super(Entry, self).save(*args, **kwargs)
```

With that, we'll move on to installing and configuring Tastypie.

1.1 Installation

Installing Tastypie is as simple as checking out the source and adding it to your project or `PYTHONPATH`.

1. Download the dependencies:
 - Python 2.4+
 - Django 1.0+ (tested on Django 1.1+)
 - `mimemparse` 0.1.3+ (<http://code.google.com/p/mimemparse/>)
 - Older versions will work, but their behavior on JSON/JSONP is a touch wonky.
 - `dateutil` (<http://labix.org/python-dateutil>)
 - **OPTIONAL** - `lxml` (<http://codespeak.net/lxml/>) if using the XML serializer
 - **OPTIONAL** - `pyyaml` (<http://pyyaml.org/>) if using the YAML serializer
 - **OPTIONAL** - `uuid` (present in 2.5+, downloadable from <http://pypi.python.org/pypi/uuid/>) if using the `ApiKey` authentication
2. Either check out `tastypie` from [GitHub](#) or to pull a release off [PyPI](#). Doing `sudo pip install django-tastypie` or `sudo easy_install django-tastypie` is all that should be required.
3. Either symlink the `tastypie` directory into your project or copy the directory in. What ever works best for you.

Note: Once `tastypie` reaches version 1.0, it will become officially available on [PyPI](#). Once that is the case, a `sudo pip install tastypie` or `sudo easy_install tastypie` should be available.

1.2 Configuration

The only mandatory configuration is adding `'tastypie'` to your `INSTALLED_APPS`. This isn't strictly necessary, as `Tastypie` has only two non-required models, but may ease usage.

You have the option to set up a number of settings (see *Tastypie Settings*) but they all have sane defaults and are not required unless you need to tweak their values.

1.3 Creating Resources

REST-style architecture talks about resources, so unsurprisingly integrating with `Tastypie` involves creating `Resource` classes. For our simple application, we'll create a file for these in `myapp/api.py`, though they can live anywhere in your application:

```
# myapp/api.py
from tastypie.resources import ModelResource
from myapp.models import Entry

class EntryResource(ModelResource):
    class Meta:
        queryset = Entry.objects.all()
        resource_name = 'entry'
```

This class, by virtue of being a `ModelResource` subclass, will introspect all non-relational fields on the `Entry` model and create it's own `ApiFields` that map to those fields, much like the way Django's `ModelForm` class introspects.

Note: The `resource_name` within the `Meta` class is optional. If not provided, it is automatically generated off the classname, removing any instances of `Resource` and lowercasing the string. So `EntryResource` would become just `entry`.

We've included the `resource_name` attribute in this example for clarity, especially when looking at the URLs, but you should feel free to omit it if you're comfortable with the automatic behavior.

1.4 Hooking Up The Resource(s)

Now that we have our `EntryResource`, we can hook it up in our `URLconf`. To do this, we simply instantiate the resource in our `URLconf` and hook up its `urls`:

```
# urls.py
from django.conf.urls.defaults import *
from myapp.api import EntryResource

entry_resource = EntryResource()

urlpatterns = patterns('',
    # The normal jazz here...
    (r'^blog/', include('myapp.urls')),
    (r'^api/', include(entry_resource.urls)),
)
```

Now it's just a matter of firing up server (`./manage.py runserver`) and going to `http://127.0.0.1:8000/api/entry/?format=json`. You should get back a list of `Entry`-like objects.

Note: The `?format=json` is an override required to make things look decent in the browser (accept headers vary between browsers). `Tastypie` properly handles the `Accept` header. So the following will work properly:

```
curl -H "Accept: application/json" http://127.0.0.1:8000/api/entry/
```

But if you're sure you want something else (or want to test in a browser), `Tastypie` lets you specify `?format=...` when you really want to force a certain type.

At this point, a bunch of other URLs are also available. Try out any/all of the following (assuming you have at least three records in the database):

- `http://127.0.0.1:8000/api/entry/?format=json`
- `http://127.0.0.1:8000/api/entry/1/?format=json`
- `http://127.0.0.1:8000/api/entry/schema/?format=json`
- `http://127.0.0.1:8000/api/entry/set/1;3/?format=json`

However, if you try sending a `POST/PUT/DELETE` to the resource, you find yourself getting “401 Unauthorized” errors. For safety, `Tastypie` ships with the `authorization` class (“what are you allowed to do”) set to `ReadOnlyAuthorization`. This makes it safe to expose on the web, but prevents us from doing `POST/PUT/DELETE`. Let's enable those:

```
# myapp/api.py
from tastypie.authorization import Authorization
from tastypie.resources import ModelResource
from myapp.models import Entry
```

```
class EntryResource(ModelResource):
    class Meta:
        queryset = Entry.objects.all()
        resource_name = 'entry'
        authorization= Authorization()
```

Warning: This is now great for testing in development but **VERY INSECURE**. You should never put a Resource like this out on the internet. Please spend some time looking at the authentication/authorization classes available in Tastypie.

With just nine lines of code, we have a full working REST interface to our Entry model. In addition, full GET/POST/PUT/DELETE support is already there, so it's possible to really work with all of the data. Well, *almost*.

You see, you'll note that not quite all of our data is there. Markedly absent is the user field, which is a ForeignKey to Django's User model. Tastypie does **NOT** introspect related data because it has no way to know how you want to represent that data.

And since that relation isn't there, any attempt to POST/PUT new data will fail, because no user is present, which is a required field on the model.

This is easy to fix, but we'll need to flesh out our API a little more.

1.5 Creating More Resources

In order to handle our user relation, we'll need to create a UserResource and tell the EntryResource to use it. So we'll modify `myapp/api.py` to match the following code:

```
# myapp/api.py
from django.contrib.auth.models import User
from tastypie import fields
from tastypie.resources import ModelResource
from myapp.models import Entry

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'user'

class EntryResource(ModelResource):
    user = fields.ForeignKey(UserResource, 'user')

    class Meta:
        queryset = Entry.objects.all()
        resource_name = 'entry'
```

We simply created a new ModelResource subclass called UserResource. Then we added a field to EntryResource that specified that the user field points to a UserResource for that data.

Now we should be able to get all of the fields back in our response. But since we have another full, working resource on our hands, we should hook that up to our API as well. And there's a better way to do it.

1.6 Adding To The Api

Tastypie ships with an `Api` class, which lets you bind multiple `Resources` together to form a coherent API. Adding it to the mix is simple.

We'll go back to our `URLconf` (`urls.py`) and change it to match the following:

```
# urls.py
from django.conf.urls.defaults import *
from tastypie.api import Api
from myapp.api import EntryResource, UserResource

v1_api = Api(api_name='v1')
v1_api.register(UserResource())
v1_api.register(EntryResource())

urlpatterns = patterns('',
    # The normal jazz here...
    (r'^blog/', include('myapp.urls')),
    (r'^api/', include(v1_api.urls)),
)
```

Note that we're now creating an `Api` instance, registering our `EntryResource` and `UserResource` instances with it and that we've modified the `urls` to now point to `v1_api.urls`.

This makes even more data accessible, so if we start up the `runserver` again, the following URLs should work:

- `http://127.0.0.1:8000/api/v1/?format=json`
- `http://127.0.0.1:8000/api/v1/user/?format=json`
- `http://127.0.0.1:8000/api/v1/user/1/?format=json`
- `http://127.0.0.1:8000/api/v1/user/schema/?format=json`
- `http://127.0.0.1:8000/api/v1/user/set/1;3/?format=json`
- `http://127.0.0.1:8000/api/v1/entry/?format=json`
- `http://127.0.0.1:8000/api/v1/entry/1/?format=json`
- `http://127.0.0.1:8000/api/v1/entry/schema/?format=json`
- `http://127.0.0.1:8000/api/v1/entry/set/1;3/?format=json`

Additionally, the representations out of `EntryResource` will now include the `user` field and point to an endpoint like `/api/v1/users/1/` to access that user's data. And full `POST/PUT` delete support should now work.

But there's several new problems. One is that our new `UserResource` leaks too much data, including fields like `email`, `password`, `is_active` and `is_staff`. Another is that we may not want to allow end users to alter `User` data. Both of these problems are easily fixed as well.

1.7 Limiting Data And Access

Cutting out the `email`, `password`, `is_active` and `is_staff` fields is easy to do. We simply modify our `UserResource` code to match the following:

```
class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
```

```
resource_name = 'user'  
excludes = ['email', 'password', 'is_active', 'is_staff', 'is_superuser']
```

The `excludes` directive tells `UserResource` which fields not to include in the output. If you'd rather whitelist fields, you could do:

```
class UserResource(ModelResource):  
    class Meta:  
        queryset = User.objects.all()  
        resource_name = 'user'  
        fields = ['username', 'first_name', 'last_name', 'last_login']
```

Now that the undesirable fields are no longer included, we can look at limiting access. This is also easy and involves making our `UserResource` look like:

```
class UserResource(ModelResource):  
    class Meta:  
        queryset = User.objects.all()  
        resource_name = 'user'  
        excludes = ['email', 'password', 'is_active', 'is_staff', 'is_superuser']  
        allowed_methods = ['get']
```

Now only HTTP GET requests will be allowed on `/api/v1/user/` endpoints. If you require more granular control, both `list_allowed_methods` and `detail_allowed_methods` options are supported.

1.8 Beyond The Basics

We now have a full working API for our application. But Tastypie supports many more features, like:

- *Authentication / Authorization*
- *Caching*
- *Throttling*
- *Resources* (filtering & sorting)
- *Serialization*

Tastypie is also very easy to override and extend. For some common patterns and approaches, you should refer to the *Tastypie Cookbook* documentation.

INTERACTING WITH THE API

Now that you've got a shiny new REST-style API in place, let's demonstrate how to interact with it. We'll assume that you have `cURL` installed on your system (generally available on most modern Mac & Linux machines), but any tool that allows you to control headers & bodies on requests will do.

We'll assume that we're interacting with the following Tastypie code:

```
# myapp/api/resources.py
from django.contrib.auth.models import User
from tastypie.authorization import Authorization
from tastypie import fields
from tastypie.resources import ModelResource, ALL, ALL_WITH_RELATIONS
from myapp.models import Entry

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'user'
        excludes = ['email', 'password', 'is_active', 'is_staff', 'is_superuser']
        filtering = {
            'username': ALL,
        }

class EntryResource(ModelResource):
    user = fields.ForeignKey(UserResource, 'user')

    class Meta:
        queryset = Entry.objects.all()
        resource_name = 'entry'
        authorization = Authorization()
        filtering = {
            'user': ALL_WITH_RELATIONS,
            'pub_date': ['exact', 'lt', 'lte', 'gte', 'gt'],
        }

# urls.py
from django.conf.urls.defaults import *
from tastypie.api import Api
from myapp.api.resources import EntryResource, UserResource

v1_api = Api(api_name='v1')
v1_api.register(UserResource())
```

```
v1_api.register(EntryResource())

urlpatterns = patterns('',
    # The normal jazz here...
    (r'^blog/', include('myapp.urls')),
    (r'^api/', include(v1_api.urls)),
)
```

Let's fire up a shell & start exploring the API!

2.1 Front Matter

Tastypie tries to treat all clients & all serialization types as equally as possible. It also tries to be a good 'Net citizen & respects the HTTP method used as well as the `Accepts` headers sent. Between these two, you control all interactions with Tastypie through relatively few endpoints.

Warning: Should you try these URLs in your browser, be warned you **WILL** need to append `?format=json` (or `xml` or `yaml`) to the URL. Your browser requests `application/xml` before `application/json`, so you'll always get back XML if you don't specify it. That's also why it's recommended that you explore via `curl`, because you avoid your browser's opinionated requests & get something closer to what any programmatic clients will get.

2.2 Fetching Data

Since reading data out of an API is a very common activity (and the easiest type of request to make), we'll start there. Tastypie tries to expose various parts of the API & interlink things within the API (HATEOAS).

2.2.1 Api-Wide

We'll start at the highest level:

```
curl http://localhost:8000/api/v1/
```

You'll get back something like:

```
{
  "entry": {
    "list_endpoint": "/api/v1/entry/",
    "schema": "/api/v1/entry/schema/"
  },
  "user": {
    "list_endpoint": "/api/v1/user/",
    "schema": "/api/v1/user/schema/"
  }
}
```

This lists out all the different `Resource` classes you registered in your `URLconf` with the API. Each one is listed by the `resource_name` you gave it and provides the `list_endpoint` & the `schema` for the resource.

Note that these links try to direct you to other parts of the API, to make exploration/discovery easier. We'll use these URLs in the next several sections.

To demonstrate another format, you could run the following to get the XML variant of the same information:

```
curl -H "Accept: application/xml" http://localhost:8000/api/v1/
```

To which you'd receive:

```
<?xml version="1.0" encoding="utf-8"?>
<response>
  <entry type="hash">
    <list_endpoint>/api/v1/entry/</list_endpoint>
    <schema>/api/v1/entry/schema/</schema>
  </entry>
  <user type="hash">
    <list_endpoint>/api/v1/user/</list_endpoint>
    <schema>/api/v1/user/schema/</schema>
  </user>
</response>
```

We'll stick to JSON for the rest of this document, but using XML should be OK to do at any time.

2.2.2 Inspecting The Resource's Schema

Since the api-wide view gave us a schema URL, let's inspect that next. We'll use the entry resource. Again, a simple GET request by curl:

```
curl http://localhost:8000/api/v1/entry/schema/
```

This time, we get back a lot more data:

```
{
  "default_format": "application/json",
  "fields": {
    "body": {
      "help_text": "Unicode string data. Ex: \"Hello World\"",
      "nullable": false,
      "readonly": false,
      "type": "string"
    },
    "id": {
      "help_text": "Unicode string data. Ex: \"Hello World\"",
      "nullable": false,
      "readonly": false,
      "type": "string"
    },
    "pub_date": {
      "help_text": "A date & time as a string. Ex: \"2010-11-10T03:07:43\"",
      "nullable": false,
      "readonly": false,
      "type": "datetime"
    },
    "resource_uri": {
      "help_text": "Unicode string data. Ex: \"Hello World\"",
      "nullable": false,
      "readonly": true,
      "type": "string"
    },
    "slug": {
      "help_text": "Unicode string data. Ex: \"Hello World\"",
```

```
        "nullable": false,
        "readonly": false,
        "type": "string"
    },
    "title": {
        "help_text": "Unicode string data. Ex: \"Hello World\"",
        "nullable": false,
        "readonly": false,
        "type": "string"
    },
    "user": {
        "help_text": "A single related resource. Can be either a URI or set of nested resource data",
        "nullable": false,
        "readonly": false,
        "type": "related"
    }
},
"filtering": {
    "pub_date": ["exact", "lt", "lte", "gte", "gt"],
    "user": 2
}
}
```

This lists out the `default_format` this resource responds with, the fields on the resource & the filtering options available. This information can be used to prepare the other aspects of the code for the data it can obtain & ways to filter the resources.

2.2.3 Getting A Collection Of Resources

Let's get down to fetching live data. From the api-wide view, we'll hit the `list_endpoint` for entry:

```
curl http://localhost:8000/api/v1/entry/
```

We get back data that looks like:

```
{
  "meta": {
    "limit": 20,
    "next": null,
    "offset": 0,
    "previous": null,
    "total_count": 3
  },
  "objects": [{
    "body": "Welcome to my blog!",
    "id": "1",
    "pub_date": "2011-05-20T00:46:38",
    "resource_uri": "/api/v1/entry/1/",
    "slug": "first-post",
    "title": "First Post",
    "user": "/api/v1/user/1/"
  },
  {
    "body": "Well, it's been awhile and I still haven't updated. ",
    "id": "2",
    "pub_date": "2011-05-21T00:46:58",
    "resource_uri": "/api/v1/entry/2/"
  }
}]
```



```

    "slug": "second-post",
    "title": "Second Post",
    "user": "/api/v1/user/1/"
  },
  {
    "body": "I'm really excited to get started with this new blog. It's gonna be great!",
    "id": "3",
    "pub_date": "2011-05-20T00:47:30",
    "resource_uri": "/api/v1/entry/3/",
    "slug": "my-blog",
    "title": "My Blog",
    "user": "/api/v1/user/2/"
  }
]
}

```

Some things to note:

- By default, you get a paginated set of objects (20 per page is the default).
- In the meta, you get a previous & next. If available, these are URIs to the previous & next pages.
- You get a list of resources/objects under the `objects` key.
- Each resources/object has a `resource_uri` field that points to the detail view for that object.
- The foreign key to User is represented as a URI by default. If you're looking for the full UserResource to be embedded in this view, you'll need to add `full=True` to the `fields.ToOneField`.

If you want to skip paginating, simply run:

```
curl http://localhost:8000/api/v1/entry/?limit=0
```

Be warned this will return all objects, so it may be a CPU/IO-heavy operation on large datasets.

Let's try filtering on the resource. Since we know we can filter on the user, we'll fetch all posts by the daniel user with:

```
curl http://localhost:8000/api/v1/entry/?user__username=daniel
```

We get back what we asked for:

```

{
  "meta": {
    "limit": 20,
    "next": null,
    "offset": 0,
    "previous": null,
    "total_count": 2
  },
  "objects": [{
    "body": "Welcome to my blog!",
    "id": "1",
    "pub_date": "2011-05-20T00:46:38",
    "resource_uri": "/api/v1/entry/1/",
    "slug": "first-post",
    "title": "First Post",
    "user": "/api/v1/user/1/"
  },
  {
    "body": "Well, it's been awhile and I still haven't updated. ",
    "id": "2",

```

```
        "pub_date": "2011-05-21T00:46:58",
        "resource_uri": "/api/v1/entry/2/",
        "slug": "second-post",
        "title": "Second Post",
        "user": "/api/v1/user/1/"
    }
}
```

Where there were three posts before, now there are only two.

2.2.4 Getting A Detail Resource

Since each resource/object in the list view had a `resource_uri`, let's explore what's there:

```
curl http://localhost:8000/api/v1/entry/1/
```

We get back a similar set of data that we received from the list view:

```
{
  "body": "Welcome to my blog!",
  "id": "1",
  "pub_date": "2011-05-20T00:46:38",
  "resource_uri": "/api/v1/entry/1/",
  "slug": "first-post",
  "title": "First Post",
  "user": "/api/v1/user/1/"
}
```

Where this proves useful (for example) is present in the data we got back. We know the URI of the `User` associated with this blog entry. Let's run:

```
curl http://localhost:8000/api/v1/user/1/
```

Without ever seeing any aspect of the `UserResource` & just following the URI given, we get back:

```
{
  "date_joined": "2011-05-20T00:42:14.990617",
  "first_name": "",
  "id": "1",
  "last_login": "2011-05-20T00:44:57.510066",
  "last_name": "",
  "resource_uri": "/api/v1/user/1/",
  "username": "daniel"
}
```

You can do a similar fetch using the following Javascript/jQuery (though be wary of same-domain policy):

```
$.ajax({
  url: 'http://localhost:8000/api/v1/user/1/',
  type: 'GET',
  accepts: 'application/json',
  dataType: 'json'
})
```

2.2.5 Selecting A Subset Of Resources

Sometimes you may want back more than one record, but not an entire list view nor do you want to do multiple requests. Tastypie includes a “set” view, which lets you cherry-pick the objects you want. For example, if we just want the first & third `Entry` resources, we’d run:

```
curl "http://localhost:8000/api/v1/entry/set/1;3/"
```

Note: Quotes are needed in this case because of the semicolon delimiter between primary keys. Without the quotes, bash tries to split it into two statements. No extraordinary quoting will be necessary in your application (unless your API client is written in bash :D).

And we get back just those two objects:

```
{
  "objects": [{
    "body": "Welcome to my blog!",
    "id": "1",
    "pub_date": "2011-05-20T00:46:38",
    "resource_uri": "/api/v1/entry/1/",
    "slug": "first-post",
    "title": "First Post",
    "user": "/api/v1/user/1/"
  },
  {
    "body": "I'm really excited to get started with this new blog. It's gonna be great!",
    "id": "3",
    "pub_date": "2011-05-20T00:47:30",
    "resource_uri": "/api/v1/entry/3/",
    "slug": "my-blog",
    "title": "My Blog",
    "user": "/api/v1/user/2/"
  }
  ]
}
```

Note that, like the list view, you get back a list of objects. Unlike the list view, there is **NO** pagination applied to these objects. You asked for them, you’re going to get them all.

2.3 Sending Data

Tastypie also gives you full write capabilities in the API. Since the `EntryResource` has the `no-limits Authentication & Authorization` on it, we can freely write data.

Warning: Note that this is a huge security hole as well. Don’t put unauthorized write-enabled resources on the Internet, because someone will trash your data.

This is why `ReadOnlyAuthorization` is the default in Tastypie & why you must override to provide more access.

The good news is that there are no new URLs to learn. The “list” & “detail” URLs we’ve been using to fetch data *ALSO* support the `POST/PUT/DELETE` HTTP methods.

2.3.1 Creating A New Resource (POST)

Let's add a new entry. To create new data, we'll switch from GET requests to the familiar POST request.

Note: Tastypie encourages “round-trippable” data, which means the data you can GET should be able to be POST/PUT'd back to recreate the same object.

If you're ever in question about what you should send, do a GET on another object & see what Tastypie thinks it should look like.

To create new resources/objects, you will POST to the list endpoint of a resource. Trying to POST to a detail endpoint has a different meaning in the REST mindset (meaning to add a resource as a child of a resource of the same type).

As with all Tastypie requests, the headers we request are important. Since we've been using primarily JSON throughout, let's send a new entry in JSON format:

```
curl --dump-header - -H "Content-Type: application/json" -X POST --data '{"body": "This will prbbly b
```

The `Content-Type` header here informs Tastypie that we're sending it JSON. We send the data as a JSON-serialized body (**NOT** as form-data in the form of URL parameters). What we get back is the following response:

```
HTTP/1.0 201 CREATED
Date: Fri, 20 May 2011 06:48:36 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Type: text/html; charset=utf-8
Location: http://localhost:8000/api/v1/entry/4/
```

You'll also note that we get a correct HTTP status code back (201) & a `Location` header, which gives us the URI to our newly created resource.

Passing `--dump-header -` is important, because it gives you all the headers as well as the status code. When things go wrong, this will be useful information to help with debugging. For instance, if we send a request without a `user`:

```
curl --dump-header - -H "Content-Type: application/json" -X POST --data '{"body": "This will prbbly b
```

We get back:

```
HTTP/1.0 400 BAD REQUEST
Date: Fri, 20 May 2011 06:53:02 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Type: text/html; charset=utf-8
```

The `'user'` field has no data and doesn't allow a default or null value.

You can do a similar POST using the following Javascript/jQuery (though be wary of same-domain policy):

```
# This may require the ``json2.js`` library for older browsers.
var data = JSON.stringify({
  "body": "This will prbbly be my 1st post.",
  "pub_date": "2011-05-22T00:46:38",
  "slug": "another-post",
  "title": "Another Post"
});

$.ajax({
  url: 'http://localhost:8000/api/v1/entry/',
  type: 'POST',
  contentType: 'application/json',
```

```

data: data,
dataType: 'json',
processData: false
})

```

2.3.2 Updating An Existing Resource (PUT)

You might have noticed that we made some typos when we submitted the POST request. We can fix this using a PUT request to the detail endpoint (modify this instance of a resource):

```
curl --dump-header - -H "Content-Type: application/json" -X PUT --data '{"body": "This will probably
```

After fixing up the body, we get back:

```

HTTP/1.0 204 NO CONTENT
Date: Fri, 20 May 2011 07:13:21 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Length: 0
Content-Type: text/html; charset=utf-8

```

We get a 204 status code, meaning our update was successful. We don't get a Location header back because we did the PUT on a detail URL, which presumably did not change.

2.3.3 Updating A Whole Collection Of Resources (PUT)

You can also, in rare circumstances, update an entire collection of objects. By sending a PUT request to the list view of a resource, you can replace the entire collection.

Warning: This deletes all of the objects first, then creates the objects afresh. This is done because determining which objects are the same is actually difficult to get correct in the general case for all people.

Send a request like:

```
curl --dump-header - -H "Content-Type: application/json" -X PUT --data '{"objects": [{"body": "Welcom
```

And you'll get back a response like:

```

HTTP/1.0 204 NO CONTENT
Date: Fri, 20 May 2011 07:13:21 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Length: 0
Content-Type: text/html; charset=utf-8

```

2.4 Deleting Data

No CRUD setup would be complete without the ability to delete resources/objects. Deleting also requires significantly less complicated requests than POST/PUT.

2.4.1 Deleting A Single Resource

We've decided that we don't like the entry we added & edited earlier. Let's delete it (but leave the other objects alone):

```
curl --dump-header - -H "Content-Type: application/json" -X DELETE http://localhost:8000/api/v1/entry/
```

Once again, we get back the “Accepted” response of a 204:

```
HTTP/1.0 204 NO CONTENT
Date: Fri, 20 May 2011 07:28:01 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Length: 0
Content-Type: text/html; charset=utf-8
```

If we request that resource, we get a 410 to show it’s no longer there:

```
curl --dump-header - http://localhost:8000/api/v1/entry/4/
```

```
HTTP/1.0 410 GONE
Date: Fri, 20 May 2011 07:29:02 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Type: text/html; charset=utf-8
```

Additionally, if we try to run the DELETE again (using the same original command), we get the “Gone” response again:

```
HTTP/1.0 410 GONE
Date: Fri, 20 May 2011 07:30:00 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Type: text/html; charset=utf-8
```

2.4.2 Deleting A Whole Collection Of Resources

Finally, it’s possible to remove an entire collection of resources. This is as destructive as it sounds. Once again, we use the DELETE method, this time on the entire list endpoint:

```
curl --dump-header - -H "Content-Type: application/json" -X DELETE http://localhost:8000/api/v1/entry/
```

As a response, we get:

```
HTTP/1.0 204 NO CONTENT
Date: Fri, 20 May 2011 07:32:51 GMT
Server: WSGIServer/0.1 Python/2.7
Content-Length: 0
Content-Type: text/html; charset=utf-8
```

Hitting the list view:

```
curl --dump-header - http://localhost:8000/api/v1/entry/
```

Gives us a 200 but no objects:

```
{
  "meta": {
    "limit": 20,
    "next": null,
    "offset": 0,
    "previous": null,
    "total_count": 0
  },
  "objects": []
}
```

2.5 You Did It!

That's a whirlwind tour of interacting with a Tastypie API. There's additional functionality present, such as:

- `POST/PUT` the other supported content-types
- More filtering/`order_by/limit/offset` tricks
- Using overridden `URLconfs` to support complex or non-PK lookups
- Authentication

But this grounds you in the basics & hopefully clarifies usage/debugging better.

TASTYPIE SETTINGS

This is a comprehensive list of the settings Tastypie recognizes.

3.1 API_LIMIT_PER_PAGE

Optional

This setting controls the default number of records Tastypie will show in a list view.

This is only used when a user does not specify a `limit` GET parameter and the `Resource` subclass has not overridden the number to be shown.

An example:

```
API_LIMIT_PER_PAGE = 50
```

If you don't want to limit the number of records by default, you can set this setting to 0:

```
API_LIMIT_PER_PAGE = 0
```

Defaults to 20.

3.2 TASTYPIE_FULL_DEBUG

Optional

This setting controls what the behavior is when an unhandled exception occurs.

If set to `True` and `settings.DEBUG = True`, the standard Django technical 500 is displayed.

If not set or set to `False`, Tastypie will return a serialized response. If `settings.DEBUG` is `True`, you'll get the actual exception message plus a traceback. If `settings.DEBUG` is `False`, Tastypie will call `mail_admins()` and provide a canned error message (which you can override with `TASTYPIE_CANNED_ERROR`) in the response.

An example:

```
TASTYPIE_FULL_DEBUG = True
```

Defaults to `False`.

3.3 TASTYPIE_CANNED_ERROR

Optional

This setting allows you to override the canned error response when an unhandled exception is raised and `settings.DEBUG` is `False`.

An example:

```
TASTYPIE_CANNED_ERROR = "Oops, we broke it!"
```

Defaults to `"Sorry, this request could not be processed. Please try again later."`.

3.4 TASTYPIE_ALLOW_MISSING_SLASH

Optional

This setting allows your URLs to be missing the final slash. Useful for integrating with other systems.

You must also have `settings.APPEND_SLASH = False` so that Django does not emit HTTP 302 redirects.

Warning: This setting causes the `Resource.get_multiple()` method to fail. If you need this method, you will have to override the `URLconf` to meet your needs.

An example:

```
TASTYPIE_ALLOW_MISSING_SLASH = True
```

Defaults to `False`.

3.5 TASTYPIE_DATETIME_FORMATTING

Optional

This setting allows you to globally choose what format your datetime/date/time data will be formatted in. Valid options are `iso-8601` & `rfc-2822`.

An example:

```
TASTYPIE_DATETIME_FORMATTING = 'rfc-2822'
```

Defaults to `iso-8601`.

USING TASTYPIE WITH NON-ORM DATA SOURCES

Much of this documentation demonstrates the use of Tastypie with Django's ORM. You might think that Tastypie depended on the ORM, when in fact, it was purpose-built to handle non-ORM data. This documentation should help you get started providing APIs using other data sources.

Virtually all of the code that makes Tastypie actually process requests & return data is within the `Resource` class. `ModelResource` is actually a light wrapper around `Resource` that provides ORM-specific access. The methods that `ModelResource` overrides are the same ones you'll need to override when hooking up your data source.

4.1 Approach

When working with `Resource`, many things are handled for you. All the authentication/authorization/caching/serialization/throttling bits should work as normal and Tastypie can support all the REST-style methods. Schemas & discovery views all work the same as well.

What you don't get out of the box are the fields you're choosing to expose & the lowest level data access methods. If you want a full read-write API, there are nine methods you need to implement. They are:

- `get_resource_uri`
- `get_object_list`
- `obj_get_list`
- `obj_get`
- `obj_create`
- `obj_update`
- `obj_delete_list`
- `obj_delete`
- `rollback`

If read-only is all you're exposing, you can cut that down to four methods to override.

4.2 Using Riak for MessageResource

As an example, we'll take integrating with Riak (a Dynamo-like NoSQL store) since it has both a simple API and demonstrate what hooking up to a non-relational datastore looks like:

```
# We need a generic object to shove data in/get data from.
# Riak generally just tosses around dictionaries, so we'll lightly
# wrap that.
class RiakObject(object):
    def __init__(self, initial=None):
        self.__dict__['_data'] = {}

        if hasattr(initial, 'items'):
            self.__dict__['_data'] = initial

    def __getattr__(self, name):
        return self._data.get(name, None)

    def __setattr__(self, name, value):
        self.__dict__['_data'][name] = value

    def to_dict(self):
        return self._data

class MessageResource(Resource):
    # Just like a Django 'Form' or 'Model', we're defining all the
    # fields we're going to handle with the API here.
    uuid = fields.CharField(attribute='uuid')
    user_uuid = fields.CharField(attribute='user_uuid')
    message = fields.CharField(attribute='message')
    created = fields.IntegerField(attribute='created')

    class Meta:
        resource_name = 'riak'
        object_class = RiakObject
        authorization = Authorization()

    # Specific to this resource, just to get the needed Riak bits.
    def _client(self):
        return riak.RiakClient()

    def _bucket(self):
        client = self._client()
        # Note that we're hard-coding the bucket to use. Fine for
        # example purposes, but you'll want to abstract this.
        return client.bucket('messages')

    # The following methods will need overriding regardless of your
    # data source.
    def get_resource_uri(self, bundle_or_obj):
        kwargs = {
            'resource_name': self._meta.resource_name,
        }

        if isinstance(bundle_or_obj, Bundle):
            kwargs['pk'] = bundle_or_obj.obj.uuid
        else:
```

```

        kwargs['pk'] = bundle_or_obj.uuid

    if self._meta.api_name is not None:
        kwargs['api_name'] = self._meta.api_name

    return self._build_reverse_url("api_dispatch_detail", kwargs=kwargs)

def get_object_list(self, request):
    query = self._client().add('messages')
    query.map("function(v) { var data = JSON.parse(v.values[0].data); return [[v.key, data]]; }")
    results = []

    for result in query.run():
        new_obj = RiakObject(initial=result[1])
        new_obj.uuid = result[0]
        results.append(new_obj)

    return results

def obj_get_list(self, request=None, **kwargs):
    # Filtering disabled for brevity...
    return self.get_object_list(request)

def obj_get(self, request=None, **kwargs):
    bucket = self._bucket()
    message = bucket.get(kwargs['pk'])
    return RiakObject(initial=message.get_data())

def obj_create(self, bundle, request=None, **kwargs):
    bundle.obj = RiakObject(initial=kwargs)
    bundle = self.full_hydrate(bundle)
    bucket = self._bucket()
    new_message = bucket.new(bundle.obj.uuid, data=bundle.obj.to_dict())
    new_message.store()
    return bundle

def obj_update(self, bundle, request=None, **kwargs):
    return self.obj_create(bundle, request, **kwargs)

def obj_delete_list(self, request=None, **kwargs):
    bucket = self._bucket()

    for key in bucket.get_keys():
        obj = bucket.get(key)
        obj.delete()

def obj_delete(self, request=None, **kwargs):
    bucket = self._bucket()
    obj = bucket.get(kwargs['pk'])
    obj.delete()

def rollback(self, bundles):
    pass

```

This represents a full, working, Riak-powered API endpoint. All REST-style actions (GET/POST/PUT/DELETE) all work correctly. The only shortcut taken in this example was skipping filter-ability, as adding in the MapReduce bits would have decreased readability.

All said and done, just nine methods needed overriding, eight of which were highly specific to how data access is

done.

TOOLS

Here are some tools that might help in interacting with the API that Tastypie provides:

5.1 Browser

5.1.1 JSONView

- Firefox - <https://addons.mozilla.org/en-US/firefox/addon/jsonview/>
- Chrome - <https://chrome.google.com/webstore/detail/chklaanhfefbnpoihckbnefhakgolnmc>

A plugin (actually two different ones that closely mirror each other) that nicely reformats JSON data in the browser.

5.2 Python

5.2.1 Slumber

<http://slumber.in/>

Slumber is a small Python library that makes it easy to access & work with APIs. It works for many others, but works especially well with Tastypie.

5.2.2 json.tool

Included with Python, this tool makes reformatting JSON easy. For example:

```
$ curl http://localhost:8000/api/v1/note/ | python -m json.tool
```

Will return nicely reformatted data like:

```
{
  "meta": {
    "total_count": 1
  },
  "objects": [
    {
      "content": "Hello world!",
      "user": "/api/v1/user/1/"
    }
  ]
}
```

```
    ]  
}
```

5.3 Javascript

5.3.1 backbone-tastypie

<https://github.com/PaulUithol/backbone-tastypie>

A small layer that makes Backbone & Tastypie plan nicely together.

5.3.2 backbone-relational

<https://github.com/PaulUithol/Backbone-relational/>

Allows Backbone to work with relational data, like the kind of data Tastypie provides.

RESOURCES

In terms of a REST-style architecture, a “resource” is a collection of similar data. This data could be a table of a database, a collection of other resources or a similar form of data storage. In Tastypie, these resources are generally intermediaries between the end user & objects, usually Django models. As such, `Resource` (and its model-specific twin `ModelResource`) form the heart of Tastypie’s functionality.

6.1 Quick Start

A sample resource definition might look something like:

```
from django.contrib.auth.models import User
from tastypie import fields
from tastypie.authorization import DjangoAuthorization
from tastypie.resources import ModelResource, ALL, ALL_WITH_RELATIONS
from myapp.models import Entry

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']

class EntryResource(ModelResource):
    user = fields.ForeignKey(UserResource, 'user')

    class Meta:
        queryset = Entry.objects.all()
        list_allowed_methods = ['get', 'post']
        detail_allowed_methods = ['get', 'post', 'put', 'delete']
        resource_name = 'myapp/entry'
        authorization = DjangoAuthorization()
        filtering = {
            'slug': ALL,
            'user': ALL_WITH_RELATIONS,
            'created': ['exact', 'range', 'gt', 'gte', 'lt', 'lte'],
        }
```

6.2 Why Class-Based?

Using class-based resources make it easier to extend/modify the code to meet your needs. APIs are rarely a one-size-fits-all problem space, so Tastypie tries to get the fundamentals right and provide you with enough hooks to customize things to work your way.

As is standard, this raises potential problems for thread-safety. Tastypie has been designed to minimize the possibility of data “leaking” between threads. This does however sometimes introduce some small complexities & you should be careful not to store state on the instances if you’re going to be using the code in a threaded environment.

6.3 Why Resource vs. ModelResource?

Make no mistake that Django models are far and away the most popular source of data. However, in practice, there are many times where the ORM isn’t the data source. Hooking up things like a NoSQL store (see *Using Tastypie With Non-ORM Data Sources*), a search solution like Haystack or even managed filesystem data are all good use cases for `Resource` knowing nothing about the ORM.

6.4 Flow Through The Request/Response Cycle

Tastypie can be thought of as a set of class-based views that provide the API functionality. As such, many part of the request/response cycle are standard Django behaviors. For instance, all routing/middleware/response-handling aspects are the same as a typical Django app. Where it differs is in the view itself.

As an example, we’ll walk through what a GET request to a list endpoint (say `/api/v1/user/?format=json`) looks like:

- The `Resource.urls` are checked by Django’s url resolvers.
- On a match for the list view, `Resource.wrap_view('dispatch_list')` is called. `wrap_view` provides basic error handling & allows for returning serialized errors.
- Because `dispatch_list` was passed to `wrap_view`, `Resource.dispatch_list` is called next. This is a thin wrapper around `Resource.dispatch`.
- `dispatch` does a bunch of heavy lifting. It ensures:
 - the requested HTTP method is in `allowed_methods` (`method_check`),
 - the class has a method that can handle the request (`get_list`),
 - the user is authenticated (`is_authenticated`),
 - the user is authorized (`is_authorized`),
 - & the user has not exceeded their throttle (`throttle_check`).

At this point, `dispatch` actually calls the requested method (`get_list`).

- `get_list` does the actual work of the API. It does:
 - A fetch of the available objects via `Resource.obj_get_list`. In the case of `ModelResource`, this builds the ORM filters to apply (`ModelResource.build_filters`). It then gets the `QuerySet` via `ModelResource.get_object_list` (which performs `Resource.apply_authorization_limits` to possibly limit the set the user can work with) and applies the built filters to it.
 - It then sorts the objects based on user input (`ModelResource.apply_sorting`).

- Then it paginates the results using the supplied `Paginator` & pulls out the data to be serialized.
- The objects in the page have `full_dehydrate` applied to each of them, causing Tastypie to translate the raw object data into the fields the endpoint supports.
- Finally, it calls `Resource.create_response`.
- `create_response` is a shortcut method that:
 - Determines the desired response format (`Resource.determine_format`),
 - Serializes the data given to it in the proper format,
 - And returns a Django `HttpResponse` (200 OK) with the serialized data.
- We bubble back up the call stack to `dispatch`. The last thing `dispatch` does is potentially store that a request occurred for future throttling (`Resource.log_throttled_access`) then either returns the `HttpResponse` or wraps whatever data came back in a response (so Django doesn't freak out).

Processing on other endpoints or using the other HTTP methods results in a similar cycle, usually differing only in what “actual work” method gets called (which follows the format of “<http_method>_<list_or_detail>”). In the case of POST/PUT, the `hydrate` cycle additionally takes place and is used to take the user data & convert it to raw data for storage.

6.5 Why Resource URIs?

Resource URIs play a heavy role in how Tastypie delivers data. This can seem very different from other solutions which simply inline related data. Though Tastypie can inline data like that (using `full=True` on the field with the relation), the default is to provide URIs.

URIs are useful because it results in smaller payloads, letting you fetch only the data that is important to you. You can imagine an instance where an object has thousands of related items that you may not be interested in.

URIs are also very cache-able, because the data at each endpoint is less likely to frequently change.

And URIs encourage proper use of each endpoint to display the data that endpoint covers.

Ideology aside, you should use whatever suits you. If you prefer fewer requests & fewer endpoints, use of `full=True` is available, but be aware of the consequences of each approach.

6.6 Advanced Data Preparation

Not all data can be easily pull off an object/model attribute. And sometimes, you (or the client) may need to send data that doesn't neatly fit back into the data model on the server side. For this, Tastypie has the “dehydrate/hydrate” cycle.

6.6.1 The Dehydrate Cycle

Tastypie uses a “dehydrate” cycle to prepare data for serialization, which is to say that it takes the raw, potentially complicated data model & turns it into a (generally simpler) processed data structure for client consumption. This usually means taking a complex data object & turning it into a dictionary of simple data types.

Broadly speaking, this takes the `bundle.obj` instance & builds `bundle.data`, which is what is actually serialized.

The cycle looks like:

- Put the data model into a `Bundle` instance, which is then passed through the various methods.

- Run through all fields on the Resource, letting each field perform its own `dehydrate` method on the bundle.
- While processing each field, look for a `dehydrate_<fieldname>` method on the Resource. If it's present, call it with the bundle.
- Finally, after all fields are processed, if the `dehydrate` method is present on the Resource, it is called & given the entire bundle.

The goal of this cycle is to populate the `bundle.data` dictionary with data suitable for serialization. With the exception of the `alter_*` methods (as hooks to manipulate the overall structure), this cycle controls what is actually handed off to be serialized & sent to the client.

Per-field dehydrate

Each field (even custom `ApiField` subclasses) has its own `dehydrate` method. If it knows how to access data (say, given the `attribute` kwarg), it will attempt to populate values itself.

The return value is put in the `bundle.data` dictionary (by the Resource) with the `fieldname` as the key.

`dehydrate_FOO`

Since not all data may be ready for consumption based on just attribute access (or may require an advanced lookup/calculation), this hook enables you to fill in data or massage whatever the field generated.

Note: The `FOO` here is not literal. Instead, it is a placeholder that should be replaced with the `fieldname` in question.

Defining these methods is especially common when denormalizing related data, providing statistics or filling in unrelated data.

A simple example:

```
class MyResource(ModelResource):
    # The 'title' field is already added to the class by 'ModelResource'
    # and populated off 'Note.title'. But we want allcaps titles...

    class Meta:
        queryset = Note.objects.all()

    def dehydrate_title(self, bundle):
        return bundle.data['title'].upper()
```

A complex example:

```
class MyResource(ModelResource):
    # As is, this is just an empty field. Without the 'dehydrate_rating'
    # method, no data would be populated for it.
    rating = fields.FloatField(readonly=True)

    class Meta:
        queryset = Note.objects.all()

    def dehydrate_rating(self, bundle):
        total_score = 0.0

        # Make sure we don't have to worry about "divide by zero" errors.
        if not bundle.obj.rating_set.count():
```

```

    return rating

    # We'll run over all the ``Rating`` objects & calculate an average.
    for rating in bundle.obj.rating_set.all():
        total_score += rating.rating

    return total_score / bundle.obj.rating_set.count()

```

The return value is updated in the `bundle.data`. You should avoid altering `bundle.data` here if you can help it.

dehydrate

The `dehydrate` method takes a now fully-populated `bundle.data` & make any last alterations to it. This is useful for when a piece of data might depend on more than one field, if you want to shove in extra data that isn't worth having its own field or if you want to dynamically remove things from the data to be returned.

A simple example:

```

class MyResource(ModelResource):
    class Meta:
        queryset = Note.objects.all()

    def dehydrate(self, bundle):
        # Include the request IP in the bundle.
        bundle.data['request_ip'] = bundle.request.META.get('REMOTE_ADDR')
        return bundle

```

A complex example:

```

class MyResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        excludes = ['email', 'password', 'is_staff', 'is_superuser']

    def dehydrate(self, bundle):
        # If they're requesting their own record, add in their email address.
        if bundle.request.user.pk == bundle.obj.pk:
            # Note that there isn't an ``email`` field on the ``Resource``.
            # By this time, it doesn't matter, as the built data will no
            # longer be checked against the fields on the ``Resource``.
            bundle.data['email'] = bundle.obj.email

        return bundle

```

This method should return a `bundle`, whether it modifies the existing one or creates a whole new one. You can even remove any/all data from the `bundle.data` if you wish.

6.6.2 The Hydrate Cycle

Tastypie uses a “hydrate” cycle to take serialized data from the client and turn it into something the data model can use. This is the reverse process from the `dehydrate` cycle. In fact, by default, Tastypie’s serialized data should be “round-trip-able”, meaning the data that comes out should be able to be fed back in & result in the same original data model. This usually means taking a dictionary of simple data types & turning it into a complex data object.

Broadly speaking, this takes the recently-deserialized `bundle.data` dictionary & builds `bundle.obj` (but does NOT save it).

The cycle looks like:

- Put the data from the client into a `Bundle` instance, which is then passed through the various methods.
- If the `hydrate` method is present on the `Resource`, it is called & given the entire bundle.
- Then run through all fields on the `Resource`, look for a `hydrate_<fieldname>` method on the `Resource`. If it's present, call it with the bundle.
- Finally after all other processing is done, while processing each field, let each field perform its own `hydrate` method on the bundle.

The goal of this cycle is to populate the `bundle.obj` data model with data suitable for saving/persistence. Again, with the exception of the `alter_*` methods (as hooks to manipulate the overall structure), this cycle controls what how the data from the client is interpreted & placed on the data model.

hydrate

The `hydrate` method allows you to make final changes to the `bundle.obj`. This includes things like prepopulating fields you don't expose over the API, recalculating related data or mangling data.

Example:

```
class MyResource(ModelResource):
    # The ``title`` field is already added to the class by ``ModelResource``
    # and populated off ``Note.title``. We'll use that title to build a
    # ``Note.slug`` as well.

    class Meta:
        queryset = Note.objects.all()

    def hydrate(self, bundle):
        # Don't change existing slugs.
        # In reality, this would be better implemented at the ``Note.save``
        # level, but is for demonstration.
        if not bundle.obj.pk:
            bundle.obj.slug = slugify(bundle.data['title'])

        return bundle
```

This method should return a `bundle`, whether it modifies the existing one or creates a whole new one. You can even remove any/all data from the `bundle.obj` if you wish.

hydrate_FOO

Data from the client may not map directly onto the data model or might need augmentation. This hook lets you take that data & convert it.

Note: The `FOO` here is not literal. Instead, it is a placeholder that should be replaced with the fieldname in question.

A simple example:

```
class MyResource(ModelResource):
    # The ``title`` field is already added to the class by ``ModelResource``
    # and populated off ``Note.title``. But we want lowercase titles...

    class Meta:
```

```

queryset = Note.objects.all()

def hydrate_title(self, bundle):
    return bundle.data['title'].lower()

```

The return value is updated in the `bundle.obj`.

Per-field hydrate

Each field (even custom `ApiField` subclasses) has its own `hydrate` method. If it knows how to access data (say, given the `attribute` kwarg), it will attempt to take data from the `bundle.data` & assign it on the data model.

The return value is put in the `bundle.obj` attribute for that fieldname.

6.7 Reverse “Relationships”

Unlike Django’s ORM, Tastypie does not automatically create reverse relations. This is because there is substantial technical complexity involved, as well as perhaps unintentionally exposing related data in an incorrect way to the end user of the API.

However, it is still possible to create reverse relations. Instead of handing the `ToOneField` or `ToManyField` a class, pass them a string that represents the full path to the desired class. Implementing a reverse relationship looks like so:

```

# myapp/api/resources.py
from tastypie import fields
from tastypie.resources import ModelResource
from myapp.models import Note, Comment

class NoteResource(ModelResource):
    comments = fields.ToManyField('myapp.api.resources.CommentResource', 'comments')

    class Meta:
        queryset = Note.objects.all()

class CommentResource(ModelResource):
    note = fields.ToOneField(NoteResource, 'notes')

    class Meta:
        queryset = Comment.objects.all()

```

Warning: Unlike Django, you can’t use just the class name (i.e. `'CommentResource'`), even if it’s in the same module. Tastypie (intentionally) lacks a construct like the `AppCache` which makes that sort of thing work in Django. Sorry.

Tastypie also supports self-referential relations. If you assume we added the appropriate self-referential `ForeignKey` to the `Note` model, implementing a similar relation in Tastypie would look like:

```

# myapp/api/resources.py
from tastypie import fields
from tastypie.resources import ModelResource
from myapp.models import Note

```

```
class NoteResource(ModelResource):
    sub_notes = fields.ToManyField('self', 'notes')

    class Meta:
        queryset = Note.objects.all()
```

6.8 Resource Options (AKA Meta)

The inner `Meta` class allows for class-level configuration of how the `Resource` should behave. The following options are available:

6.8.1 serializer

Controls which serializer class the `Resource` should use. Default is `tastypie.serializers.Serializer()`.

6.8.2 authentication

Controls which authentication class the `Resource` should use. Default is `tastypie.authentication.Authentication()`.

6.8.3 authorization

Controls which authorization class the `Resource` should use. Default is `tastypie.authorization.ReadOnlyAuthorization()`.

6.8.4 validation

Controls which validation class the `Resource` should use. Default is `tastypie.validation.Validation()`.

6.8.5 paginator_class

Controls which paginator class the `Resource` should use. Default is `tastypie.paginator.Paginator()`.

Note: This is different than the other options in that you supply a class rather than an instance. This is done because the `Paginator` has some per-request initialization options.

6.8.6 cache

Controls which cache class the `Resource` should use. Default is `tastypie.cache.NoCache()`.

6.8.7 throttle

Controls which throttle class the Resource should use. Default is `tastypie.throttle.BaseThrottle()`.

6.8.8 allowed_methods

Controls what list & detail REST methods the Resource should respond to. Default is `None`, which means delegate to the more specific `list_allowed_methods` & `detail_allowed_methods` options.

You may specify a list like `['get', 'post', 'put', 'delete', 'patch']` as a shortcut to prevent having to specify the other options.

6.8.9 list_allowed_methods

Controls what list REST methods the Resource should respond to. Default is `['get', 'post', 'put', 'delete', 'patch']`.

6.8.10 detail_allowed_methods

Controls what detail REST methods the Resource should respond to. Default is `['get', 'post', 'put', 'delete', 'patch']`.

6.8.11 limit

Controls what how many results the Resource will show at a time. Default is either the `API_LIMIT_PER_PAGE` setting (if provided) or 20 if not specified.

6.8.12 api_name

An override for the Resource to use when generating resource URLs. Default is `None`.

6.8.13 resource_name

An override for the Resource to use when generating resource URLs. Default is `None`.

If not provided, the Resource or `ModelResource` will attempt to name itself. This means a lowercase version of the classname preceding the word `Resource` if present (i.e. `SampleContentResource` would become `samplecontent`).

6.8.14 default_format

Specifies the default serialization format the Resource should use if one is not requested (usually by the `Accept` header or `format` GET parameter). Default is `application/json`.

6.8.15 filtering

Provides a list of fields that the `Resource` will accept client filtering on. Default is `{}`.

Keys should be the fieldnames as strings while values should be a list of accepted filter types.

6.8.16 ordering

Specifies the what fields the `Resource` should allow ordering on. Default is `[]`.

Values should be the fieldnames as strings. When provided to the `Resource` by the `order_by` GET parameter, you can specify either the `fieldname` (ascending order) or `-fieldname` (descending order).

6.8.17 object_class

Provides the `Resource` with the object that serves as the data source. Default is `None`.

In the case of `ModelResource`, this is automatically populated by the `queryset` option and is the model class.

6.8.18 queryset

Provides the `Resource` with the set of Django models to respond with. Default is `None`.

Unused by `Resource` but present for consistency.

Warning: If you place any callables in this, they'll only be evaluated once (when the `Meta` class is instantiated). This especially affects things that are date/time related. Please see the [:ref:cookbook](#) for a way around this.

6.8.19 fields

Controls what introspected fields the `Resource` should include. A whitelist of fields. Default is `[]`.

6.8.20 excludes

Controls what introspected fields the `Resource` should *NOT* include. A blacklist of fields. Default is `[]`.

6.8.21 include_resource_uri

Specifies if the `Resource` should include an extra field that displays the detail URL (within the api) for that resource. Default is `True`.

6.8.22 include_absolute_url

Specifies if the `Resource` should include an extra field that displays the `get_absolute_url` for that object (on the site proper). Default is `False`.

6.8.23 `always_return_data`

Specifies all HTTP methods (except DELETE) should return a serialized form of the data. Default is False.

If False, `HttpNoContent` (204) is returned on POST/PUT with an empty body & a `Location` header of where to request the full resource.

If True, `HttpAccepted` (202) is returned on POST/PUT with a body containing all the data in a serialized form.

6.9 Basic Filtering

`ModelResource` provides a basic Django ORM filter interface. Simply list the resource fields which you'd like to filter on and the allowed expression in a *filtering* property of your resource's Meta class:

```
from tastypie.constants import ALL, ALL_WITH_RELATIONS

class MyResource(ModelResource):
    class Meta:
        filtering = {
            "slug": ('exact', 'startswith',),
            "title": ALL,
        }
```

Valid filtering values are: Django ORM filters (e.g. `startswith`, `exact`, `lte`, etc. or the `ALL` or `ALL_WITH_RELATIONS` constants defined in `tastypie.constants`.

These filters will be extracted from URL query strings using the same double-underscore syntax as the Django ORM:

```
/api/v1/myresource/?slug=myslug
/api/v1/myresource/?slug__startswith=test
```

6.10 Advanced Filtering

If you need to filter things other than ORM resources or wish to apply additional constraints (e.g. text filtering using *django-haystack* <<http://haystacksearch.org>> rather than simple database queries) your `Resource` may define a custom `build_filters()` method which allows you to filter the queryset before processing a request:

```
from haystack.query import SearchQuerySet

class MyResource(Resource):
    def build_filters(self, filters=None):
        if filters is None:
            filters = {}

        orm_filters = super(MyResource, self).build_filters(filters)

        if "q" in filters:
            sqs = SearchQuerySet().auto_query(filters['q'])

            orm_filters["pk__in"] = [i.pk for i in sqs]

        return orm_filters
```

6.11 Resource Methods

Handles the data, request dispatch and responding to requests.

Serialization/deserialization is handled “at the edges” (i.e. at the beginning/end of the request/response cycle) so that everything internally is Python data structures.

This class tries to be non-model specific, so it can be hooked up to other data sources, such as search results, files, other data, etc.

6.11.1 `wrap_view`

`Resource.wrap_view` (*self*, *view*)

Wraps methods so they can be called in a more functional way as well as handling exceptions better.

Note that if `BadRequest` or an exception with a `response` attr are seen, there is special handling to either present a message back to the user or return the response traveling with the exception.

6.11.2 `base_urls`

`Resource.base_urls` (*self*)

The standard URLs this `Resource` should respond to. These include the list, detail, schema & multiple endpoints by default.

Should return a list of individual URLconf lines (**NOT** wrapped in `patterns`).

6.11.3 `override_urls`

`Resource.override_urls` (*self*)

A hook for adding your own URLs or overriding the default URLs. Useful for adding custom endpoints or overriding the built-in ones (from `base_urls`).

Should return a list of individual URLconf lines (**NOT** wrapped in `patterns`).

6.11.4 `urls`

`Resource.urls` (*self*)

Property

The endpoints this `Resource` responds to. A combination of `base_urls` & `override_urls`.

Mostly a standard URLconf, this is suitable for either automatic use when registered with an `Api` class or for including directly in a URLconf should you choose to.

6.11.5 `determine_format`

`Resource.determine_format` (*self*, *request*)

Used to determine the desired format.

Largely relies on `tastypie.utils.mime.determine_format` but here as a point of extension.

6.11.6 `serialize`

`Resource.serialize` (*self, request, data, format, options=None*)

Given a request, data and a desired format, produces a serialized version suitable for transfer over the wire.

Mostly a hook, this uses the `Serializer` from `Resource._meta`.

6.11.7 `deserialize`

`Resource.deserialize` (*self, request, data, format='application/json'*)

Given a request, data and a format, deserializes the given data.

It relies on the request properly sending a `CONTENT_TYPE` header, falling back to `application/json` if not provided.

Mostly a hook, this uses the `Serializer` from `Resource._meta`.

6.11.8 `alter_list_data_to_serialize`

`Resource.alter_list_data_to_serialize` (*self, request, data*)

A hook to alter list data just before it gets serialized & sent to the user.

Useful for restructuring/renaming aspects of the what's going to be sent.

Should accommodate for a list of objects, generally also including meta data.

6.11.9 `alter_detail_data_to_serialize`

`Resource.alter_detail_data_to_serialize` (*self, request, data*)

A hook to alter detail data just before it gets serialized & sent to the user.

Useful for restructuring/renaming aspects of the what's going to be sent.

Should accommodate for receiving a single bundle of data.

6.11.10 `alter_deserialized_list_data`

`Resource.alter_deserialized_list_data` (*self, request, data*)

A hook to alter list data just after it has been received from the user & gets deserialized.

Useful for altering the user data before any hydration is applied.

6.11.11 `alter_deserialized_detail_data`

`Resource.alter_deserialized_detail_data` (*self, request, data*)

A hook to alter detail data just after it has been received from the user & gets deserialized.

Useful for altering the user data before any hydration is applied.

6.11.12 `dispatch_list`

`Resource.dispatch_list` (*self*, *request*, ***kwargs*)

A view for handling the various HTTP methods (GET/POST/PUT/DELETE) over the entire list of resources.

Relies on `Resource.dispatch` for the heavy-lifting.

6.11.13 `dispatch_detail`

`Resource.dispatch_detail` (*self*, *request*, ***kwargs*)

A view for handling the various HTTP methods (GET/POST/PUT/DELETE) on a single resource.

Relies on `Resource.dispatch` for the heavy-lifting.

6.11.14 `dispatch`

`Resource.dispatch` (*self*, *request_type*, *request*, ***kwargs*)

Handles the common operations (allowed HTTP method, authentication, throttling, method lookup) surrounding most CRUD interactions.

6.11.15 `remove_api_resource_names`

`Resource.remove_api_resource_names` (*self*, *url_dict*)

Given a dictionary of regex matches from a URLconf, removes `api_name` and/or `resource_name` if found.

This is useful for converting URLconf matches into something suitable for data lookup. For example:

```
Model.objects.filter(**self.remove_api_resource_names(matches))
```

6.11.16 `method_check`

`Resource.method_check` (*self*, *request*, *allowed=None*)

Ensures that the HTTP method used on the request is allowed to be handled by the resource.

Takes an `allowed` parameter, which should be a list of lowercase HTTP methods to check against. Usually, this looks like:

```
# The most generic lookup.
self.method_check(request, self._meta.allowed_methods)

# A lookup against what's allowed for list-type methods.
self.method_check(request, self._meta.list_allowed_methods)

# A useful check when creating a new endpoint that only handles
# GET.
self.method_check(request, ['get'])
```

6.11.17 `is_authorized`

`Resource.is_authorized` (*self*, *request*, *object=None*)

Handles checking of permissions to see if the user has authorization to GET, POST, PUT, or DELETE this resource. If `object` is provided, the authorization backend can apply additional row-level permissions checking.

6.11.18 `is_authenticated`

`Resource.is_authenticated` (*self*, *request*)

Handles checking if the user is authenticated and dealing with unauthenticated users.

Mostly a hook, this uses class assigned to `authentication` from `Resource._meta`.

6.11.19 `throttle_check`

`Resource.throttle_check` (*self*, *request*)

Handles checking if the user should be throttled.

Mostly a hook, this uses class assigned to `throttle` from `Resource._meta`.

6.11.20 `log_throttled_access`

`Resource.log_throttled_access` (*self*, *request*)

Handles the recording of the user's access for throttling purposes.

Mostly a hook, this uses class assigned to `throttle` from `Resource._meta`.

6.11.21 `build_bundle`

`Resource.build_bundle` (*self*, *obj=None*, *data=None*, *request=None*)

Given either an object, a data dictionary or both, builds a `Bundle` for use throughout the `dehydrate/hydrate` cycle.

If no object is provided, an empty object from `Resource._meta.object_class` is created so that attempts to access `bundle.obj` do not fail.

6.11.22 `build_filters`

`Resource.build_filters` (*self*, *filters=None*)

Allows for the filtering of applicable objects.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.23 `apply_sorting`

`Resource.apply_sorting` (*self*, *obj_list*, *options=None*)

Allows for the sorting of objects being returned.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.24 `get_resource_uri`

`Resource.get_resource_uri` (*self*, *bundle_or_obj*)

This needs to be implemented at the user level.

A `return reverse("api_dispatch_detail", kwargs={'resource_name': self.resource_name, 'pk': object.id})` should be all that would be needed.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.25 `get_resource_list_uri`

`Resource.get_resource_list_uri` (*self*)

Returns a URL specific to this resource's list endpoint.

6.11.26 `get_via_uri`

`Resource.get_via_uri` (*self*, *uri*, *request=None*)

This pulls apart the salient bits of the URI and populates the resource via a `obj_get`.

Optionally accepts a `request`.

If you need custom behavior based on other portions of the URI, simply override this method.

6.11.27 `full_dehydrate`

`Resource.full_dehydrate` (*self*, *bundle*)

Given a bundle with an object instance, extract the information from it to populate the resource.

6.11.28 `dehydrate`

`Resource.dehydrate` (*self*, *bundle*)

A hook to allow a final manipulation of data once all fields/methods have built out the dehydrated data.

Useful if you need to access more than one dehydrated field or want to annotate on additional data.

Must return the modified bundle.

6.11.29 `full_hydrate`

`Resource.full_hydrate` (*self*, *bundle*)

Given a populated bundle, distill it and turn it back into a full-fledged object instance.

6.11.30 `hydrate`

`Resource.hydrate` (*self*, *bundle*)

A hook to allow a final manipulation of data once all fields/methods have built out the hydrated data.

Useful if you need to access more than one hydrated field or want to annotate on additional data.

Must return the modified bundle.

6.11.31 `hydrate_m2m`

`Resource.hydrate_m2m` (*self*, *bundle*)

Populate the ManyToMany data on the instance.

6.11.32 `build_schema`

`Resource.build_schema` (*self*)

Returns a dictionary of all the fields on the resource and some properties about those fields.

Used by the `schema/` endpoint to describe what will be available.

6.11.33 `dehydrate_resource_uri`

`Resource.dehydrate_resource_uri` (*self*, *bundle*)

For the automatically included `resource_uri` field, dehydrate the URI for the given bundle.

Returns empty string if no URI can be generated.

6.11.34 `generate_cache_key`

`Resource.generate_cache_key` (*self*, **args*, ***kwargs*)

Creates a unique-enough cache key.

This is based off the current `api_name/resource_name/args/kwags`.

6.11.35 `get_object_list`

`Resource.get_object_list` (*self*, *request*)

A hook to allow making returning the list of available objects.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.36 `apply_authorization_limits`

`Resource.apply_authorization_limits` (*self*, *request*, *object_list*)

Allows the `Authorization` class to further limit the object list. Also a hook to customize per `Resource`.

Calls `Authorization.apply_limits` if available.

6.11.37 `can_create`

`Resource.can_create` (*self*)

Checks to ensure `post` is within `allowed_methods`.

6.11.38 `can_update`

`Resource.can_update` (*self*)

Checks to ensure `put` is within `allowed_methods`.

Used when hydrating related data.

6.11.39 `can_delete`

`Resource.can_delete` (*self*)

Checks to ensure `delete` is within `allowed_methods`.

6.11.40 `apply_filters`

`Resource.apply_filters` (*self*, *request*, *applicable_filters*)

A hook to alter how the filters are applied to the object list.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.41 `obj_get_list`

`Resource.obj_get_list` (*self*, *request=None*, ***kwargs*)

Fetches the list of objects available on the resource.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.42 `cached_obj_get_list`

`Resource.cached_obj_get_list` (*self*, *request=None*, ***kwargs*)

A version of `obj_get_list` that uses the cache as a means to get commonly-accessed data faster.

6.11.43 `obj_get`

`Resource.obj_get` (*self*, *request=None*, ***kwargs*)

Fetches an individual object on the resource.

This needs to be implemented at the user level. If the object can not be found, this should raise a `NotFound` exception.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.44 `cached_obj_get`

`Resource.cached_obj_get` (*self*, *request=None*, ***kwargs*)

A version of `obj_get` that uses the cache as a means to get commonly-accessed data faster.

6.11.45 `obj_create`

`Resource.obj_create` (*self*, *bundle*, *request=None*, ***kwargs*)

Creates a new object based on the provided data.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.46 `obj_update`

`Resource.obj_update` (*self*, *bundle*, *request=None*, ***kwargs*)

Updates an existing object (or creates a new object) based on the provided data.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.47 `obj_delete_list`

`Resource.obj_delete_list` (*self*, *request=None*, ***kwargs*)

Deletes an entire list of objects.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.48 `obj_delete`

`Resource.obj_delete` (*self*, *request=None*, ***kwargs*)

Deletes a single object.

This needs to be implemented at the user level.

`ModelResource` includes a full working version specific to Django's `Models`.

6.11.49 `create_response`

`Resource.create_response` (*self*, *request*, *data*, *response_class=HttpResponse*, ***response_kwargs*)

Extracts the common “which-format/serialize/return-response” cycle.

Mostly a useful shortcut/hook.

6.11.50 `is_valid`

`Resource.is_valid` (*self*, *bundle*, *request=None*)

Handles checking if the data provided by the user is valid.

Mostly a hook, this uses class assigned to `validation` from `Resource._meta`.

If validation fails, an error is raised with the error messages serialized inside it.

6.11.51 `rollback`

`Resource.rollback` (*self*, *bundles*)

Given the list of bundles, delete all objects pertaining to those bundles.

This needs to be implemented at the user level. No exceptions should be raised if possible.

`ModelResource` includes a full working version specific to Django’s `Models`.

6.11.52 `get_list`

`Resource.get_list` (*self*, *request*, ***kwargs*)

Returns a serialized list of resources.

Calls `obj_get_list` to provide the data, then handles that result set and serializes it.

Should return a `HttpResponse` (200 OK).

6.11.53 `get_detail`

`Resource.get_detail` (*self*, *request*, ***kwargs*)

Returns a single serialized resource.

Calls `cached_obj_get/obj_get` to provide the data, then handles that result set and serializes it.

Should return a `HttpResponse` (200 OK).

6.11.54 `put_list`

`Resource.put_list` (*self*, *request*, ***kwargs*)

Replaces a collection of resources with another collection.

Calls `delete_list` to clear out the collection then `obj_create` with the provided the data to create the new collection.

Return `HttpNoContent` (204 No Content) if `Meta.always_return_data = False` (default).

Return `HttpAccepted` (202 Accepted) if `Meta.always_return_data = True`.

6.11.55 `put_detail`

Resource.`put_detail` (*self*, *request*, ***kwargs*)

Either updates an existing resource or creates a new one with the provided data.

Calls `obj_update` with the provided data first, but falls back to `obj_create` if the object does not already exist.

If a new resource is created, return `HttpCreated` (201 Created). If `Meta.always_return_data = True`, there will be a populated body of serialized data.

If an existing resource is modified and `Meta.always_return_data = False` (default), return `HttpNoContent` (204 No Content). If an existing resource is modified and `Meta.always_return_data = True`, return `HttpAccepted` (202 Accepted).

6.11.56 `post_list`

Resource.`post_list` (*self*, *request*, ***kwargs*)

Creates a new resource/object with the provided data.

Calls `obj_create` with the provided data and returns a response with the new resource's location.

If a new resource is created, return `HttpCreated` (201 Created). If `Meta.always_return_data = True`, there will be a populated body of serialized data.

6.11.57 `post_detail`

Resource.`post_detail` (*self*, *request*, ***kwargs*)

Creates a new subcollection of the resource under a resource.

This is not implemented by default because most people's data models aren't self-referential.

If a new resource is created, return `HttpCreated` (201 Created).

6.11.58 `delete_list`

Resource.`delete_list` (*self*, *request*, ***kwargs*)

Destroys a collection of resources/objects.

Calls `obj_delete_list`.

If the resources are deleted, return `HttpNoContent` (204 No Content).

6.11.59 `delete_detail`

Resource.`delete_detail` (*self*, *request*, ***kwargs*)

Destroys a single resource/object.

Calls `obj_delete`.

If the resource is deleted, return `HttpNoContent` (204 No Content). If the resource did not exist, return `HttpNotFound` (404 Not Found).

6.11.60 patch_list

Resource.**patch_list** (*self*, *request*, ***kwargs*)

Updates a collection in-place.

The exact behavior of PATCH to a list resource is still the matter of some debate in REST circles, and the PATCH RFC isn't standard. So the behavior this method implements (described below) is something of a stab in the dark. It's mostly cribbed from GData, with a smattering of ActiveResource-isms and maybe even an original idea or two.

The PATCH format is one that's similar to the response returned from a GET on a list resource:

```
{
  "objects": [{object}, {object}, ...],
  "deleted_objects": ["URI", "URI", "URI", ...],
}
```

For each object in `objects`:

- If the dict does not have a `resource_uri` key then the item is considered “new” and is handled like a POST to the resource list.
- If the dict has a `resource_uri` key and the `resource_uri` refers to an existing resource then the item is a update; it's treated like a PATCH to the corresponding resource detail.
- If the dict has a `resource_uri` but the resource *doesn't* exist, then this is considered to be a create-via-PUT.

Each entry in `deleted_objects` refers to a resource URI of an existing resource to be deleted; each is handled like a DELETE to the relevant resource.

In any case:

- If there's a resource URI it *must* refer to a resource of this type. It's an error to include a URI of a different resource.
- PATCH is all or nothing. If a single sub-operation fails, the entire request will fail and all resources will be rolled back.

6.11.61 patch_detail

Resource.**patch_detail** (*self*, *request*, ***kwargs*)

Updates a resource in-place.

Calls `obj_update`.

If the resource is updated, return `HttpAccepted` (202 Accepted). If the resource did not exist, return `HttpNotFound` (404 Not Found).

6.11.62 get_schema

Resource.**get_schema** (*self*, *request*, ***kwargs*)

Returns a serialized form of the schema of the resource.

Calls `build_schema` to generate the data. This method only responds to HTTP GET.

Should return a `HttpResponse` (200 OK).

6.11.63 `get_multiple`

`Resource.get_multiple` (*self*, *request*, ***kwargs*)

Returns a serialized list of resources based on the identifiers from the URL.

Calls `obj_get` to fetch only the objects requested. This method only responds to HTTP GET.

Should return a `HttpResponse` (200 OK).

6.12 ModelResource Methods

A subclass of `Resource` designed to work with Django's `Models`.

This class will introspect a given `Model` and build a field list based on the fields found on the model (excluding relational fields).

Given that it is aware of Django's ORM, it also handles the CRUD data operations of the resource.

6.12.1 `should_skip_field`

`ModelResource.should_skip_field` (*cls*, *field*)

Class method

Given a Django model field, return if it should be included in the contributed `ApiFields`.

6.12.2 `api_field_from_django_field`

`ModelResource.api_field_from_django_field` (*cls*, *f*, *default=CharField*)

Class method

Returns the field type that would likely be associated with each Django type.

6.12.3 `get_fields`

`ModelResource.get_fields` (*cls*, *fields=None*, *excludes=None*)

Class method

Given any explicit fields to include and fields to exclude, add additional fields based on the associated model.

6.12.4 `check_filtering`

`ModelResource.check_filtering` (*self*, *field_name*, *filter_type='exact'*, *filter_bits=None*)

Given a field name, a optional filter type and an optional list of additional relations, determine if a field can be filtered on.

If a filter does not meet the needed conditions, it should raise an `InvalidFilterError`.

If the filter meets the conditions, a list of attribute names (not field names) will be returned.

6.12.5 build_filters

`ModelResource.build_filters` (*self*, *filters=None*)

Given a dictionary of filters, create the necessary ORM-level filters.

Keys should be resource fields, **NOT** model fields.

Valid values are either a list of Django filter types (i.e. [`'startswith'`, `'exact'`, `'lte'`]), the `ALL` constant or the `ALL_WITH_RELATIONS` constant.

At the declarative level:

```
filtering = {
    'resource_field_name': ['exact', 'startswith', 'endswith', 'contains'],
    'resource_field_name_2': ['exact', 'gt', 'gte', 'lt', 'lte', 'range'],
    'resource_field_name_3': ALL,
    'resource_field_name_4': ALL_WITH_RELATIONS,
    ...
}
```

Accepts the filters as a dict. None by default, meaning no filters.

6.12.6 apply_sorting

`ModelResource.apply_sorting` (*self*, *obj_list*, *options=None*)

Given a dictionary of options, apply some ORM-level sorting to the provided `QuerySet`.

Looks for the `order_by` key and handles either ascending (just the field name) or descending (the field name with a `-` in front).

The field name should be the resource field, **NOT** model field.

6.12.7 apply_filters

`ModelResource.apply_filters` (*self*, *request*, *applicable_filters*)

An ORM-specific implementation of `apply_filters`.

The default simply applies the `applicable_filters` as `**kwargs`, but should make it possible to do more advanced things.

6.12.8 get_object_list

`ModelResource.get_object_list` (*self*, *request*)

A ORM-specific implementation of `get_object_list`.

Returns a `QuerySet` that may have been limited by other overrides.

6.12.9 obj_get_list

`ModelResource.obj_get_list` (*self*, *filters=None*, ***kwargs*)

A ORM-specific implementation of `obj_get_list`.

Takes an optional `filters` dictionary, which can be used to narrow the query.

6.12.10 `obj_get`

`ModelResource.obj_get` (*self*, ***kwargs*)

A ORM-specific implementation of `obj_get`.

Takes optional `kwargs`, which are used to narrow the query to find the instance.

6.12.11 `obj_create`

`ModelResource.obj_create` (*self*, *bundle*, ***kwargs*)

A ORM-specific implementation of `obj_create`.

6.12.12 `obj_update`

`ModelResource.obj_update` (*self*, *bundle*, ***kwargs*)

A ORM-specific implementation of `obj_update`.

6.12.13 `obj_delete_list`

`ModelResource.obj_delete_list` (*self*, ***kwargs*)

A ORM-specific implementation of `obj_delete_list`.

Takes optional `kwargs`, which can be used to narrow the query.

6.12.14 `obj_delete`

`ModelResource.obj_delete` (*self*, ***kwargs*)

A ORM-specific implementation of `obj_delete`.

Takes optional `kwargs`, which are used to narrow the query to find the instance.

6.12.15 `rollback`

`ModelResource.rollback` (*self*, *bundles*)

A ORM-specific implementation of `rollback`.

Given the list of bundles, delete all models pertaining to those bundles.

6.12.16 `save_related`

`ModelResource.save_related` (*self*, *bundle*)

Handles the saving of related non-M2M data.

Calling assigning `child.parent = parent` & then calling `Child.save` isn't good enough to make sure the parent is saved.

To get around this, we go through all our related fields & call `save` on them if they have related, non-M2M data. M2M data is handled by the `ModelResource.save_m2m` method.

6.12.17 `save_m2m`

`ModelResource.save_m2m` (*self*, *bundle*)

Handles the saving of related M2M data.

Due to the way Django works, the M2M data must be handled after the main instance, which is why this isn't a part of the main `save` bits.

Currently slightly inefficient in that it will clear out the whole relation and recreate the related data as needed.

6.12.18 `get_resource_uri`

`ModelResource.get_resource_uri` (*self*, *bundle_or_obj*)

Handles generating a resource URI for a single resource.

Uses the model's `pk` in order to create the URI.

BUNDLES

7.1 What Are Bundles?

Bundles are a small abstraction that allow Tastypie to pass data between resources. This allows us not to depend on passing `request` to every single method (especially in places where this would be overkill). It also allows resources to work with data coming into the application paired together with an unsaved instance of the object in question. Finally, it aids in keeping Tastypie more thread-safe.

Think of it as package of user data & an object instance (either of which are optionally present).

7.2 Attributes

All data within a bundle can be optional, especially depending on how it's being used. If you write custom code using `Bundle`, make sure appropriate guards are in place.

7.2.1 `obj`

Either a Python object or `None`.

Usually a Django model, though it may/may not have been saved already.

7.2.2 `data`

Always a plain Python dictionary of data. If not provided, it will be empty.

7.2.3 `request`

Either the Django `request` that's part of the issued request or an empty `HttpRequest` if it wasn't provided.

7.2.4 `related_obj`

Either another "parent" Python object or `None`.

Useful when handling one-to-many relations. Used in conjunction with `related_name`.

7.2.5 `related_name`

Either a Python string name of an attribute or `None`.

Useful when handling one-to-many relations. Used in conjunction with `related_obj`.

API

In terms of a REST-style architecture, the “api” is a collection of resources. In TastyPie, the `Api` gathers together the `Resources` & provides a nice way to use them as a set. It handles many of the `URLconf` details for you, provides a helpful “top-level” view to show what endpoints are available & some extra URL resolution juice.

8.1 Quick Start

A sample api definition might look something like (usually located in a `URLconf`):

```
from tastypie.api import Api
from myapp.api.resources import UserResource, EntryResource

v1_api = Api(api_name='v1')
v1_api.register(UserResource())
v1_api.register(EntryResource())

# Standard bits...
urlpatterns = patterns('',
    (r'^api/', include(v1_api.urls)),
)
```

8.2 Api Methods

Implements a registry to tie together the various resources that make up an API.

Especially useful for navigation, HATEOAS and for providing multiple versions of your API.

Optionally supplying `api_name` allows you to name the API. Generally, this is done with version numbers (i.e. `v1`, `v2`, etc.) but can be named any string.

8.2.1 register

`Api.register(self, resource, canonical=True):`

Registers an instance of a `Resource` subclass with the API.

Optionally accept a `canonical` argument, which indicates that the resource being registered is the canonical variant. Defaults to `True`.

8.2.2 unregister

Api.unregister(self, resource_name):

If present, unregisters a resource from the API.

8.2.3 canonical_resource_for

Api.canonical_resource_for(self, resource_name):

Returns the canonical resource for a given `resource_name`.

override_urls

Api.override_urls(self):

A hook for adding your own URLs or overriding the default URLs. Useful for adding custom endpoints or overriding the built-in ones.

Should return a list of individual URLconf lines (**NOT** wrapped in `patterns`).

8.2.4 urls

Api.urls(self):

Property

Provides URLconf details for the `Api` and all registered `Resources` beneath it.

8.2.5 top_level

Api.top_level(self, request, api_name=None):

A view that returns a serialized list of all resources registers to the `Api`. Useful for discovery.

RESOURCE FIELDS

When designing an API, an important component is defining the representation of the data you're presenting. Like Django models, you can control the representation of a `Resource` using fields. There are a variety of fields for various types of data.

9.1 Quick Start

For the impatient:

```
import datetime
from tastypie import fields
from tastypie.resources import Resource
from myapp.api.resources import ProfileResource, NoteResource

class PersonResource(Resource):
    name = fields.CharField(attribute='name')
    age = fields.IntegerField(attribute='years_old', null=True)
    created = fields.DateTimeField(readonly=True, help_text='When the person was created', default=datetime.datetime.now)
    is_active = fields.BooleanField(default=True)
    profile = fields.ToOneField(ProfileResource, 'profile')
    notes = fields.ToManyField(NoteResource, 'notes', full=True)
```

9.2 Standard Data Fields

All standard data fields have a common base class `ApiField` which handles the basic implementation details.

Note: You should not use the `ApiField` class directly. Please use one of the subclasses that is more correct for your data.

9.2.1 Common Field Options

All `ApiField` objects accept the following options.

attribute

`ApiField.attribute`

A string naming an instance attribute of the object wrapped by the Resource. The attribute will be accessed during the `dehydrate` or or written during the `hydrate`.

Defaults to `None`, meaning data will be manually accessed.

default

`ApiField.default`

Provides default data when the object being `dehydrated/hydrated` has no data on the field.

Defaults to `tastypie.fields.NOT_PROVIDED`.

null

`ApiField.null`

Indicates whether or not a `None` is allowable data on the field. Defaults to `False`.

blank

`ApiField.blank`

Indicates whether or not data may be omitted on the field. Defaults to `False`.

This is useful for allowing the user to omit data that you can populate based on the request, such as the `user` or `site` to associate a record with.

readonly

`ApiField.readonly`

Indicates whether the field is used during the `hydrate` or not. Defaults to `False`.

unique

`ApiField.unique`

Indicates whether the field is a unique identifier for the object.

help_text

`ApiField.help_text`

A human-readable description of the field exposed at the schema level. Defaults to the per-Field definition.

9.2.2 Field Types

9.2.3 BooleanField

A boolean field.

Covers both `models.BooleanField` and `models.NullBooleanField`.

9.2.4 CharField

A text field of arbitrary length.

Covers both `models.CharField` and `models.TextField`.

9.2.5 DateField

A date field.

9.2.6 DateTimeField

A datetime field.

9.2.7 DecimalField

A decimal field.

9.2.8 DictField

A dictionary field.

9.2.9 FileField

A file-related field.

Covers both `models.FileField` and `models.ImageField`.

9.2.10 FloatField

A floating point field.

9.2.11 IntegerField

An integer field.

Covers `models.IntegerField`, `models.PositiveIntegerField`, `models.PositiveSmallIntegerField` and `models.SmallIntegerField`.

9.2.12 ListField

A list field.

9.2.13 TimeField

A time field.

9.3 Relationship Fields

Provides access to data that is related within the database.

The `RelatedField` base class is not intended for direct use but provides functionality that `ToOneField` and `ToManyField` build upon.

The contents of this field actually point to another `Resource`, rather than the related object. This allows the field to represent its data in different ways.

The abstractions based around this are “leaky” in that, unlike the other fields provided by `tastypie`, these fields don’t handle arbitrary objects very well. The subclasses use Django’s ORM layer to make things go, though there is no ORM-specific code at this level.

9.3.1 Common Field Options

In addition to the common attributes for all `ApiField`, relationship fields accept the following.

`to`

`RelatedField.to`

The `to` argument should point to a `Resource` class, NOT to a `Model`. Required.

`full`

`RelatedField.full`

Indicates how the related `Resource` will appear post-dehydrate. If `False`, the related `Resource` will appear as a URL to the endpoint of that resource. If `True`, the result of the sub-resource’s `dehydrate` will be included in full.

`related_name`

`RelatedField.related_name`

Used to help automatically populate reverse relations when creating data. Defaults to `None`.

In order for this option to work correctly, there must be a field on the other `Resource` with this as an attribute/`instance_name`. Usually this just means adding a reflecting `ToOneField` pointing back.

Example:

```

class EntryResource(ModelResource):
    authors = fields.ToManyField('path.to.api.resources.AuthorResource', 'author_set', related_name=

class Meta:
    queryset = Entry.objects.all()
    resource_name = 'entry'

class AuthorResource(ModelResource):
    entry = fields.ToOneField(EntryResource, 'entry')

class Meta:
    queryset = Author.objects.all()
    resource_name = 'author'

```

9.3.2 Field Types

ToOneField

Provides access to related data via foreign key.

This subclass requires Django's ORM layer to work properly.

OneToOneField

An alias to ToOneField for those who prefer to mirror `django.db.models`.

ForeignKey

An alias to ToOneField for those who prefer to mirror `django.db.models`.

ToManyField

Provides access to related data via a join table.

This subclass requires Django's ORM layer to work properly.

This field also has special behavior when dealing with `attribute` in that it can take a callable. For instance, if you need to filter the reverse relation, you can do something like:

```
subjects = fields.ToManyField(SubjectResource, attribute=lambda bundle: Subject.objects.filter(notes=
```

Note that the `hydrate` portions of this field are quite different than any other field. `hydrate_m2m` actually handles the data and relations. This is due to the way Django implements M2M relationships.

ManyToManyField

An alias to ToManyField for those who prefer to mirror `django.db.models`.

OneToManyField

An alias to ToManyField for those who prefer to mirror `django.db.models`.

CACHING

When adding an API to your site, it's important to understand that most consumers of the API will not be people, but instead machines. This means that the traditional “fetch-read-click” cycle is no longer measured in minutes but in seconds or milliseconds.

As such, caching is a very important part of the deployment of your API. Tastypie ships with two classes to make working with caching easier. These caches store at the object level, reducing access time on the database.

However, it's worth noting that these do *NOT* cache serialized representations. For heavy traffic, we'd encourage the use of a caching proxy, especially [Varnish](#), as it shines under this kind of usage. It's far faster than Django views and already neatly handles most situations.

10.1 Usage

Using these classes is simple. Simply provide them (or your own class) as a `Meta` option to the `Resource` in question. For example:

```
from django.contrib.auth.models import User
from tastypie.cache import SimpleCache
from tastypie.resources import ModelResource

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']
        # Add it here.
        cache = SimpleCache()
```

10.2 Caching Options

Tastypie ships with the following `Cache` classes:

10.2.1 NoCache

The no-op cache option, this does no caching but serves as an api-compatible plug. Very useful for development.

10.2.2 SimpleCache

This option does basic object caching, attempting to find the object in the cache & writing the object to the cache. It uses Django's current `CACHE_BACKEND` to store cached data.

10.3 Implementing Your Own Cache

Implementing your own `Cache` class is as simple as subclassing `NoCache` and overriding the `get` & `set` methods. For example, a json-backed cache might look like:

```
import json
from django.conf import settings
from tastypie.cache import NoCache

class JSONCache(NoCache):
    def _load(self):
        data_file = open(settings.TASTYPIE_JSON_CACHE, 'r')
        return json.load(data_file)

    def _save(self, data):
        data_file = open(settings.TASTYPIE_JSON_CACHE, 'w')
        return json.dump(data, data_file)

    def get(self, key):
        data = self._load()
        return data.get(key, None)

    def set(self, key, value, timeout=60):
        data = self._load()
        data[key] = value
        self._save(data)
```

Note that this is *NOT* necessarily an optimal solution, but is simply demonstrating how one might go about implementing your own `Cache`.

VALIDATION

Validation allows you to ensure that the data being submitted by the user is appropriate for storage. This can range from simple type checking on up to complex validation that compares different fields together.

If the data is valid, an empty dictionary is returned and processing continues as normal. If the data is invalid, a dictionary of error messages (keys being the field names, values being a list of error messages). This will be immediately returned to the user, serialized in the format they requested.

11.1 Usage

Using these classes is simple. Simply provide them (or your own class) as a `Meta` option to the `Resource` in question. For example:

```
from django.contrib.auth.models import User
from tastypie.validation import Validation
from tastypie.resources import ModelResource

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']
        # Add it here.
        validation = Validation()
```

11.2 Validation Options

Tastypie ships with the following `Validation` classes:

11.2.1 Validation

The no-op validation option, the data submitted is always considered to be valid.

This is the default class hooked up to `Resource/ModelResource`.

11.2.2 FormValidation

A more complex form of validation, this class accepts a `form_class` argument to its constructor. You supply a Django Form (or `ModelForm`, though `save` will never get called) and Tastypie will verify the data in the `Bundle` against the form.

This class **DOES NOT** alter the data sent, only verifies it. If you want to alter the data, please use the `CleanDataFormValidation` class instead.

Warning: Data in the bundle must line up with the fieldnames in the `Form`. If they do not, you'll need to either munge the data or change your form.

Usage looks like:

```
from django import forms

class NoteForm(forms.Form):
    title = forms.CharField(max_length=100)
    slug = forms.CharField(max_length=50)
    content = forms.CharField(required=False, widget=forms.Textarea)
    is_active = forms.BooleanField()

form = FormValidation(form_class=NoteForm)
```

11.2.3 CleanedDataFormValidation

Similar to the `FormValidation` class, this uses a Django Form to handle validation. **However**, it will use the `form.cleaned_data` to replace the bundle data sent by user! Usage is identical to `FormValidation`.

11.3 Implementing Your Own Validation

Implementing your own `Validation` classes is a simple process. The constructor can take whatever `**kwargs` it needs (if any). The only other method to implement is the `is_valid` method:

```
from tastypie.validation import Validation

class AwesomeValidation(Validation):
    def is_valid(self, bundle, request=None):
        if not bundle.data:
            return {'__all__': 'Not quite what I had in mind.'}

        errors = {}

        for key, value in bundle.data.items():
            if not isinstance(value, basestring):
                continue

            if not 'awesome' in value:
                errors[key] = ['NOT ENOUGH AWESOME. NEEDS MORE.']

        return errors
```

Under this validation, every field that's a string is checked for the word 'awesome'. If it's not in the string, it's an error.

AUTHENTICATION / AUTHORIZATION

Authentication & authorization make up the components needed to verify that a certain user has access to the API and what they can do with it.

Authentication answers the question “can they see this data?” This usually involves requiring credentials, such as an API key or username/password.

Authorization answers the question “what objects can they modify?” This usually involves checking permissions, but is open to other implementations.

12.1 Usage

Using these classes is simple. Simply provide them (or your own class) as a `Meta` option to the `Resource` in question. For example:

```
from django.contrib.auth.models import User
from tastypie.authentication import BasicAuthentication
from tastypie.authorization import DjangoAuthorization
from tastypie.resources import ModelResource

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']
        # Add it here.
        authentication = BasicAuthentication()
        authorization = DjangoAuthorization()
```

12.2 Authentication Options

Tastypie ships with the following `Authentication` classes:

12.2.1 Authentication

The no-op authentication option, the client is always allowed through. Very useful for development and read-only APIs.

12.2.2 BasicAuthentication

This authentication scheme uses HTTP Basic Auth to check a user's credentials. The username is their `django.contrib.auth.models.User` username (assuming it is present) and their password should also correspond to that entry.

Warning: If you're using Apache & `mod_wsgi`, you will need to enable `WSGIPassAuthorization On`. See [this post](#) for details.

12.2.3 ApiKeyAuthentication

As an alternative to requiring sensitive data like a password, the `ApiKeyAuthentication` allows you to collect just username & a machine-generated api key. Tastypie ships with a special `Model` just for this purpose, so you'll need to ensure `tastypie` is in `INSTALLED_APPS`.

Tastypie includes a signal function you can use to auto-create `ApiKey` objects. Hooking it up looks like:

```
from django.contrib.auth.models import User
from django.db import models
from tastypie.models import create_api_key

models.signals.post_save.connect(create_api_key, sender=User)
```

12.2.4 DigestAuthentication

This authentication scheme uses HTTP Digest Auth to check a user's credentials. The username is their `django.contrib.auth.models.User` username (assuming it is present) and their password should be their machine-generated api key. As with `ApiKeyAuthentication`, `tastypie` should be included in `INSTALLED_APPS`.

Warning: If you're using Apache & `mod_wsgi`, you will need to enable `WSGIPassAuthorization On`. See [this post](#) for details (even though it only mentions Basic auth).

12.2.5 OAuthAuthentication

Handles OAuth, which checks a user's credentials against a separate service. Currently verifies against OAuth 1.0a services.

This does *NOT* provide OAuth authentication in your API, strictly consumption.

Warning: If you're used to in-browser OAuth flow (click a "Sign In" button, get redirected, login on remote service, get redirected back), this isn't the same. Most prominently, expecting that would cause API clients to have to use tools like `mechanize` to fill in forms, which would be difficult.

This authentication expects that you're already followed some sort of OAuth flow & that the credentials (Nonce/token/etc) are simply being passed to it. It merely checks that the credentials are valid. No requests are made to remote services as part of this authentication class.

12.3 Authorization Options

Tastypie ships with the following `Authorization` classes:

12.3.1 Authorization

The no-op authorization option, no permissions checks are performed.

Warning: This is a potentially dangerous option, as it means *ANY* recognized user can modify *ANY* data they encounter in the API. Be careful who you trust.

12.3.2 ReadOnlyAuthorization

This authorization class only permits reading data, regardless of what the `Resource` might think is allowed. This is the default `Authorization` class and the safe option.

12.3.3 DjangoAuthorization

The most advanced form of authorization, this checks the permission a user has granted to them (via `django.contrib.auth.models.Permission`). In conjunction with the admin, this is a very effective means of control.

12.4 Implementing Your Own Authentication/Authorization

Implementing your own `Authentication/Authorization` classes is a simple process. `Authentication` has two methods to override (one of which is optional but recommended to be customized) and `Authorization` has just one required method and one optional method:

```
from tastypie.authentication import Authentication
from tastypie.authorization import Authorization

class SillyAuthentication(Authentication):
    def is_authenticated(self, request, **kwargs):
        if 'daniel' in request.user.username:
            return True

        return False

    # Optional but recommended
    def get_identifier(self, request):
        return request.user.username

class SillyAuthorization(Authorization):
    def is_authorized(self, request, object=None):
        if request.user.date_joined.year == 2010:
            return True
        else:
            return False

    # Optional but useful for advanced limiting, such as per user.
    def apply_limits(self, request, object_list):
        if request and hasattr(request, 'user'):
            return object_list.filter(author__username=request.user.username)

        return object_list.none()
```

Under this scheme, only users with 'daniel' in their username will be allowed in, and only those who joined the site in 2010 will be allowed to affect data.

If the optional `apply_limits` method is included, each user that fits the above criteria will only be able to access their own records.

SERIALIZATION

Serialization can be one of the most contentious areas of an API. Everyone has their own requirements, their own preferred output format & the desire to have control over what is returned.

As a result, TastyPie ships with a serializer that tries to meet the basic needs of most use cases, and the flexibility to go outside of that when you need to.

The default `Serializer` supports the following formats:

- json
- jsonp
- xml
- yaml
- html
- plist (see <http://explorapp.com/biplist/>)

13.1 Usage

Using this class is simple. It is the default option on all `Resource` classes unless otherwise specified. The following code is a no-op, but demonstrate how you could use your own serializer:

```
from django.contrib.auth.models import User
from tastypie.resources import ModelResource
from tastypie.serializers import Serializer

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']
        # Add it here.
        serializer = Serializer()
```

13.2 Implementing Your Own Serializer

There are several different use cases here. We'll cover simple examples of wanting a tweaked format & adding a different format.

To tweak a format, simply override its `to_<format>` & `from_<format>` methods. So adding the server time to all output might look like so:

```
import time
from tastypie.serializers import Serializer

class CustomJSONSerializer(Serializer):
    def to_json(self, data, options=None):
        options = options or {}

        data = self.to_simple(data, options)

        # Add in the current time.
        data['requested_time'] = time.time()

        return simplejson.dumps(data, cls=json.DjangoJSONEncoder, sort_keys=True)

    def from_json(self, content):
        data = simplejson.loads(content)

        if 'requested_time' in data:
            # Log the request here...
            pass

        return data
```

In the case of adding a different format, let's say you want to add a CSV output option to the existing set. Your `Serializer` subclass might look like:

```
import csv
import StringIO
from tastypie.serializers import Serializer

class CSVSerializer(Serializer):
    formats = ['json', 'jsonp', 'xml', 'yaml', 'html', 'plist', 'csv']
    content_types = {
        'json': 'application/json',
        'jsonp': 'text/javascript',
        'xml': 'application/xml',
        'yaml': 'text/yaml',
        'html': 'text/html',
        'plist': 'application/x-plist',
        'csv': 'text/csv',
    }

    def to_csv(self, data, options=None):
        options = options or {}
        data = self.to_simple(data, options)
        raw_data = StringIO.StringIO()
        # Untested, so this might not work exactly right.
        for item in data:
            writer = csv.DictWriter(raw_data, item.keys(), extrasaction='ignore')
            writer.write(item)
        return raw_data

    def from_csv(self, content):
        raw_data = StringIO.StringIO(content)
```

```

data = []
# Untested, so this might not work exactly right.
for item in csv.DictReader(raw_data):
    data.append(item)
return data

```

13.3 Serializer Methods

A swappable class for serialization.

This handles most types of data as well as the following output formats:

- * json
- * jsonp
- * xml
- * yaml
- * html
- * plist

It was designed to make changing behavior easy, either by overriding the various format methods (i.e. `to_json`), by changing the `formats/content_types` options or by altering the other hook methods.

13.3.1 `get_mime_for_format`

`Serializer.get_mime_for_format(self, format):`

Given a format, attempts to determine the correct MIME type.

If not available on the current `Serializer`, returns `application/json` by default.

13.3.2 `format_datetime`

`Serializer.format_datetime(data):`

A hook to control how datetimes are formatted.

Can be overridden at the `Serializer` level (`datetime_formatting`) or globally (via `settings.TASTYPIE_DATETIME_FORMATTING`).

Default is `iso-8601`, which looks like “2010-12-16T03:02:14”.

13.3.3 `format_date`

`Serializer.format_date(data):`

A hook to control how dates are formatted.

Can be overridden at the `Serializer` level (`datetime_formatting`) or globally (via `settings.TASTYPIE_DATETIME_FORMATTING`).

Default is `iso-8601`, which looks like “2010-12-16”.

13.3.4 `format_time`

`Serializer.format_time(data)`:

A hook to control how times are formatted.

Can be overridden at the `Serializer` level (`datetime_formatting`) or globally (via `settings.TASTYPIE_DATETIME_FORMATTING`).

Default is `iso-8601`, which looks like “03:02:14”.

13.3.5 `serialize`

`Serializer.serialize(self, bundle, format='application/json', options={})`:

Given some data and a format, calls the correct method to serialize the data and returns the result.

13.3.6 `deserialize`

`Serializer.deserialize(self, content, format='application/json')`:

Given some data and a format, calls the correct method to deserialize the data and returns the result.

13.3.7 `to_simple`

`Serializer.to_simple(self, data, options)`:

For a piece of data, attempts to recognize it and provide a simplified form of something complex.

This brings complex Python data structures down to native types of the serialization format(s).

13.3.8 `to_etree`

`Serializer.to_etree(self, data, options=None, name=None, depth=0)`:

Given some data, converts that data to an `etree.Element` suitable for use in the XML output.

13.3.9 `from_etree`

`Serializer.from_etree(self, data)`:

Not the smartest deserializer on the planet. At the request level, it first tries to output the deserialized subelement called “object” or “objects” and falls back to deserializing based on hinted types in the XML element attribute “type”.

13.3.10 `to_json`

`Serializer.to_json(self, data, options=None)`:

Given some Python data, produces JSON output.

13.3.11 from_json

Serializer.from_json(self, content):

Given some JSON data, returns a Python dictionary of the decoded data.

13.3.12 to_jsonp

Serializer.to_jsonp(self, data, options=None):

Given some Python data, produces JSON output wrapped in the provided callback.

13.3.13 to_xml

Serializer.to_xml(self, data, options=None):

Given some Python data, produces XML output.

13.3.14 from_xml

Serializer.from_xml(self, content):

Given some XML data, returns a Python dictionary of the decoded data.

13.3.15 to_yaml

Serializer.to_yaml(self, data, options=None):

Given some Python data, produces YAML output.

13.3.16 from_yaml

Serializer.from_yaml(self, content):

Given some YAML data, returns a Python dictionary of the decoded data.

13.3.17 to_plist

Serializer.to_plist(self, data, options=None):

Given some Python data, produces binary plist output.

13.3.18 from_plist

Serializer.from_plist(self, content):

Given some binary plist data, returns a Python dictionary of the decoded data.

13.3.19 `to_html`

`Serializer.to_html(self, data, options=None):`

Reserved for future usage.

The desire is to provide HTML output of a resource, making an API available to a browser. This is on the TODO list but not currently implemented.

13.3.20 `from_html`

`Serializer.from_html(self, content):`

Reserved for future usage.

The desire is to handle form-based (maybe Javascript?) input, making an API available to a browser. This is on the TODO list but not currently implemented.

THROTTLING

Sometimes, the client on the other end may request data too frequently or you have a business use case that dictates that the client should be limited to a certain number of requests per hour.

For this, TastyPie includes throttling as a way to limit the number of requests in a timeframe.

14.1 Usage

To specify a throttle, add the `Throttle` class to the `Meta` class on the `Resource`:

```
from django.contrib.auth.models import User
from tastypie.resources import ModelResource
from tastypie.throttle import BaseThrottle

class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()
        resource_name = 'auth/user'
        excludes = ['email', 'password', 'is_superuser']
        # Add it here.
        throttle = BaseThrottle(throttle_at=100)
```

14.2 Throttle Options

Each of the `Throttle` classes accepts the following initialization arguments:

- `throttle_at` - the number of requests at which the user should be throttled. Default is 150 requests.
- `timeframe` - the length of time (in seconds) in which the user make up to the `throttle_at` requests. Default is 3600 seconds (1 hour).
- `expiration` - the length of time to retain the times the user has accessed the api in the cache. Default is 604800 (1 week).

TastyPie ships with the following `Throttle` classes:

14.2.1 BaseThrottle

The no-op throttle option, this does no throttling but implements much of the common logic and serves as an api-compatible plug. Very useful for development.

14.2.2 CacheThrottle

This uses just the cache to manage throttling. Fast but prone to cache misses and/or cache restarts.

14.2.3 CacheDBThrottle

A write-through option that uses the cache first & foremost, but also writes through to the database to persist access times. Useful for logging client accesses & with RAM-only caches.

14.3 Implementing Your Own Throttle

Writing a `Throttle` class is not quite as simple as the other components. There are two important methods, `should_be_throttled` & `accessed`. The `should_be_throttled` method dictates whether or not the client should be throttled. The `accessed` method allows for the recording of the hit to the API.

An example of a subclass might be:

```
import random
from tastypie.throttle import BaseThrottle

class RandomThrottle(BaseThrottle):
    def should_be_throttled(self, identifier, **kwargs):
        if random.randint(0, 10) % 2 == 0:
            return True

        return False

    def accessed(self, identifier, **kwargs):
        pass
```

This throttle class would pick a random number between 0 & 10. If the number is even, their request is allowed through; otherwise, their request is throttled & rejected.

TASTYPIE COOKBOOK

15.1 Adding Custom Values

You might encounter cases where you wish to include additional data in a response which is not obtained from a field or method on your model. You can easily extend the `dehydrate()` method to provide additional values:

```
class MyModelResource(Resource):
    class Meta:
        qs = MyModel.objects.all()

    def dehydrate(self, bundle):
        bundle.data['custom_field'] = "Whatever you want"
        return bundle
```

15.2 Pre-Request Alterations To The Queryset

A common pattern is needing to limit a queryset by something that changes per-request, for instance the date/time. You can accomplish this by lightly modifying `get_object_list()`:

```
class MyResource(ModelResource):
    class Meta:
        queryset = MyObject.objects.all()

    def get_object_list(self, request):
        return super(MyResource, self).get_object_list(request).filter(start_date__gte=datetime.datetime.now())
```

15.3 Using Your Resource In Regular Views

In addition to using your resource classes to power the API, you can also use them to write other parts of your application, such as your views. For instance, if you wanted to encode user information in the page for some Javascript's use, you could do the following:

```
# views.py
from django.shortcuts import render_to_response
from myapp.api.resources import UserResource

def user_detail(request, username):
    ur = UserResource()
```

```
user = ur.obj_get(username=username)

# Other things get prepped to go into the context then...

ur_bundle = ur.build_bundle(obj=user, request=request)
return render_to_response('myapp/user_detail.html', {
    # Other things here.
    "user_json": ur.serialize(None, ur.full_dehydrate(ur_bundle), 'application/json'),
})
```

15.4 Using Non-PK Data For Your URLs

By convention, `ModelResource`'s usually expose the detail endpoints utilizing the primary key of the `Model` they represent. However, this is not a strict requirement. Each URL can take other named URLconf parameters that can be used for the lookup.

For example, if you want to expose `User` resources by username, you can do something like the following:

```
# myapp/api/resources.py
class UserResource(ModelResource):
    class Meta:
        queryset = User.objects.all()

    def override_urls(self):
        return [
            url(r"^(?P<resource_name>%s)/(?P<username>[\w\d_-]+)/$" % self._meta.resource_name, self)
        ]
```

The added URLconf matches before the standard URLconf included by default & matches on the username provided in the URL.

15.5 Nested Resources

You can also do “nested resources” (resources within another related resource) by lightly overriding the `override_urls` method & adding on a new method to handle the children:

```
class ParentResource(ModelResource):
    children = fields.ToManyField(ChildResource, 'children')

    def override_urls(self):
        return [
            url(r"^(?P<resource_name>%s)/(?P<pk>[\w/\-]+)/children%s$" % (self._meta.resource_name,
            ])

    def get_children(self, request, **kwargs):
        try:
            obj = self.cached_obj_get(request=request, **self.remove_api_resource_names(kwargs))
        except ObjectDoesNotExist:
            return HttpGone()
        except MultipleObjectsReturned:
            return HttpMultipleChoices("More than one resource is found at this URI.")

        child_resource = ChildResource()
        return child_resource.get_detail(request, parent_id=obj.pk)
```

Another alternative approach is to override the dispatch method:

```
# myapp/api/resources.py
class EntryResource(ModelResource):
    user = fields.ForeignKey(UserResource, 'user')

    class Meta:
        queryset = Entry.objects.all()
        resource_name = 'entry'

    def dispatch(self, request_type, request, **kwargs):
        username = kwargs.pop('username')
        kwargs['user'] = get_object_or_404(User, username=username)
        return super(EntryResource, self).dispatch(request_type, request, **kwargs)

# urls.py
from django.conf.urls.defaults import *
from myapp.api import EntryResource

entry_resource = EntryResource()

urlpatterns = patterns('',
    # The normal jazz here, then...
    (r'^api/(?P<username>\w+)/', include(entry_resource.urls)),
)
```

15.6 Adding Search Functionality

Another common request is being able to integrate search functionality. This approach uses [Haystack](#), though you could hook it up to any search technology. We leave the CRUD methods of the resource alone, choosing to add a new endpoint at `/api/v1/notes/search/`:

```
from django.conf.urls.defaults import *
from django.core.paginator import Paginator, InvalidPage
from django.http import Http404
from haystack.query import SearchQuerySet
from tastypie.resources import ModelResource
from tastypie.utils import trailing_slash
from notes.models import Note

class NoteResource(ModelResource):
    class Meta:
        queryset = Note.objects.all()
        resource_name = 'notes'

    def override_urls(self):
        return [
            url(r"^(?P<resource_name>%s)/search%s$" % (self._meta.resource_name, trailing_slash()),
            ]

    def get_search(self, request, **kwargs):
        self.method_check(request, allowed=['get'])
        self.is_authenticated(request)
        self.throttle_check(request)
```

```
# Do the query.
sqs = SearchQuerySet().models(Note).load_all().auto_query(request.GET.get('q', ''))
paginator = Paginator(sqs, 20)

try:
    page = paginator.page(int(request.GET.get('page', 1)))
except InvalidPage:
    raise Http404("Sorry, no results on that page.")

objects = []

for result in page.object_list:
    bundle = self.build_bundle(obj=result.object, request=request)
    bundle = self.full_dehydrate(bundle)
    objects.append(bundle)

object_list = {
    'objects': objects,
}

self.log_throttled_access(request)
return self.create_response(request, object_list)
```

15.7 Creating per-user resources

One might want to create an API which will require every user to authenticate and every user will be working only with objects associated with him. Let's see how to implement it for two basic operations: listing and creation of an object.

For listing we want to list only objects for which 'user' field matches 'request.user'. This could be done by applying filter in `apply_authorization_limits` method of your resource.

For creating we'd have to wrap `obj_create` method of `ModelResource`. Then the resulting code will look something like:

```
# myapp/api/resources.py
class EnvironmentResource(ModelResource):
    class Meta:
        queryset = Environment.objects.all()
        resource_name = 'environment'
        list_allowed_methods = ['get', 'post']
        authentication = ApiKeyAuthentication()
        authorization = Authorization()

    def obj_create(self, bundle, request=None, **kwargs):
        return super(EnvironmentResource, self).obj_create(bundle, request, user=request.user)

    def apply_authorization_limits(self, request, object_list):
        return object_list.filter(user=request.user)
```

15.8 camelCase JSON Serialization

The convention in the world of Javascript has standardized on camelCase, where Tastypie uses underscore syntax, which can lead to "ugly" looking code in Javascript. You can create a custom serializer that emits values in camelCase

instead:

```

from tastypie.serializers import Serializer

class CamelCaseJSONSerializer(Serializer):
    formats = ['json']
    content_types = {
        'json': 'application/json',
    }

    def to_json(self, data, options=None):
        # Changes underscore_separated names to camelCase names to go from python convention to java
        data = self.to_simple(data, options)

    def underscoreToCamel(match):
        return match.group()[0] + match.group()[2].upper()

    def camelize(data):
        if isinstance(data, dict):
            new_dict = {}
            for key, value in data.items():
                new_key = re.sub(r"[a-z]_[a-z]", underscoreToCamel, key)
                new_dict[new_key] = camelize(value)
            return new_dict
        if isinstance(data, (list, tuple)):
            for i in range(len(data)):
                data[i] = camelize(data[i])
            return data
        return data

    camelized_data = camelize(data)

    return simplejson.dumps(camelized_data, sort_keys=True)

    def from_json(self, content):
        # Changes camelCase names to underscore_separated names to go from javascript convention to p
        data = simplejson.loads(content)

    def camelToUnderscore(match):
        return match.group()[0] + "_" + match.group()[1].lower()

    def underscoreize(data):
        if isinstance(data, dict):
            new_dict = {}
            for key, value in data.items():
                new_key = re.sub(r"[a-z][A-Z]", camelToUnderscore, key)
                new_dict[new_key] = underscoreize(value)
            return new_dict
        if isinstance(data, (list, tuple)):
            for i in range(len(data)):
                data[i] = underscoreize(data[i])
            return data
        return data

    underscored_data = underscoreize(data)

    return underscored_data

```


DEBUGGING TASTYPIE

There are some common problems people run into when using Tastypie for the first time. Some of the common problems and things to try appear below.

16.1 “I’m getting XML output in my browser but I want JSON output!”

This is actually not a bug and JSON support is present in your `Resource`. This issue is that Tastypie respects the `Accept` header your browser sends. Most browsers send something like:

```
Accept: application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,*/*;q=0.5
```

Note that `application/xml` comes first, which is a format that Tastypie handles by default, hence why you receive XML.

If you use `curl` from the command line, you should receive JSON by default:

```
curl http://localhost:8000/api/v1/
```

If you want JSON in the browser, simply append `?format=json` to your URL. Tastypie always respects this override first, before it falls back to the `Accept` header.

16.2 “What’s the format for a POST or PUT?”

You can view full schema for your resource through *Inspecting The Resource’s Schema*.

In general, Tastypie will accept resources in the same format as it gives you. This means that you can see what any POST or PUT should look like by performing a GET of that resource.

Creating a duplicate of an entry, using Python and `Requests`:

```
import requests
import json

response = requests.get('http://localhost:8000/api/v1/entry/1/')
event = json.loads(response.content)

del event['id'] # We want the server to assign a new id

response = requests.post('http://localhost:8000/api/v1/entry/',
                        data=json.dumps(event),
                        headers={'content-type': 'application/json'})
```

The new event should be almost identical, with the exception of readonly fields. This method may fail if your model has a unique constraint, or otherwise fails validation.

This is less likely to happen on PUT, except for application logic changes (e.g. a *last_update* field). The following two `curl` commands replace and entry with an copy:

```
curl -H 'Accept: application/json' 'http://localhost:8000/api/v1/entry/1/' | \  
curl -H 'Content-Type: application/json' -X PUT --data @- "http://localhost:8000/api/v1/entry/1/"
```

You can do this over an entire collection as well:

```
curl -H 'Accept: application/json' 'http://localhost:8000/api/v1/entry/?limit=0' | \  
curl -H 'Content-Type: application/json' -X PUT --data @- "http://localhost:8000/api/v1/entry/"
```

SITES USING TASTYPIE

The following sites are a partial list of people using TastyPie. I'm always interested in adding more sites, so please find me ([daniellindsley](#)) via IRC or start a mailing list thread.

17.1 LJWorld Marketplace

- <http://www2.ljworld.com/marketplace/api/v1/?format=json>

17.2 Forkinit

Read-only API access to recipes.

- <http://forkinit.com/>
- <http://forkinit.com/api/v1/?format=json>

17.3 Read The Docs

A hosted documentation site, primarily for Python docs. General purpose read-write access.

- <http://readthedocs.org/>
- <http://readthedocs.org/api/v1/?format=json>

17.4 Luzme

An e-book search site that lets you fetch pricing information.

- <http://luzme.com/>
- <http://luzme.readthedocs.org/en/latest/>

17.5 Politifact

To power their mobile (iPhone/Android/Playbook) applications.

- <http://www.politifact.com/mobile/>

GETTING HELP

There are two primary ways of getting help. We have a [mailing list](http://groups.google.com/group/django-tastypie/) hosted at Google (<http://groups.google.com/group/django-tastypie/>) and an IRC channel (#tastypie on irc.freenode.net) to get help, want to bounce idea or generally shoot the breeze.

QUICK START

1. Add `tastypie` to `INSTALLED_APPS`.
2. Create an `api` directory in your app with a bare `__init__.py`.
3. Create an `<my_app>/api/resources.py` file and place the following in it:

```
from tastypie.resources import ModelResource
from my_app.models import MyModel

class MyModelResource(ModelResource):
    class Meta:
        queryset = MyModel.objects.all()
        allowed_methods = ['get']
```

4. In your root `URLconf`, add the following code (around where the admin code might be):

```
from tastypie.api import Api
from my_app.api.resources import MyModelResource

v1_api = Api(api_name='v1')
v1_api.register(MyModelResource())

urlpatterns = patterns('',
    # ...more URLconf bits here...
    # Then add:
    (r'^api/', include(v1_api.urls)),
)
```

5. Hit `http://localhost:8000/api/v1/?format=json` in your browser!

REQUIREMENTS

Tastypie requires the following modules. If you use [Pip](#), you can install the necessary bits via the included `requirements.txt`:

20.1 Required

- Python 2.5+
- Django 1.2+
- `mimemagic` 0.1.3+ (<http://code.google.com/p/mimemagic/>)
 - Older versions will work, but their behavior on JSON/JSONP is a touch wonky.
- `dateutil` (<http://labix.org/python-dateutil>) ≥ 1.5 , < 2.0

20.2 Optional

- `python_digest` (<https://bitbucket.org/akoha/python-digest/>)
- `lxml` (<http://codespeak.net/lxml/>) if using the XML serializer
- `pyyaml` (<http://pyyaml.org/>) if using the YAML serializer
- `biplist` (<http://explorapp.com/biplist/>) if using the binary plist serializer

PYTHON MODULE INDEX

t

`tastypie.fields, ??`