
Python Serverless Microframework for AWS

Release 1.15.1

Jun 18, 2020

1	Quickstart and Tutorial	3
1.1	Credentials	3
1.2	Creating Your Project	4
1.3	Deploying	4
1.4	Next Steps	5
2	Tutorial: URL Parameters	7
3	Tutorial: Error Messages	9
4	Tutorial: Additional Routing	11
5	Tutorial: Request Metadata	13
6	Tutorial: Request Content Types	17
7	Tutorial: Customizing the HTTP Response	19
8	Tutorial: GZIP compression for json	21
9	Tutorial: CORS Support	23
10	Tutorial: Policy Generation	25
10.1	Manually Providing Policies	26
10.2	Experimental Status	27
11	Tutorial: Using Custom Authentication	29
11.1	API Key	29
11.2	Using AWS IAM	29
11.3	Using Amazon Cognito User Pools	29
11.4	Using Custom Authorizers	30
12	Tutorial: Local Mode	31
13	Deleting Your App	33
14	Topics	35
14.1	Routing	35
14.2	Views	37
14.3	Configuration File	41

14.4	Multifile Support	48
14.5	Logging	49
14.6	SDK Generation	51
14.7	Chalice Stages	53
14.8	App Packaging	55
14.9	Python Version Support	58
14.10	Changing Python Runtime Versions	59
14.11	AWS CloudFormation Support	59
14.12	Terraform Support	63
14.13	Authorization	65
14.14	Lambda Event Sources	68
14.15	Pure Lambda Functions	72
14.16	Blueprints	73
14.17	Websockets	76
14.18	Continuous Deployment (CD)	78
14.19	Experimental APIs	80
15	API Reference	83
15.1	Chalice	83
15.2	Request	88
15.3	Response	89
15.4	Authorization	89
15.5	APIGateway	91
15.6	WebsocketAPI	92
15.7	CORS	93
15.8	Event Sources	94
15.9	Blueprints	98
15.10	Websockets	98
16	Tutorials	101
16.1	Websocket Tutorials	101
17	Upgrade Notes	121
17.1	Upgrade Notes	121
17.2	Indices and tables	127



AWS Chalice allows you to quickly create and deploy applications that use Amazon API Gateway and AWS Lambda. It provides:

- A command line tool for creating, deploying, and managing your app
- A familiar and easy to use API for declaring views in python code
- Automatic IAM policy generation

```
$ pip install chalice
$ chalice new-project helloworld && cd helloworld
$ cat app.py

from chalice import Chalice

app = Chalice(app_name="helloworld")

@app.route("/")
def index():
    return {"hello": "world"}

$ chalice deploy
...
https://endpoint/dev

$ curl https://endpoint/api
{"hello": "world"}
```

Up and running in less than 30 seconds.

Quickstart and Tutorial

In this tutorial, you'll use the `chalice` command line utility to create and deploy a basic REST API. This quickstart uses Python 3.7, but AWS Chalice supports all versions of python supported by AWS Lambda, which includes python2.7, python3.6, python3.7, python3.8. We recommend you use a version of Python 3. You can find the latest versions of python on the [Python download page](#).

To install Chalice, we'll first create and activate a virtual environment in python3.7:

```
$ python3 --version
Python 3.7.3
$ python3 -m venv venv37
$ . venv37/bin/activate
```

Next we'll install Chalice using `pip`:

```
$ python3 -m pip install chalice
```

You can verify you have chalice installed by running:

```
$ chalice --help
Usage: chalice [OPTIONS] COMMAND [ARGS]...
...
```

1.1 Credentials

Before you can deploy an application, be sure you have credentials configured. If you have previously configured your machine to run `boto3` (the AWS SDK for Python) or the AWS CLI then you can skip this section.

If this is your first time configuring credentials for AWS you can follow these steps to quickly get started:

```
$ mkdir ~/.aws
$ cat >> ~/.aws/config
[default]
aws_access_key_id=YOUR_ACCESS_KEY_HERE
aws_secret_access_key=YOUR_SECRET_ACCESS_KEY
region=YOUR_REGION (such as us-west-2, us-west-1, etc)
```

If you want more information on all the supported methods for configuring credentials, see the [boto3 docs](#).

1.2 Creating Your Project

The next thing we'll do is use the `chalice` command to create a new project:

```
$ chalice new-project helloworld
```

This will create a `helloworld` directory. Cd into this directory. You'll see several files have been created for you:

```
$ cd helloworld
$ ls -la
drwxr-xr-x  .chalice
-rw-r--r--  app.py
-rw-r--r--  requirements.txt
```

You can ignore the `.chalice` directory for now, the two main files we'll focus on is `app.py` and `requirements.txt`.

Let's take a look at the `app.py` file:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/')
def index():
    return {'hello': 'world'}
```

The `new-project` command created a sample app that defines a single view, `/`, that when called will return the JSON body `{"hello": "world"}`.

1.3 Deploying

Let's deploy this app. Make sure you're in the `helloworld` directory and run `chalice deploy`:

```
$ chalice deploy
...
Initiating first time deployment...
https://qxea58oupc.execute-api.us-west-2.amazonaws.com/api/
```

You now have an API up and running using API Gateway and Lambda:

```
$ curl https://qxea58oupc.execute-api.us-west-2.amazonaws.com/api/
{"hello": "world"}
```

Try making a change to the returned dictionary from the `index()` function. You can then redeploy your changes by running `chalice deploy`.

For the rest of these tutorials, we'll be using `httpie` instead of `curl` (<https://github.com/jakubroztocil/httpie>) to test our API. You can install `httpie` using `pip install httpie`, or if you're on Mac, you can run `brew install httpie`. The Github link has more information on installation instructions. Here's an example of using `httpie` to request the root resource of the API we just created. Note that the command name is `http`:

```
$ http https://qxea58oupc.execute-api.us-west-2.amazonaws.com/api/
HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 18
```



```
Content-Type: application/json
Date: Mon, 30 May 2016 17:55:50 GMT
X-Cache: Miss from cloudfront

{
  "hello": "world"
}
```

Additionally, the API Gateway endpoints will be shortened to `https://endpoint/api/` for brevity. Be sure to substitute `https://endpoint/api/` for the actual endpoint that the chalice CLI displays when you deploy your API (it will look something like `https://abcdefg.execute-api.us-west-2.amazonaws.com/api/`).

1.4 Next Steps

You've now created your first app using chalice.

The next few sections will build on this quickstart section and introduce you to additional features including: URL parameter capturing, error handling, advanced routing, current request metadata, and automatic policy generation.

Tutorial: URL Parameters

Now we're going to make a few changes to our `app.py` file that demonstrate additional capabilities provided by the python serverless microframework for AWS.

Our application so far has a single view that allows you to make an HTTP GET request to `/`. Now let's suppose we want to capture parts of the URI:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

CITIES_TO_STATE = {
    'seattle': 'WA',
    'portland': 'OR',
}

@app.route('/')
def index():
    return {'hello': 'world'}

@app.route('/cities/{city}')
def state_of_city(city):
    return {'state': CITIES_TO_STATE[city]}
```

In the example above, we've now added a `state_of_city` view that allows a user to specify a city name. The view function takes the city name and returns name of the state the city is in. Notice that the `@app.route` decorator has a URL pattern of `/cities/{city}`. This means that the value of `{city}` is captured and passed to the view function. You can also see that the `state_of_city` takes a single argument. This argument is the name of the city provided by the user. For example:

```
GET /cities/seattle --> state_of_city('seattle')
GET /cities/portland --> state_of_city('portland')
```

Now that we've updated our `app.py` file with this new view function, let's redeploy our application. You can run `chalice deploy` from the `helloworld` directory and it will deploy your application:

```
$ chalice deploy
```

Let's try it out. Note the examples below use the `http` command from the `httpie` package. You can install this using `pip install httpie`:

```
$ http https://endpoint/api/cities/seattle
HTTP/1.1 200 OK
```

```
{
  "state": "WA"
}

$ http https://endpoint/api/cities/portland
HTTP/1.1 200 OK

{
  "state": "OR"
}
```

Notice what happens if we try to request a city that's not in our `CITIES_TO_STATE` map:

```
$ http https://endpoint/api/cities/vancouver
HTTP/1.1 500 Internal Server Error
Content-Type: application/json
X-Cache: Error from cloudfront

{
  "Code": "ChaliceViewError",
  "Message": "ChaliceViewError: An internal server error occurred."
}
```

In the next section, we'll see how to fix this and provide better error messages.

Tutorial: Error Messages

In the example above, you'll notice that when our app raised an uncaught exception, a 500 internal server error was returned.

In this section, we're going to show how you can debug and improve these error messages.

The first thing we're going to look at is how we can debug this issue. By default, debugging is turned off, but you can enable debugging to get more information:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')
app.debug = True
```

The `app.debug = True` enables debugging for your app. Save this file and redeploy your changes:

```
$ chalice deploy
...
https://endpoint/api/
```

Now, when you request the same URL that returned an internal server error, you'll get back the original stack trace:

```
$ http https://endpoint/api/cities/vancouver
Traceback (most recent call last):
  File "/var/task/chalice/app.py", line 304, in _get_view_function_response
    response = view_function(*function_args)
  File "/var/task/app.py", line 18, in state_of_city
    return {'state': CITIES_TO_STATE[city]}
KeyError: u'vancouver'
```

We can see that the error is caused from an uncaught `KeyError` resulting from trying to access the `vancouver` key.

Now that we know the error, we can fix our code. What we'd like to do is catch this exception and instead return a more helpful error message to the user. Here's the updated code:

```
from chalice import BadRequestError

@app.route('/cities/{city}')
def state_of_city(city):
    try:
        return {'state': CITIES_TO_STATE[city]}
    except KeyError:
        raise BadRequestError("Unknown city '%s', valid choices are: %s" % (
            city, ', '.join(CITIES_TO_STATE.keys())))
```

Save and deploy these changes:

```
$ chalice deploy
$ http https://endpoint/api/cities/vancouver
HTTP/1.1 400 Bad Request

{
  "Code": "BadRequestError",
  "Message": "Unknown city 'vancouver', valid choices are: portland, seattle"
}
```

We can see now that we have received a `Code` and `Message` key, with the message being the value we passed to `BadRequestError`. Whenever you raise a `BadRequestError` from your view function, the framework will return an HTTP status code of 400 along with a JSON body with a `Code` and `Message`. There are a few additional exceptions you can raise from your python code:

```
* BadRequestError - return a status code of 400
* UnauthorizedError - return a status code of 401
* ForbiddenError - return a status code of 403
* NotFoundError - return a status code of 404
* ConflictError - return a status code of 409
* UnprocessableEntityError - return a status code of 422
* TooManyRequestsError - return a status code of 429
* ChaliceViewError - return a status code of 500
```

You can import these directly from the `chalice` package:

```
from chalice import UnauthorizedError
```

Tutorial: Additional Routing

So far, our examples have only allowed GET requests. It's actually possible to support additional HTTP methods. Here's an example of a view function that supports PUT:

```
@app.route('/resource/{value}', methods=['PUT'])
def put_test(value):
    return {"value": value}
```

We can test this method using the `http` command:

```
$ http PUT https://endpoint/api/resource/foo
HTTP/1.1 200 OK

{
    "value": "foo"
}
```

Note that the `methods` kwarg accepts a list of methods. Your view function will be called when any of the HTTP methods you specify are used for the specified resource. For example:

```
@app.route('/myview', methods=['POST', 'PUT'])
def myview():
    pass
```

The above view function will be called when either an HTTP POST or PUT is sent to `/myview`.

Alternatively if you do not want to share the same view function across multiple HTTP methods for the same route url, you may define separate view functions to the same route url but have the view functions differ by HTTP method. For example:

```
@app.route('/myview', methods=['POST'])
def myview_post():
    pass

@app.route('/myview', methods=['PUT'])
def myview_put():
    pass
```

This setup will route all HTTP POST's to `/myview` to the `myview_post()` view function and route all HTTP PUT's to `/myview` to the `myview_put()` view function. It is also important to note that the view functions **must** have unique names. For example, both view functions cannot be named `myview()`.

In the next section we'll go over how you can introspect the given request in order to differentiate between various HTTP methods.

Tutorial: Request Metadata

In the examples above, you saw how to create a view function that supports an HTTP PUT request as well as a view function that supports both POST and PUT via the same view function. However, there's more information we might need about a given request:

- In a PUT/POST, you frequently send a request body. We need some way of accessing the contents of the request body.
- For view functions that support multiple HTTP methods, we'd like to detect which HTTP method was used so we can have different code paths for PUTs vs. POSTs.

All of this and more is handled by the current request object that the `chalice` library makes available to each view function when it's called.

Let's see an example of this. Suppose we want to create a view function that allowed you to PUT data to an object and retrieve that data via a corresponding GET. We could accomplish that with the following view function:

```
from chalice import NotFoundError

OBJECTS = {}

@app.route('/objects/{key}', methods=['GET', 'PUT'])
def myobject(key):
    request = app.current_request
    if request.method == 'PUT':
        OBJECTS[key] = request.json_body
    elif request.method == 'GET':
        try:
            return {key: OBJECTS[key]}
        except KeyError:
            raise NotFoundError(key)
```

Save this in your `app.py` file and rerun `chalice deploy`. Now, you can make a PUT request to `/objects/your-key` with a request body, and retrieve the value of that body by making a subsequent GET request to the same resource. Here's an example of its usage:

```
# First, trying to retrieve the key will return a 404.
$ http GET https://endpoint/api/objects/mykey
HTTP/1.1 404 Not Found

{
  "Code": "NotFoundError",
  "Message": "mykey"
}
```

```
# Next, we'll create that key by sending a PUT request.
$ echo '{"foo": "bar"}' | http PUT https://endpoint/api/objects/mykey
HTTP/1.1 200 OK

null

# And now we no longer get a 404, we instead get the value we previously
# put.
$ http GET https://endpoint/api/objects/mykey
HTTP/1.1 200 OK

{
  "mykey": {
    "foo": "bar"
  }
}
```

You might see a problem with storing the objects in a module level `OBJECTS` variable. We address this in the next section.

The `app.current_request` object also has the following properties.

- `current_request.query_params` - A dict of the query params for the request.
- `current_request.headers` - A dict of the request headers.
- `current_request.uri_params` - A dict of the captured URI params.
- `current_request.method` - The HTTP method (as a string).
- `current_request.json_body` - The parsed JSON body (`json.loads(raw_body)`)
- `current_request.raw_body` - The raw HTTP body as bytes.
- `current_request.context` - A dict of additional context information
- `current_request.stage_vars` - Configuration for the API Gateway stage

Don't worry about the context and stage_vars for now. We haven't discussed those concepts yet. The `current_request` object also has a `to_dict` method, which returns all the information about the current request as a dictionary. Let's use this method to write a view function that returns everything it knows about the request:

```
@app.route('/introspect')
def introspect():
    return app.current_request.to_dict()
```

Save this to your `app.py` file and redeploy with `chalice deploy`. Here's an example of hitting the `/introspect` URL. Note how we're sending a query string as well as a custom `X-TestHeader` header:

```
$ http 'https://endpoint/api/introspect?query1=value1&query2=value2' 'X-TestHeader: Foo'
HTTP/1.1 200 OK

{
  "context": {
    "apiId": "apiId",
    "httpMethod": "GET",
    "identity": {
      "accessKey": null,
      "accountId": null,
      "apiKey": null,
      "caller": null,
      "cognitoAuthenticationProvider": null,
```

```
        "cognitoAuthenticationType": null,
        "cognitoIdentityId": null,
        "cognitoIdentityPoolId": null,
        "sourceIp": "1.1.1.1",
        "userAgent": "HTTPie/0.9.3",
        "userArn": null
    },
    "requestId": "request-id",
    "resourceId": "resourceId",
    "resourcePath": "/introspect",
    "stage": "dev"
},
"headers": {
    "accept": "*/*",
    ...
    "x-testheader": "Foo"
},
"method": "GET",
"query_params": {
    "query1": "value1",
    "query2": "value2"
},
"raw_body": null,
"stage_vars": null,
"uri_params": null
}
```

Tutorial: Request Content Types

The default behavior of a view function supports a request body of `application/json`. When a request is made with a `Content-Type` of `application/json`, the `app.current_request.json_body` attribute is automatically set for you. This value is the parsed JSON body.

You can also configure a view function to support other content types. You can do this by specifying the `content_types` parameter value to your `app.route` function. This parameter is a list of acceptable content types. Here's an example of this feature:

```
import sys

from chalice import Chalice
from urllib.parse import urlparse, parse_qs

app = Chalice(app_name='helloworld')

@app.route('/', methods=['POST'],
            content_types=['application/x-www-form-urlencoded'])
def index():
    parsed = parse_qs(app.current_request.raw_body.decode())
    return {
        'states': parsed.get('states', [])
    }
```

There's a few things worth noting in this view function. First, we've specified that we only accept the `application/x-www-form-urlencoded` content type. If we try to send a request with `application/json`, we'll now get a 415 Unsupported Media Type response:

```
$ http POST https://endpoint/api/ states=WA states=CA --debug
...
>>> requests.request(**{'allow_redirects': False,
    'headers': {'Accept': 'application/json',
                'Content-Type': 'application/json'},
    ...

HTTP/1.1 415 Unsupported Media Type

{
    "message": "Unsupported Media Type"
}
```

If we use the `--form` argument, we can see the expected behavior of this view function because `httpie` sets the

Content-Type header to application/x-www-form-urlencoded:

```
$ http --form POST https://endpoint/api/formtest states=WA states=CA --debug
...
>>> requests.request(**{'allow_redirects': False,
    'headers': {'Content-Type': 'application/x-www-form-urlencoded; charset=utf-8'},
    ...
HTTP/1.1 200 OK
{
    "states": [
        "WA",
        "CA"
    ]
}
```

The second thing worth noting is that `app.current_request.json_body` **is only available for the application/json content type**. In our example above, we used `app.current_request.raw_body` to access the raw body bytes:

```
parsed = parse_qs(app.current_request.raw_body)
```

`app.current_request.json_body` is set to `None` whenever the Content-Type is not `application/json`. This means that you will need to use `app.current_request.raw_body` and parse the request body as needed.

Tutorial: Customizing the HTTP Response

The return value from a chalice view function is serialized as JSON as the response body returned back to the caller. This makes it easy to create rest APIs that return JSON response bodies.

Chalice allows you to control this behavior by returning an instance of a chalice specific `Response` class. This behavior allows you to:

- Specify the status code to return
- Specify custom headers to add to the response
- Specify response bodies that are not `application/json`

Here's an example of this:

```
from chalice import Chalice, Response

app = Chalice(app_name='custom-response')

@app.route('/')
def index():
    return Response(body='hello world!',
                    status_code=200,
                    headers={'Content-Type': 'text/plain'})
```

This will result in a plain text response body:

```
$ http https://endpoint/api/
HTTP/1.1 200 OK
Content-Length: 12
Content-Type: text/plain

hello world!
```

Tutorial: GZIP compression for json

The return value from a chalice view function is serialized as JSON as the response body returned back to the caller. This makes it easy to create rest APIs that return JSON response bodies.

Chalice allows you to control this behavior by returning an instance of a chalice specific `Response` class. This behavior allows you to:

- Add `application/json` to `binary_types`
- Specify the status code to return
- Specify custom header `Content-Type: application/json`
- Specify custom header `Content-Encoding: gzip`

Here's an example of this:

```
import json
import gzip
from chalice import Chalice, Response

app = Chalice(app_name='compress-response')
app.api.binary_types.append('application/json')

@app.route('/')
def index():
    blob = json.dumps({'hello': 'world'}).encode('utf-8')
    payload = gzip.compress(blob)
    custom_headers = {
        'Content-Type': 'application/json',
        'Content-Encoding': 'gzip'
    }
    return Response(body=payload,
                    status_code=200,
                    headers=custom_headers)
```


Tutorial: CORS Support

You can specify whether a view supports CORS by adding the `cors=True` parameter to your `@app.route()` call. By default this value is `False`. Global CORS can be set by setting `app.api.cors = True`.

```
@app.route('/supports-cors', methods=['PUT'], cors=True)
def supports_cors():
    return {}
```

Setting `cors=True` has similar behavior to enabling CORS using the AWS Console. This includes:

- Injecting the `Access-Control-Allow-Origin: *` header to your responses, including all error responses you can return.
- Automatically adding an `OPTIONS` method to support preflighting requests.

The preflight request will return a response that includes:

- `Access-Control-Allow-Origin: *`
- The `Access-Control-Allow-Methods` header will return a list of all HTTP methods you've called out in your view function. In the example above, this will be `PUT, OPTIONS`.
- `Access-Control-Allow-Headers: Content-Type, X-Amz-Date, Authorization, X-API-Key, X-Amz-Security-Token`.

If more fine grained control of the CORS headers is desired, set the `cors` parameter to an instance of `CORSConfig` instead of `True`. The `CORSConfig` object can be imported from the `chalice` package it's constructor takes the following keyword arguments that map to CORS headers:

Argument	Type	Header
<code>allow_origin</code>	<code>str</code>	<code>Access-Control-Allow-Origin</code>
<code>allow_headers</code>	<code>list</code>	<code>Access-Control-Allow-Headers</code>
<code>expose_headers</code>	<code>list</code>	<code>Access-Control-Expose-Headers</code>
<code>max_age</code>	<code>int</code>	<code>Access-Control-Max-Age</code>
<code>allow_credentials</code>	<code>bool</code>	<code>Access-Control-Allow-Credentials</code>

Code sample defining more CORS headers:

```
from chalice import CORSConfig
cors_config = CORSConfig(
    allow_origin='https://foo.example.com',
    allow_headers=['X-Special-Header'],
    max_age=600,
    expose_headers=['X-Special-Header'],
    allow_credentials=True
)
```

```
@app.route('/custom-cors', methods=['GET'], cors=cors_config)
def supports_custom_cors():
    return {'cors': True}
```

There's a couple of things to keep in mind when enabling cors for a view:

- An OPTIONS method for preflighting is always injected. Ensure that you don't have OPTIONS in the methods=[...] list of your view function.
- Even though the Access-Control-Allow-Origin header can be set to a string that is a space separated list of origins, this behavior does not work on all clients that implement CORS. You should only supply a single origin to the CORSConfig object. If you need to supply multiple origins you will need to define a custom handler for it that accepts OPTIONS requests and matches the Origin header against a whitelist of origins. If the match is successful then return just their Origin back to them in the Access-Control-Allow-Origin header.

Example:

```
from chalice import Chalice, Response

app = Chalice(app_name='multipleorigincors')

_ALLOWED_ORIGINS = set([
    'http://allowed1.example.com',
    'http://allowed2.example.com',
])

@app.route('/cors_multiple_origins', methods=['GET', 'OPTIONS'])
def supports_cors_multiple_origins():
    method = app.current_request.method
    if method == 'OPTIONS':
        headers = {
            'Access-Control-Allow-Method': 'GET,OPTIONS',
            'Access-Control-Allow-Origin': ','.join(_ALLOWED_ORIGINS),
            'Access-Control-Allow-Headers': 'X-Some-Header',
        }
        origin = app.current_request.headers.get('origin', '')
        if origin in _ALLOWED_ORIGINS:
            headers.update({'Access-Control-Allow-Origin': origin})
        return Response(
            body=None,
            headers=headers,
        )
    elif method == 'GET':
        return 'Foo'
```

- Every view function must explicitly enable CORS support.

The last point will change in the future. See [this issue](#) for more information.

Tutorial: Policy Generation

In the previous section we created a basic rest API that allowed you to store JSON objects by sending the JSON in the body of an HTTP PUT request to `/objects/{name}`. You could then retrieve objects by sending a GET request to `/objects/{name}`.

However, there's a problem with the code we wrote:

```
OBJECTS = {  
}  
  
@app.route('/objects/{key}', methods=['GET', 'PUT'])  
def myobject(key):  
    request = app.current_request  
    if request.method == 'PUT':  
        OBJECTS[key] = request.json_body  
    elif request.method == 'GET':  
        try:  
            return {key: OBJECTS[key]}  
        except KeyError:  
            raise NotFoundError(key)
```

We're storing the key value pairs in a module level `OBJECTS` variable. We can't rely on local storage like this persisting across requests.

A better solution would be to store this information in Amazon S3. To do this, we're going to use `boto3`, the AWS SDK for Python. First, install `boto3`:

```
$ pip install boto3
```

Next, add `boto3` to your `requirements.txt` file:

```
$ pip freeze | grep boto3 >> requirements.txt
```

The `requirements.txt` file should be in the same directory that contains your `app.py` file. Next, let's update our view code to use `boto3`:

```
import json  
import boto3  
from botocore.exceptions import ClientError  
  
from chalice import NotFoundError  
  
S3 = boto3.client('s3', region_name='us-west-2')  
BUCKET = 'your-bucket-name'
```

```
@app.route('/objects/{key}', methods=['GET', 'PUT'])
def s3objects(key):
    request = app.current_request
    if request.method == 'PUT':
        S3.put_object(Bucket=BUCKET, Key=key,
                      Body=json.dumps(request.json_body))
    elif request.method == 'GET':
        try:
            response = S3.get_object(Bucket=BUCKET, Key=key)
            return json.loads(response['Body'].read())
        except ClientError as e:
            raise NotFoundError(key)
```

Make sure to change BUCKET with the name of an S3 bucket you own. Redeploy your changes with `chalice deploy`. Now, whenever we make a PUT request to `/objects/keyname`, the data send will be stored in S3. Any subsequent GET requests will retrieve this data from S3.

10.1 Manually Providing Policies

IAM permissions can be auto generated, provided manually or can be pre-created and explicitly configured. To use a pre-configured IAM role ARN for chalice, add these two keys to your chalice configuration. Setting `manage_iam_role` to false tells Chalice to not attempt to generate policies and create IAM role.

```
"manage_iam_role":false
"iam_role_arn":"arn:aws:iam::<account-id>:role/<role-name>"
```

Whenever your application is deployed using chalice, the auto generated policy is written to disk at `<projectdir>/.chalice/policy.json`. When you run the `chalice deploy` command, you can also specify the `--no-autogen-policy` option. Doing so will result in the chalice CLI loading the `<projectdir>/.chalice/policy.json` file and using that file as the policy for the IAM role. You can manually edit this file and specify `--no-autogen-policy` if you'd like to have full control over what IAM policy to associate with the IAM role.

You can also run the `chalice gen-policy` command from your project directory to print the auto generated policy to stdout. You can then use this as a starting point for your policy.

```
$ chalice gen-policy
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListAllMyBuckets"
      ],
      "Resource": [
        "*"
      ],
      "Effect": "Allow",
      "Sid": "9155de6ad1d74e4c8b1448255770e60c"
    }
  ]
}
```

10.2 Experimental Status

The automatic policy generation is still in the early stages, it should be considered experimental. You can always disable policy generation with `--no-autogen-policy` for complete control.

Additionally, you will be prompted for confirmation whenever the auto policy generator detects actions that it would like to add or remove:

```
$ chalice deploy
Updating IAM policy.

The following action will be added to the execution policy:

s3:ListBucket

Would you like to continue? [Y/n]:
```

Tutorial: Using Custom Authentication

AWS API Gateway routes can be authenticated in multiple ways:

- API Key
- AWS IAM
- Cognito User Pools
- Custom Auth Handler

11.1 API Key

```
@app.route('/authenticated', methods=['GET'], api_key_required=True)
def authenticated():
    return {"secure": True}
```

Only requests sent with a valid *X-API-Key* header will be accepted.

11.2 Using AWS IAM

```
authorizer = IAMAuthorizer()

@app.route('/iam-role', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

11.3 Using Amazon Cognito User Pools

To integrate with cognito user pools, you can use the `CognitoUserPoolAuthorizer` object:

```
authorizer = CognitoUserPoolAuthorizer(
    'MyPool', header='Authorization',
    provider_arns=['arn:aws:cognito:...:userpool/name'])

@app.route('/user-pools', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

Note, earlier versions of chalice also have an `app.define_authorizer` method as well as an `authorizer_name` argument on the `@app.route(...)` method. This approach is deprecated in favor of `CognitoUserPoolAuthorizer` and the `authorizer` argument in the `@app.route(...)` method. `app.define_authorizer` will be removed in future versions of chalice.

11.4 Using Custom Authorizers

To integrate with custom authorizers, you can use the `CustomAuthorizer` method on the `app` object. You'll need to set the `authorizer_uri` to the URI of your lambda function.

```
authorizer = CustomAuthorizer(
    'MyCustomAuth', header='Authorization',
    authorizer_uri=('arn:aws:apigateway:region:lambda:path/2015-03-31'
                   '/functions/arn:aws:lambda:region:account-id:'
                   'function:FunctionName/invocations'))

@app.route('/custom-auth', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

Tutorial: Local Mode

As you develop your application, you may want to experiment locally before deploying your changes. You can use `chalice local` to spin up a local HTTP server you can use for testing.

For example, if we have the following `app.py` file:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/')
def index():
    return {'hello': 'world'}
```

We can run `chalice local` to test this API locally:

```
$ chalice local
Serving on localhost:8000
```

We can override the port using:

```
$ chalice local --port=8080
```

We can now test our API using `localhost:8000`:

```
$ http localhost:8000/
HTTP/1.0 200 OK
Content-Length: 18
Content-Type: application/json
Date: Thu, 27 Oct 2016 20:08:43 GMT
Server: BaseHTTP/0.3 Python/3.7.3

{
  "hello": "world"
}
```

The `chalice local` command *does not* assume the role associated with your lambda function, so you'll need to use an `AWS_PROFILE` that has sufficient permissions to your AWS resources used in your `app.py`.

Deleting Your App

You can use the `chalice delete` command to delete your app. Similar to the `chalice deploy` command, you can specify which chalice stage to delete. By default it will delete the `dev` stage:

```
$ chalice delete --stage dev
Deleting Rest API: duvw4kwyl3
Deleting function aws:arn:lambda:region:123456789:helloworld-dev
Deleting IAM Role helloworld-dev
```


14.1 Routing

The `Chalice.route()` method is used to construct which routes you want to create for your API. The concept is the same mechanism used by `Flask` and `bottle`. You decorate a function with `@app.route(...)`, and whenever a user requests that URL, the function you've decorated is called. For example, suppose you deployed this app:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/')
def index():
    return {'view': 'index'}

@app.route('/a')
def a():
    return {'view': 'a'}

@app.route('/b')
def b():
    return {'view': 'b'}
```

If you go to `https://endpoint/`, the `index()` function would be called. If you went to `https://endpoint/a` and `https://endpoint/b`, then the `a()` and `b()` function would be called, respectively.

Note: Do not end your route paths with a trailing slash. If you do this, the `chalice deploy` command will raise a validation error.

You can also create a route that captures part of the URL. This captured value will then be passed in as arguments to your view function:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/users/{name}')
```

```
def users(name):  
    return {'name': name}
```

If you then go to `https://endpoint/users/james`, then the view function will be called as: `users('james')`. The parameters are passed as keyword parameters based on the name as they appear in the URL. The argument names for the view function must match the name of the captured argument:

```
from chalice import Chalice  
  
app = Chalice(app_name='helloworld')  
  
@app.route('/a/{first}/b/{second}')  
def users(first, second):  
    return {'first': first, 'second': second}
```

14.1.1 Other Request Metadata

The route path can only contain `[a-zA-Z0-9._-]` chars and curly braces for parts of the URL you want to capture. You do not need to model other parts of the request you want to capture, including headers and query strings. Within a view function, you can introspect the current request using the `app.current_request` attribute. This also means you cannot control the routing based on query strings or headers. Here's an example for accessing query string data in a view function:

```
from chalice import Chalice  
  
app = Chalice(app_name='helloworld')  
  
@app.route('/users/{name}')  
def users(name):  
    result = {'name': name}  
    if app.current_request.query_params.get('include-greeting') == 'true':  
        result['greeting'] = 'Hello, %s' % name  
    return result
```

In the function above, if the user provides a `?include-greeting=true` in the HTTP request, then an additional greeting key will be returned:

```
$ http https://endpoint/api/users/bob  
  
{  
  "name": "bob"  
}  
  
$ http https://endpoint/api/users/bob?include-greeting=true  
  
{  
  "greeting": "Hello, bob",  
  "name": "bob"  
}
```


14.2 Views

A view function in chalice is the function attached to an `@app.route()` decorator. In the example below, `index` is the view function:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/')
def index():
    return {'view': 'index'}
```

14.2.1 View Function Parameters

A view function's parameters correspond to the number of captured URL parameters specified in the `@app.route` call. In the example above, the route `/` specifies no captured parameters so the `index` view function accepts no parameters. However, in the view function below, a single URL parameter, `{city}` is specified, so the view function must accept a single parameter:

```
from chalice import Chalice

app = Chalice(app_name='helloworld')

@app.route('/cities/{city}')
def index(city):
    return {'city': city}
```

This indicates that the value of `{city}` is variable, and whatever value is provided in the URL is passed to the `index` view function. For example:

```
GET /cities/seattle --> index('seattle')
GET /cities/portland --> index('portland')
```

If you want to access any other metadata of the incoming HTTP request, you can use the `app.current_request` property, which is an instance of the `Request` class.

14.2.2 View Function Return Values

The response returned back to the client depends on the behavior of the view function. There are several options available:

- Returning an instance of `Response`. This gives you complete control over what gets returned back to the customer.
- A bytes type response body must have a `Content-Type` header value that is present in the `app.api.binary_types` list in order to be handled properly.
- Any other return value will be serialized as JSON and sent back as the response body with content type `application/json`.
- Any subclass of `ChaliceViewError` will result in an HTTP response being returned with the status code associated with that response, and a JSON response body containing a `Code` and a `Message`. This is discussed in more detail below.

- Any other exception raised will result in a 500 HTTP response. The body of that response depends on whether debug mode is enabled.

14.2.3 Error Handling

Chalice provides a built in set of exception classes that map to common HTTP errors including:

- `BadRequestError`- returns a status code of 400
- `UnauthorizedError`- returns a status code of 401
- `ForbiddenError`- returns a status code of 403
- `NotFoundError`- returns a status code of 404
- `ConflictError`- returns a status code of 409
- `TooManyRequestsError`- returns a status code of 429
- `ChaliceViewError`- returns a status code of 500

You can raise these anywhere in your view functions and chalice will convert these to the appropriate HTTP response. The default chalice error responses will send the error back as `application/json` with the response body containing a `Code` corresponding to the exception class name and a `Message` key corresponding to the string provided when the exception was instantiated. For example:

```
from chalice import Chalice
from chalice import BadRequestError

app = Chalice(app_name="badrequest")

@app.route('/badrequest')
def badrequest():
    raise BadRequestError("This is a bad request")
```

This view function will generate the following HTTP response:

```
$ http https://endpoint/api/badrequest
HTTP/1.1 400 Bad Request

{
  "Code": "BadRequestError",
  "Message": "This is a bad request"
}
```

In addition to the built in chalice exceptions, you can use the `Response` class to customize the HTTP errors if you prefer to either not have JSON error responses or customize the JSON response body for errors. For example:

```
from chalice import Chalice, Response

app = Chalice(app_name="badrequest")

@app.route('/badrequest')
def badrequest():
    return Response(body='Plain text error message',
                    headers={'Content-Type': 'text/plain'},
                    status_code=400)
```

14.2.4 Specifying HTTP Methods

So far, our examples have only allowed GET requests. It's actually possible to support additional HTTP methods. Here's an example of a view function that supports PUT:

```
@app.route('/resource/{value}', methods=['PUT'])
def put_test(value):
    return {"value": value}
```

We can test this method using the http command:

```
$ http PUT https://endpoint/api/resource/foo
HTTP/1.1 200 OK

{
    "value": "foo"
}
```

Note that the `methods` kwarg accepts a list of methods. Your view function will be called when any of the HTTP methods you specify are used for the specified resource. For example:

```
@app.route('/myview', methods=['POST', 'PUT'])
def myview():
    pass
```

The above view function will be called when either an HTTP POST or PUT is sent to `/myview` as shown below:

```
POST /myview --> myview()
PUT /myview --> myview()
```

Alternatively if you do not want to share the same view function across multiple HTTP methods for the same route url, you may define separate view functions to the same route url but have the view functions differ by HTTP method. For example:

```
@app.route('/myview', methods=['POST'])
def myview_post():
    pass

@app.route('/myview', methods=['PUT'])
def myview_put():
    pass
```

This setup will route all HTTP POST's to `/myview` to the `myview_post()` view function and route all HTTP PUT's to `/myview` to the `myview_put()` view function as shown below:

```
POST /myview --> myview_post()
PUT /myview --> myview_put()
```

If you do chose to use separate view functions for the same route path, it is important to know:

- View functions that share the same route cannot have the same names. For example, two view functions that both share the same route path cannot both be named `view()`.
- View functions that share the same route cannot overlap in supported HTTP methods. For example if two view function both share the same route path, they both cannot contain `'PUT'` in their route `methods` list.
- View functions that share the same route path and have CORS configured cannot have differing CORS configuration. For example, if two view functions that both share the same route path, the route configuration for one of the view functions cannot set `cors=True` while having the route configuration of the other view function be set to `cors=app.CORSConfig(allow_origin='https://foo.example.com')`.

14.2.5 Binary Content

Chalice supports binary payloads through its `app.api.binary_types` list. Any type in this list is considered a binary Content-Type. Whenever a request with a Content-Type header is encountered that matches an entry in the `binary_types` list, its body will be available as a `bytes` type on the property `app.current_request.raw_body`. Similarly, in order to send binary data back in a response, simply set your Content-Type header to something present in the `binary_types` list. Note that you can override the default types by modifying the `app.api.binary_types` list at the module level.

Here is an example app which simply echos back binary content:

```
from chalice import Chalice, Response

app = Chalice(app_name="binary-response")

@app.route('/bin-echo', methods=['POST'],
           content_types=['application/octet-stream'])
def bin_echo():
    raw_request_body = app.current_request.raw_body
    return Response(body=raw_request_body,
                    status_code=200,
                    headers={'Content-Type': 'application/octet-stream'})
```

You can see this app echo back binary data sent to it:

```
$ echo -n -e "\xFE\xED" | http POST $(chalice url)bin-echo \
  Accept:application/octet-stream Content-Type:application/octet-stream | xxd
0000000: feed                                     ..
```

Note that both the Accept and Content-Type header are required. If you fail to set the Content-Type header on the request will result in a 415 `UnsupportedMediaType` error. Care must be taken when configuring what `content_types` a route accepts, they must all be valid binary types, or they must all be non-binary types. The Accept header must also be set if the data returned is to be the raw binary, if is omitted the call return a 400 Bad Request response.

For example, here is the same call as above without the Accept header:

```
$ echo -n -e "\xFE\xED" | http POST $(chalice url)bin-echo \
  Content-Type:application/octet-stream
HTTP/1.1 400 Bad Request
Connection: keep-alive
Content-Length: 270
Content-Type: application/json
Date: Sat, 27 May 2017 07:09:51 GMT

{
  "Code": "BadRequest",
  "Message": "Request did not specify an Accept header with
  application/octet-stream, The response has a Content-Type of
  application/octet-stream. If a response has a binary Content-Type then
  the request must specify an Accept header that matches."
}
```

14.2.6 Usage Recommendations

If you want to return a JSON response body, just return the corresponding python types directly. You don't need to use the `Response` class. Chalice will automatically convert this to a JSON HTTP response as a convenience for you.

Use the `Response` class when you want to return non-JSON content, or when you want to inject custom HTTP headers to your response.

For errors, raise the built in `ChaliceViewError` subclasses (e.g `BadRequestError`, `NotFoundError`, `ConflictError` etc) when you want to return a HTTP error response with a preconfigured JSON body containing a Code and Message.

Use the `Response` class when you want to customize the error responses to either return a different JSON error response body, or to return an HTTP response that's not `application/json`.

14.3 Configuration File

Whenever you create a new project using `chalice new-project`, a `.chalice` directory is created for you. In this directory is a `config.json` file that you can use to control what happens when you `chalice deploy`:

```
$ tree -a
.
-- .chalice
|   -- config.json
-- app.py
-- requirements.txt

1 directory, 3 files
```

14.3.1 Stage Specific Configuration

As of version 0.7.0 of chalice, you can specify configuration that is specific to a chalice stage as well as configuration that should be shared across all stages. See the [Chalice Stages](#) doc for more information about chalice stages.

- `stages` - This value of this key is a mapping of chalice stage name to stage configuration. Chalice assumes a default stage name of `dev`. If you run the `chalice new-project` command on chalice 0.7.0 or higher, this key along with the default `dev` key will automatically be created for you. See the examples section below for some stage specific configurations.

The following config values can either be specified per stage config or as a top level key which is not tied to a specific stage. Whenever a stage specific configuration value is needed, the `stages` mapping is checked first. If no value is found then the top level keys will be checked.

`api_gateway_endpoint_type`

The endpoint configuration of the deployed API Gateway which determines how the API will be accessed, can be `EDGE`, `REGIONAL`, `PRIVATE`. Note this value can only be set as a top level key and defaults to `EDGE`. For more information see <https://amzn.to/2LofApt>

`api_gateway_endpoint_vpce`

When configuring a Private API a VPC Endpoint id must be specified to configure a default resource policy on the API if an explicit policy is not specified. This value can be a list or a string of endpoint ids.

`api_gateway_policy_file`

A file pointing to an IAM resource policy for the REST API. If not specified chalice will autogenerate this policy when `endpoint_type` is `PRIVATE`. This filename is relative to the `.chalice` directory.

`api_gateway_stage`

The name of the API gateway stage. This will also be the URL prefix for your API (`https://endpoint/prefix/your-api`).

`autogen_policy`

A boolean value that indicates if chalice should try to automatically generate an IAM policy based on analyzing your application source code. The default value is `true`. If this value is `false` then chalice will load try to a local file in `.chalice/policy-<stage-name>.json` instead of auto-generating a policy from source code analysis.

`environment_variables`

A mapping of key value pairs. These key value pairs will be set as environment variables in your application. All environment variables must be strings. If this key is specified in both a stage specific config option as well as a top level key, the stage specific environment variables will be merged into the top level keys. See the examples section below for a concrete example.

`iam_policy_file`

When `autogen_policy` is `false`, chalice will try to load an IAM policy from disk instead of auto-generating one based on source code analysis. The default location of this file is `.chalice/policy-<stage-name>.json`, e.g. `.chalice/policy-dev.json`, `.chalice/policy-prod.json`, etc. You can change the filename by providing this `iam_policy_file` config option. This filename is relative to the `.chalice` directory.

`iam_role_arn`

If `manage_iam_role` is `false`, you must specify this value that indicates which IAM role arn to use when configuration your application. This value is only used if `manage_iam_role` is `false`.

`lambda_memory_size`

An integer representing the amount of memory, in MB, your Lambda function is given. AWS Lambda uses this memory size to infer the amount of CPU allocated to your function. The default `lambda_memory_size` value is 128. The value must be a multiple of 64 MB.

`lambda_timeout`

An integer representing the function execution time, in seconds, at which AWS Lambda should terminate the function. The default `lambda_timeout` is 60 seconds.

`layers`

A list of Lambda Layers arns. This value can be provided per stage as well as per Lambda function. See [AWS Lambda Layers Configuration](#).

`manage_iam_role`

true/false. Indicates if you want chalice to create and update the IAM role used for your application. By default, this value is `true`. However, if you have a pre-existing role you've created, you can set this value to `false` and a role will not be created or updated. `"manage_iam_role": false` means that you are responsible for managing the role and any associated policies associated with that role. If this value is `false` you must specify an `iam_role_arn`, otherwise an error is raised when you try to run `chalice deploy`.

`minimum_compression_size`

An integer value that indicates the minimum compression size to apply to the API gateway. If this key is specified in both a stage specific config option as well as a top level key, the stage specific key will override the top level key for the given stage. For more information check out the [Service Docs](#)

`reserved_concurrency`

An integer representing each function's reserved concurrency. This value can be provided per stage as well as per Lambda function. AWS Lambda reserves this value of concurrency to each lambda deployed in this stage. If the value is set to 0, invocations to this function are blocked. If the value is unset, there will be no reserved concurrency allocations. For more information, see [AWS Documentation on managing concurrency](#).

`subnet_ids`

A list of subnet ids for VPC configuration. This value can be provided per stage as well as per Lambda function. In order for this value to take effect, you must also provide the `security_group_ids` value. When both values are provided and `autogen_policy` is `True`, chalice will automatically update your IAM role with the necessary permissions to create, describe, and delete ENIs. If you are managing the IAM role policy yourself, make sure to update your permissions accordingly, as described in the [AWS Lambda VPC documentation](#).

`security_group_ids`

A list of security groups for VPC configuration. This value can be provided per stage as well as per Lambda function. In order for this value to take effect, you must also provide the `subnet_ids` value.

`tags`

A mapping of key value pairs. These key value pairs will be set as the tags on the resources running your deployed application. All tag keys and values must be strings. Similar to `environment_variables`, if a key is specified in both a stage specific config option as well as a top level key, the stage specific tags will be merged into the top level keys. By default, all chalice deployed resources are tagged with the key `'aws-chalice'` whose value is `'version={chalice-version}:stage={stage-name}:app={app-name}'`. Currently only the following chalice deployed resources are tagged: Lambda functions.

14.3.2 Lambda Specific Configuration

In addition to a chalice stage, there are also some configuration values that can be specified per Lambda function. A chalice app can have many stages, and a stage can have many Lambda functions.

You have the option to specify configuration for a lambda function across all your stages, or for a lambda function in a specific stage.

To configure per lambda configuration for a specific stage, you add a `lambda_functions` key in your stage configuration:

```
{
  "version": "2.0",
  "app_name": "app",
  "stages": {
    "dev": {
      "lambda_functions": {
        "foo": {
          "lambda_timeout": 120
        }
      }
    }
  }
}
```

To specify per lambda configuration across all stages, you add a top level `lambda_functions` key:

```
{
  "version": "2.0",
  "app_name": "app",
  "lambda_functions": {
    "foo": {
      "lambda_timeout": 120
    }
  }
}
```

Each key in the `lambda_functions` dictionary is the name of a Lambda function in your app. The value is a dictionary of configuration that will be applied to that function. These are the configuration options that can be applied per function:

- `autogen_policy`
- `environment_variables`
- `iam_policy_file`
- `iam_role_arn`
- `lambda_memory_size`
- `lambda_timeout`
- `layers`
- `manage_iam_role`
- `reserved_concurrency`
- `security_group_ids`
- `subnet_ids`
- `tags`

See the *Stage Specific Configuration* section above for a description of these config options.

In general, the name of your lambda function will correspond to the name of the function in your app. For example:

```
@app.lambda_function()
def foo(event, context):
    pass
```

To specify configuration for this function, you would use the key of `foo` in the `lambda_functions` configuration.

There is one exception to this, which is any python function decorated with the `@app.route()` decorator. Chalice uses a single Lambda function for all requests from API gateway, and this name is `api_handler`. So if you have an app like this:

```
@app.route('/')
def index(): pass

@app.route('/foo/bar')
def other_handler(): pass
```

Then to specify configuration values for the underlying lambda function, which `index()` and `other_handler()` share, you would specify:

```
{
  "lambda_functions": {
    "api_handler": {
      "subnet_ids": ["sn-1", "sn-2"],
      "security_group_ids": ["sg-10", "sg-11"],
      "layers": ["layer-arn-1", "layer-arn-2"],
    }
  }
}
```

14.3.3 Examples

Below are examples that show how you can configure your chalice app.

IAM Roles and Policies

Here's an example for configuring IAM policies across stages:

```
{
  "version": "2.0",
  "app_name": "app",
  "stages": {
    "dev": {
      "autogen_policy": true,
      "api_gateway_stage": "dev"
    },
    "beta": {
      "autogen_policy": false,
      "iam_policy_file": "beta-app-policy.json"
    },
    "prod": {
      "manage_iam_role": false,
      "iam_role_arn": "arn:aws:iam::...:role/prod-role"
    }
  }
}
```

```
}  
}
```

In this config file we're specifying three stages, `dev`, `beta`, and `prod`. In the `dev` stage, chalice will automatically generate an IAM policy based on analyzing the application source code. For the `beta` stage, chalice will load the `.chalice/beta-app-policy.json` file and use it as the policy to associate with the IAM role for that stage. In the `prod` stage, chalice won't modify any IAM roles. It will just set the IAM role for the Lambda function to be `arn:aws:iam::...:role/prod-role`.

Here's an example that show config precedence:

```
{  
  "version": "2.0",  
  "app_name": "app",  
  "api_gateway_stage": "api",  
  "stages": {  
    "dev": {  
    },  
    "beta": {  
    },  
    "prod": {  
      "api_gateway_stage": "prod",  
      "manage_iam_role": false,  
      "iam_role_arn": "arn:aws:iam::...:role/prod-role"  
    }  
  }  
}
```

In this config file, both the `dev` and `beta` stage will have an API gateway stage name of `api` because they will default to the top level `api_gateway_stage` key. However, the `prod` stage will have an API gateway stage name of `prod` because the `api_gateway_stage` is specified in `{"stages": {"prod": ...}}` mapping.

Environment Variables

In the following example, environment variables are specified both as top level keys as well as per stage. This allows us to provide environment variables that all stages should have as well as stage specific environment variables:

```
{  
  "version": "2.0",  
  "app_name": "app",  
  "environment_variables": {  
    "SHARED_CONFIG": "foo",  
    "OTHER_CONFIG": "from-top"  
  },  
  "stages": {  
    "dev": {  
      "environment_variables": {  
        "TABLE_NAME": "dev-table",  
        "OTHER_CONFIG": "dev-value"  
      }  
    },  
    "prod": {  
      "environment_variables": {  
        "TABLE_NAME": "prod-table",  
        "OTHER_CONFIG": "prod-value"  
      }  
    }  
  }  
}
```

```
}
}
```

For the above config, the dev stage will have the following environment variables set:

```
{
  "SHARED_CONFIG": "foo",
  "TABLE_NAME": "dev-table",
  "OTHER_CONFIG": "dev-value",
}
```

The prod stage will have these environment variables set:

```
{
  "SHARED_CONFIG": "foo",
  "TABLE_NAME": "prod-table",
  "OTHER_CONFIG": "prod-value",
}
```

Per Lambda Examples

Suppose we had the following chalice app:

```
from chalice import Chalice

app = Chalice(app_name='demo')

@app.lambda_function()
def foo(event, context):
    pass

@app.lambda_function()
def bar(event, context):
    pass
```

Given these two functions, we'd like to configure the functions as follows:

- Both functions should have an environment variable `OWNER` with value `dev-team`.
- The `foo` function should have an autogenerated IAM policy managed by chalice.
- The `foo` function should be run in a VPC with subnet ids `sn-1` and `sn-2`, with security groups `sg-10` and `sg-11`. Chalice should also automatically configure the IAM policy with permissions to modify EC2 network interfaces.
- The `foo` function should have two connected layers as `layer-arn-1` and `layer-arn-2`. Chalice should automatically configure the IAM policy.
- The `bar` function should use a pre-existing IAM role that was created outside of chalice. Chalice should not perform an IAM role management for the `bar` function.
- The `bar` function should have an environment variable `TABLE_NAME` with value `mytable`.

We can accomplish all this with this config file:

```
{
  "stages": {
    "dev": {
      "environment_variables": {
        "OWNER": "dev-team"
      }
    }
  }
}
```

```
}
"api_gateway_stage": "api",
"lambda_functions": {
  "foo": {
    "subnet_ids": ["sn-1", "sn-2"],
    "security_group_ids": ["sg-10", "sg-11"],
    "layers": ["layer-arn-1", "layer-arn-2"],
  },
  "bar": {
    "manage_iam_role": false,
    "iam_role_arn": "arn:aws:iam::my-role-name",
    "environment_variables": {"TABLE_NAME": "mytable"}
  }
}
},
"version": "2.0",
"app_name": "demo"
}
```

14.4 Multifile Support

The `app.py` file contains all of your view functions and route information, but you don't have to keep all of your application code in your `app.py` file.

As your application grows, you may reach out a point where you'd prefer to structure your application in multiple files. You can create a `chalicelib/` directory, and anything in that directory is recursively included in the deployment package. This means that you can have files besides just `.py` files in `chalicelib/`, including `.json` files for config, or any kind of binary assets.

Let's take a look at a few examples.

Consider the following app directory structure layout:

```
.
-- app.py
-- chalicelib
|   -- __init__.py
-- requirements.txt
```

Where `chalicelib/__init__.py` contains:

```
MESSAGE = 'world'
```

and the `app.py` file contains:

```
1 from chalice import Chalice
2 from chalicelib import MESSAGE
3
4 app = Chalice(app_name="multifile")
5
6 @app.route("/")
7 def index():
8     return {"hello": MESSAGE}
```

Note in line 2 we're importing the `MESSAGE` variable from the `chalicelib` package, which is a top level directory in our project. We've created a `chalicelib/__init__.py` file which turns the `chalicelib` directory into a python package.

We can also use this directory to store config data. Consider this app structure layout:

```
.
-- app.py
-- chaliceelib
|   -- config.json
-- requirements.txt
```

With `chaliceelib/config.json` containing:

```
{"message": "world"}
```

In our `app.py` code, we can load and use our config file:

```
1 import os
2 import json
3
4 from chalice import Chalice
5
6 app = Chalice(app_name="multifile")
7
8 filename = os.path.join(
9     os.path.dirname(__file__), 'chaliceelib', 'config.json')
10 with open(filename) as f:
11     config = json.load(f)
12
13 @app.route("/")
14 def index():
15     # We can access ``config`` here if we want.
16     return {"hello": config['message']}
```

14.5 Logging

You have several options for logging in your application. You can use any of the options available to lambda functions as outlined in the [AWS Lambda Docs](#). The simplest option is to just use print statements. Anything you print will be accessible in cloudwatch logs as well as in the output of the `chalice logs` command.

In addition to using the `stdlib` logging module directly, the framework offers a preconfigured logger designed to work nicely with Lambda. This is offered purely as a convenience, you can use `print` or the logging module directly if you prefer.

You can access this logger via the `app.log` attribute, which is a logger specifically for your application. This attribute is an instance of `logging.getLogger(your_app_name_)` that's been preconfigured with reasonable defaults:

- StreamHandler associated with `sys.stdout`.
- Log level set to `logging.ERROR` by default. You can also manually set the logging level by setting `app.log.setLevel(logging.DEBUG)`.
- A logging formatter that displays the app name, level name, and message.

14.5.1 Examples

In the following application, we're using the application logger to emit two log messages, one at `DEBUG` and one at the `ERROR` level:

```
from chalice import Chalice

app = Chalice(app_name='demolog')

@app.route('/')
def index():
    app.log.debug("This is a debug statement")
    app.log.error("This is an error statement")
    return {'hello': 'world'}
```

If we make a request to this endpoint, and then look at `chalice logs` we'll see the following log message:

```
2016-11-06 20:24:25.490000 9d2a92 demolog - ERROR - This is an error statement
```

As you can see, only the ERROR level log is emitted because the default log level is ERROR. Also note the log message formatting. This is the default format that's been automatically configured. We can make a change to set our log level to debug:

```
from chalice import Chalice

app = Chalice(app_name='demolog')
# Enable DEBUG logs.
app.log.setLevel(logging.DEBUG)

@app.route('/')
def index():
    app.log.debug("This is a debug statement")
    app.log.error("This is an error statement")
    return {'hello': 'world'}
```

Now if we make a request to the `/` URL and look at the output of `chalice logs`, we'll see the following log message:

```
2016-11-07 12:29:15.714 431786 demolog - DEBUG - This is a debug statement
2016-11-07 12:29:15.714 431786 demolog - ERROR - This is an error statement
```

As you can see here, both the debug and error log message are shown.

You can use the `-n/--name` option to view the logs for a specific lambda function. By default, the logs for the API handler lambda function are shown. This corresponds to any log statements made within an `@app.route()` call. The name option is the logical name of the lambda function. This is the name of the python function by default, or whatever name you provided as the name kwarg to the `@app.lambda_function()` call. For example, given this app:

```
from chalice import Chalice

app = Chalice(app_name='multilog')

@app.lambda_function()
def foo(event, context):
    app.log.debug("Invoking from function foo")
    return {'hello': 'world'}

@app.lambda_function(name='MyFunction')
def bar(event, context):
```

```
incr_counter()
app.log.debug("Invoking from function bar")
return {'hello': 'world'}
```

You can retrieve logs for the above function by running:

```
$ chalice logs --name foo
$ chalice logs --name MyFunction
```

14.6 SDK Generation

The `@app.route(...)` information you provide chalice allows it to create corresponding routes in API Gateway. One of the benefits of this approach is that we can leverage API Gateway's SDK generation process. Chalice offers a `chalice generate-sdk` command that will automatically generate an SDK based on your declared routes.

Note: The only supported language at this time is javascript.

Keep in mind that chalice itself does not have any logic for generating SDKs. The SDK generation happens service side in [API Gateway](#), the `chalice generate-sdk` is just a high level wrapper around that functionality.

To generate an SDK for a chalice app, run this command from the project directory:

```
$ chalice generate-sdk /tmp/sdk
```

You should now have a generated javascript sdk in `/tmp/sdk`. API Gateway includes a `README.md` as part of its SDK generation which contains details on how to use the javascript SDK.

14.6.1 Example

Suppose we have the following chalice app:

```
from chalice import Chalice

app = Chalice(app_name='sdktest')

@app.route('/', cors=True)
def index():
    return {'hello': 'world'}

@app.route('/foo', cors=True)
def foo():
    return {'foo': True}

@app.route('/hello/{name}', cors=True)
def hello_name(name):
    return {'hello': name}

@app.route('/users/{user_id}', methods=['PUT'], cors=True)
def update_user(user_id):
    return {"msg": "fake updated user", "userId": user_id}
```

Let's generate a javascript SDK and test it out in the browser. Run the following command from the project dir:

```
$ chalice generate-sdk /tmp/sdkdemo
$ cd /tmp/sdkdemo
$ ls -la
-rw-r--r--  1 jamessar  r   3227 Nov 21 17:06 README.md
-rw-r--r--  1 jamessar  r   9243 Nov 21 17:06 apigClient.js
drwxr-xr-x  6 jamessar  r    204 Nov 21 17:06 lib
```

You should now be able to follow the instructions from API Gateway in the `README.md` file. Below is a snippet that shows how the generated javascript SDK methods correspond to the `@app.route()` calls in chalice.

```
<script type="text/javascript">
  // Below are examples of how the javascript SDK methods
  // correspond to chalice @app.routes()
  var apigClient = apigClientFactory.newClient();

  // @app.route('/')
  apigClient.rootGet().then(result => {
    document.getElementById('root-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/foo')
  apigClient.fooGet().then(result => {
    document.getElementById('foo-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/hello/{name}')
  apigClient.helloNameGet({name: 'jimmy'}).then(result => {
    document.getElementById('helloname-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/users/{user_id}', methods=['PUT'])
  apigClient.usersUserIdPut({user_id: '123'}, 'body content').then(result => {
    document.getElementById('users-userid-put').innerHTML = JSON.stringify(result.data);
  });
</script>
```

Example HTML File

If you want to try out the example above, you can use the following `index.html` page to test:

```
<!DOCTYPE html>
<html lang="en">
  <head>
    <title>SDK Test</title>
    <meta charset="UTF-8">
    <meta name="viewport" content="width=device-width, initial-scale=1">
    <link rel="stylesheet" href="https://cdnjs.cloudflare.com/ajax/libs/skeleton/2.0.4/skeleton.r
    <script type="text/javascript" src="lib/axios/dist/axios.standalone.js"></script>
    <script type="text/javascript" src="lib/CryptoJS/rollups/hmac-sha256.js"></script>
    <script type="text/javascript" src="lib/CryptoJS/rollups/sha256.js"></script>
    <script type="text/javascript" src="lib/CryptoJS/components/hmac.js"></script>
    <script type="text/javascript" src="lib/CryptoJS/components/enc-base64.js"></script>
    <script type="text/javascript" src="lib/url-template/url-template.js"></script>
    <script type="text/javascript" src="lib/apiGatewayCore/sigV4Client.js"></script>
    <script type="text/javascript" src="lib/apiGatewayCore/apiGatewayClient.js"></script>
    <script type="text/javascript" src="lib/apiGatewayCore/simpleHttpClient.js"></script>
    <script type="text/javascript" src="lib/apiGatewayCore/utils.js"></script>
```



```

<script type="text/javascript" src="apigClient.js"></script>

<script type="text/javascript">
  // Below are examples of how the javascript SDK methods
  // correspond to chalice @app.routes()
  var apigClient = apigClientFactory.newClient();

  // @app.route('/')
  apigClient.rootGet().then(result => {
    document.getElementById('root-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/foo')
  apigClient.fooGet().then(result => {
    document.getElementById('foo-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/hello/{name}')
  apigClient.helloNameGet({name: 'jimmy'}).then(result => {
    document.getElementById('helloname-get').innerHTML = JSON.stringify(result.data);
  });

  // @app.route('/users/{user_id}', methods=['PUT'])
  apigClient.usersUserIdPut({user_id: '123'}, 'body content').then(result => {
    document.getElementById('users-userid-put').innerHTML = JSON.stringify(result.data);
  });
</script>
</head>
<body>
  <div><h5>result of rootGet()</h5><pre id="root-get"></pre></div>
  <div><h5>result of fooGet()</h5><pre id="foo-get"></pre></div>
  <div><h5>result of helloNameGet({name: 'jimmy'})</h5><pre id="helloname-get"></pre></div>
  <div><h5>result of usersUserIdPut({user_id: '123'})</h5><pre id="users-userid-put"></pre></div>
</body>
</html>

```

14.7 Chalice Stages

Chalice has the concept of stages, which are completely separate sets of AWS resources. When you first create a chalice project and run commands such as `chalice deploy` and `chalice url`, you don't have to specify any stage values or stage configuration. This is because chalice will use a stage named `dev` by default.

You may eventually want to have multiple stages of your application. A common configuration would be to have a `dev`, `beta` and `prod` stage. A `dev` stage would be used by developers to test out new features. Completed features would be deployed to `beta`, and the `prod` stage would be used for serving production traffic.

Chalice can help you manage this.

To create a new chalice stage, specify the `--stage` argument. If the stage does not exist yet, it will be created for you:

```
$ chalice deploy --stage prod
```

By creating a new chalice stage, a new API Gateway rest API, Lambda function, and potentially (depending on config settings) a new IAM role will be created for you.

14.7.1 Example

Let's say we have a new app:

```
$ chalice new-project myapp
$ cd myapp
$ chalice deploy
...
https://mmnkdi.execute-api.us-west-2.amazonaws.com/api/
```

We've just created our first stage, dev. We can iterate on our application and continue to run `chalice deploy` to deploy our code to the dev stage. Let's say we want to now create a prod stage. To do this, we can run:

```
$ chalice deploy --stage prod
...
https://wk9fhx.execute-api.us-west-2.amazonaws.com/api/
```

We now have two completely separate rest APIs:

```
$ chalice url --stage dev
https://mmnkdi.execute-api.us-west-2.amazonaws.com/api/

$ chalice url --stage prod
https://wk9fhx.execute-api.us-west-2.amazonaws.com/api/
```

Additionally, we can see all our deployed values by looking at the `.chalice/deployed/dev.json` or `.chalice/deployed/prod.json` files:

```
$ cat .chalice/deployed/dev.json
{
  "resources": [
    {
      "name": "api_handler",
      "resource_type": "lambda_function",
      "lambda_arn": "arn:aws:lambda:...:function:myapp-dev"
    },
    {
      "name": "rest_api",
      "resource_type": "rest_api",
      "rest_api_id": "wk9fhx",
      "rest_api_url": "https://wk9fhx.execute-api.us-west-2.amazonaws.com/api/"
    }
  ],
  "schema_version": "2.0",
  "backend": "api"
}

$ cat .chalice/deployed/prod.json
{
  "resources": [
    {
      "name": "api_handler",
      "resource_type": "lambda_function",
      "lambda_arn": "arn:aws:lambda:...:function:myapp-prod"
    },
    {
      "name": "rest_api",
      "resource_type": "rest_api",
      "rest_api_id": "mmnkdi",

```

```

    "rest_api_url": "https://mmnkdi.execute-api.us-west-2.amazonaws.com/api/"
  }
],
"schema_version": "2.0",
"backend": "api"
}

```

14.8 App Packaging

In order to deploy your Chalice app, a zip file is created that contains your application and all third party packages your application requires. This file is used by AWS Lambda and is referred to as a deployment package.

Chalice will automatically create this deployment package for you, and offers several features to make this easier to manage. Chalice allows you to clearly separate application specific modules and packages you are writing from 3rd party package dependencies.

14.8.1 App Directories

You have two options to structure application specific code/config:

- **app.py** - This file includes all your route information and is always included in the deployment package.
- **chalicelib/** - This directory (if it exists) is included in the deployment package. This is where you can add config files and additional application modules if you prefer not to have all your app code in the `app.py` file.

See [Multifile Support](#) for more info on the `chalicelib/` directory. Both the `app.py` and the `chalicelib/` directory are intended for code that you write yourself.

14.8.2 3rd Party Packages

There are two options for handling python package dependencies:

- **requirements.txt** - During the packaging process, Chalice will install any packages it finds or can build compatible wheels for. Specifically all pure python packages as well as all packages that upload wheel files for the `manylinux1_x86_64` platform will be automatically installable.
- **vendor/** - The *contents* of this directory are automatically added to the top level of the deployment package.

Chalice will also check for an optional `vendor/` directory in the project root directory. The contents of this directory are automatically included in the top level of the deployment package (see [Examples](#) for specific examples). The `vendor/` directory is helpful in these scenarios:

- You need to include custom packages or binary content that is not accessible via `pip`. These may be internal packages that aren't public.
- Wheel files are not available for a package you need from `pip`.
- A package is installable with `requirements.txt` but has optional `c` extensions. Chalice can build the dependency without the `c` extensions, but if you want better performance you can vendor a version that is compiled.

As a general rule of thumb, code that you write goes in either `app.py` or `chalicelib/`, and dependencies are either specified in `requirements.txt` or placed in the `vendor/` directory.

14.8.3 Examples

Suppose I have the following app structure:

```
.
-- app.py
-- chaliceelib
|   -- __init__.py
|   -- utils.py
-- requirements.txt
-- vendor
    -- internalpackage
        -- __init__.py
```

And the `requirements.txt` file had one requirement:

```
$ cat requirements.txt
sortedcontainers==1.5.4
```

Then the final deployment package directory structure would look like this:

```
.
-- app.py
-- chaliceelib
|   -- __init__.py
|   -- utils.py
-- internalpackage
|   -- __init__.py
-- sortedcontainers
    -- __init__.py
```

This directory structure is then zipped up and sent to AWS Lambda during the deployment process.

14.8.4 Cryptography Example

Note

Since the original version of this example was written, cryptography has released version 2.0 which does have manylinux1 wheels available. This means if you want to use cryptography in a Chalice app all you need to do is add `cryptography` or `cryptography>=2.0` in your `requirements.txt` file.

This example will use version 1.9 of Cryptography because it is a good example of a library with C extensions and no wheel files.

Below shows an example of how to use the [cryptography 1.9](#) package in a Chalice app for the `python3.6` lambda environment.

Suppose you are on a Mac or Windows and want to deploy a Chalice app that depends on the `cryptography==1.9` package. If you simply add it to your `requirements.txt` file and try to deploy it with `chalice deploy` you will get the following warning during deployment:

```
$ cat requirements.txt
cryptography==1.9
$ chalice deploy
Updating IAM policy.
Updating lambda function...
Creating deployment package.

Could not install dependencies:
cryptography==1.9
```

```
You will have to build these yourself and vendor them in
the chalice vendor folder.
```

```
Your deployment will continue but may not work correctly
if missing dependencies are not present. For more information:
http://chalice.readthedocs.io/en/latest/topics/packaging.html
```

This happened because the `cryptography` version 1.9 does not have wheel files available on PyPi, and has C extensions. Since we are not on the same platform as AWS Lambda, the compiled C extensions Chalice built were not compatible. To get around this we are going to leverage the `vendor/` directory, and build the `cryptography` package on a compatible linux system.

You can do this yourself by building `cryptography` on an Amazon Linux instance running in EC2. All of the following commands were run inside a `python 3.6` virtual environment.

- Download the source first:

```
$ pip download cryptography==1.9
```

This will download all the requirements into the current working directory. The directory should have the following contents:

```
- asn1crypto-0.22.0-py2.py3-none-any.whl
- cffi-1.10.0-cp36-cp36m-manylinux1_x86_64.whl
- cryptography-1.9.tar.gz
- idna-2.5-py2.py3-none-any.whl
- pycparser-2.17.tar.gz
- six-1.10.0-py2.py3-none-any.whl
```

This is a complete set of dependencies required for the `cryptography` package. Most of these packages have wheels that were downloaded, which means they can simply be put in the `requirements.txt` and Chalice will take care of downloading them. That leaves `cryptography` itself and `pycparser` as the only two that did not have a wheel file available for download.

- Next build the `cryptography` source package into a wheel file:

```
$ pip wheel cryptography-1.9.tar.gz
```

This will take a few seconds and build a wheel file for both `cryptography` and `pycparser`. The directory should now have two additional wheel files:

```
- cryptography-1.9-cp36-cp36m-linux_x86_64.whl
- pycparser-2.17-py2.py3-none-any.whl
```

The `cryptography` wheel file has been built with a compatible architecture for Lambda (`linux_x86_64`) and the `pycparser` has been built for any architecture which means it can also be automatically packaged by Chalice if it is listed in the `requirements.txt` file.

- Download the `cryptography` wheel file from the Amazon Linux instance and unzip it into the `vendor/` directory in the root directory of your Chalice app.

You should now have a project directory that looks like this:

```
$ tree
.
-- app.py
-- requirements.txt
```

```
-- vendor
  -- cryptography
  |   -- ... Lots of files
  |
  -- cryptography-1.9.dist-info
    -- DESCRIPTION.rst
    -- METADATA
    -- RECORD
    -- WHEEL
    -- entry_points.txt
    -- metadata.json
    -- top_level.txt
```

The `requirements.txt` file should look like this:

```
$ cat requirements.txt
cffi==1.10.0
six==1.10.0
asn1crypto==0.22.0
idna==2.5
pycparser==2.17
```

In your `app.py` file you can now import `cryptography`, and these dependencies will all get included when the `chalice deploy` command is run.

14.9 Python Version Support

Chalice supports all versions of python supported by AWS Lambda, which is currently `python2.7` and `python3.6`.

Chalice will automatically pick which version of python to use for Lambda based on the major version of python you are using. You don't have to explicitly configure which version of python you want to use. For example:

```
$ python --version
Python 3.6.1
$ chalice new-project test-versions
$ cd test-versions
$ chalice package test-package
$ grep -C 3 python test-package/sam.json
    "APIHandler": {
        "Type": "AWS::Serverless::Function",
        "Properties": {
            "Runtime": "python3.6",
            "Handler": "app.app",
            "CodeUri": "./deployment.zip",
            "Events": {

# Similarly, if we were to run "chalice deploy" we'd
# use python3.6 for the runtime.
$ chalice --debug deploy
Initiating first time deployment...
Deploying to: dev
...
"Runtime": "python3.6"
...
https://rest-api-id.execute-api.us-west-2.amazonaws.com/api/
```

In the example above, we're using python 3.6.1 so chalice automatically selects the `python3.6` runtime for lambda. If we were using python 2.7.11, chalice would automatically select `python2.7` as the runtime.

Chalice will emit a warning if the minor version does not match a python version supported by Lambda. Chalice will select the closest Lambda version in this scenario, as shown in the table below.

Local Python Version	Lambda Python Runtime
python2.7.10	python2.7
python2.7.11	python2.7
python2.7.12	python2.7
python2.7.13	python2.7
python3.3.6	python3.6
python3.4.6	python3.6
python3.5.3	python3.6
python3.6.0	python3.6
python3.6.1	python3.6

We strongly encourage you to develop your application using the same major/minor version of python you plan on using on AWS Lambda.

14.10 Changing Python Runtime Versions

The version of the python runtime to use in AWS Lambda can be reconfigured whenever you deploy your chalice app. This allows you to migrate to python3 in AWS Lambda by creating a new virtual environment that uses python3. For example, suppose you have an existing chalice app that uses python2:

```
$ python --version
Python 2.7.12
$ chalice deploy
...
https://endpoint/api
```

To upgrade the application to use python3, create a python3 virtual environment and redeploy.

```
:: $ deactivate $ python3 -m venv /tmp/venv3 $ source /tmp/venv3/bin/activate $ python --version Python 3.6.1 $
chalice deploy ...
```

14.11 AWS CloudFormation Support

When you run `chalice deploy`, chalice will deploy your application using the [AWS SDK for Python](#)). Chalice also provides functionality that allows you to manage deployments yourself using cloudformation. This is provided via the `chalice package` command.

When you run this command, chalice will generate the AWS Lambda deployment package that contains your application as well as a [Serverless Application Model \(SAM\)](#) template. You can then use a tool like the AWS CLI, or any cloudformation deployment tools you use, to deploy your chalice application.

14.11.1 Considerations

Using the `chalice package` command is useful when you don't want to use `chalice deploy` to manage your deployments. There's several reasons why you might want to do this:

- You have pre-existing infrastructure and tooling set up to manage cloudformation stacks.

- You want to integrate with other cloudformation stacks to manage all your AWS resources, including resources outside of your chalice app.
- You'd like to integrate with [AWS CodePipeline](#) to automatically deploy changes when you push to a git repo.

Keep in mind that you can't switch between `chalice deploy` and `chalice package` + CloudFormation for deploying your app.

If you choose to use `chalice package` and CloudFormation to deploy your app, you won't be able to switch back to `chalice deploy`. Running `chalice deploy` would create an entirely new set of AWS resources (API Gateway Rest API, AWS Lambda function, etc).

14.11.2 Template Merge

It's a common use case to need to modify a Chalice generated template before deployment. Often to inject extra resources, values, or configurations that are not supported directly by Chalice. It will always be the case that something on AWS is not supported by Chalice directly that a consumer may want to interact with.

The package command can now be invoked with the `--merge-template` argument:

```
$ chalice package --merge-template extras.json out
```

This `extras.json` file should be a JSON formatted file which will be deep-merged on top of the `sam.json` that is generated by Chalice.

For a simple example lets assume that we have the default new Chalice project and that `extras.json` has the following content:

```
{
  "Resources" : {
    "MusicTable" : {
      "Type" : "AWS::DynamoDB::Table",
      "Properties" : {
        "TableName" : "MusicData",
        "AttributeDefinitions" : [
          {
            "AttributeName" : "Album",
            "AttributeType" : "S"
          },
          {
            "AttributeName" : "Artist",
            "AttributeType" : "S"
          }
        ],
        "KeySchema" : [
          {
            "AttributeName" : "Album",
            "KeyType" : "HASH"
          },
          {
            "AttributeName" : "Artist",
            "KeyType" : "RANGE"
          }
        ],
        "ProvisionedThroughput" : {
          "ReadCapacityUnits" : "5",
          "WriteCapacityUnits" : "5"
        }
      }
    }
  }
}
```



```

    },
    "APIHandler": {
      "Properties": {
        "Environment": {
          "Variables": {
            "MUSIC_TABLE": {"Ref": "MusicTable"}
          }
        }
      }
    }
  }
}

```

The generated template located at `out/sam.json` will have the DynamoDB table injected into the resource section, as well as the `MUSIC_TABLE` environment variable added to the `APIHandler` Lambda function:

```

...
    "APIHandler": {
      "Type": "AWS::Serverless::Function",
      "Properties": {
        "Runtime": "python3.6",
        "Handler": "app.app",
        "CodeUri": "./deployment.zip",
        "Tags": {
          "aws-chalice": "version=1.10-:stage=dev:app=test"
        },
        "Timeout": 60,
        "MemorySize": 128,
        "Role": {
          "Fn::GetAtt": [
            "DefaultRole",
            "Arn"
          ]
        },
        "Environment": {
          "Variables": {
            "MUSIC_TABLE": "MusicData"
          }
        }
      }
    }
  },
...

```

This gives us the ability to inject arbitrary resources into our Chalice applications, and reference them from our Chalice-deployed functions. We can now rely on Chalice's policy auto-generation to generate a policy that allows DynamoDB access, inject our own policy modifications through the same `extras.json` file, or specify a custom policy using the config file.

14.11.3 Example

In this example, we'll create a chalice app and deploy it using the AWS CLI.

First install the necessary packages:

```

$ virtualenv /tmp/venv
$ . /tmp/venv/bin/activate
$ pip install chalice awscli

```

```
$ chalice new-project test-cfn-deploy
$ cd test-cfn-deploy
```

At this point we've installed chalice and the AWS CLI and we have a basic app created locally. Next we'll run the package command and look at its contents:

```
$ $ chalice package /tmp/packaged-app/
Creating deployment package.
$ ls -la /tmp/packaged-app/
-rw-r--r--  1 j      wheel  3355270 May 25 14:20 deployment.zip
-rw-r--r--  1 j      wheel    3068 May 25 14:20 sam.json

$ unzip -l /tmp/packaged-app/deployment.zip | tail -n 5
 17292  05-25-17 14:19  chalice/app.py
   283   05-25-17 14:19  chalice/__init__.py
   796   05-25-17 14:20  app.py
-----
9826899                                723 files

$ head < /tmp/packaged-app/sam.json
{
  "AWSTemplateFormatVersion": "2010-09-09",
  "Outputs": {
    "RestAPIId": {
      "Value": {
        "Ref": "RestAPI"
      }
    },
    "APIHandlerName": {
      "Value": {
```

As you can see in the above example, the package command created a directory that contained two files, a `deployment.zip` file, which is the Lambda deployment package, and a `sam.json` file, which is the SAM template that can be deployed using CloudFormation. Next we're going to use the AWS CLI to deploy our app. To this, we'll first run the `aws cloudformation package` command, which will take our `deployment.zip` file and upload to an S3 bucket we specify:

```
$ aws cloudformation package \
  --template-file /tmp/packaged-app/sam.json \
  --s3-bucket myapp-bucket \
  --output-template-file /tmp/packaged-app/packaged.yaml
```

Now we can deploy our app using the `aws cloudformation deploy` command:

```
$ aws cloudformation deploy \
  --template-file /tmp/packaged-app/packaged.yaml \
  --stack-name test-cfn-stack \
  --capabilities CAPABILITY_IAM
Waiting for changeset to be created..
Waiting for stack create/update to complete
Successfully created/updated stack - test-cfn-stack
```

This will take a few minutes to complete, but once it's done, the endpoint url will be available as an output:

```
$ aws cloudformation describe-stacks --stack-name test-cfn-stack \
  --query "Stacks[0].Outputs[?OutputKey=='EndpointURL'][0].OutputValue"
"https://abc29hkq0i.execute-api.us-west-2.amazonaws.com/api/"

$ http "https://abc29hkq0i.execute-api.us-west-2.amazonaws.com/api/"
```

```
HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 18
Content-Type: application/json
...
{
  "hello": "world"
}
```

14.12 Terraform Support

When you run `chalice deploy`, chalice will deploy your application using the [AWS SDK for Python](#)). Chalice also provides functionality that allows you to manage deployments yourself using terraform. This is provided via the `chalice package --pkg-format terraform` command.

When you run this command, chalice will generate the AWS Lambda deployment package that contains your application and a [Terraform](#) configuration file. You can then use the terraform cli to deploy your chalice application.

14.12.1 Considerations

Using the `chalice package` command is useful when you don't want to use `chalice deploy` to manage your deployments. There's several reasons why you might want to do this:

- You have pre-existing infrastructure and tooling set up to manage Terraform deployments.
- You want to integrate with other Terraform resources to manage all your application resources, including resources outside of your chalice app.
- You'd like to integrate with [AWS CodePipeline](#) to automatically deploy changes when you push to a git repo.

Keep in mind that you can't switch between `chalice deploy` and `chalice package` + Terraform for deploying your app.

If you choose to use `chalice package` and Terraform to deploy your app, you won't be able to switch back to `chalice deploy`. Running `chalice deploy` would create an entirely new set of AWS resources (API Gateway Rest API, AWS Lambda function, etc).

14.12.2 Example

In this example, we'll create a chalice app and deploy it using the AWS CLI.

First install the necessary packages:

```
$ virtualenv /tmp/venv
$ . /tmp/venv/bin/activate
$ pip install chalice awscli
$ chalice new-project test-tf-deploy
$ cd test-tf-deploy
```

At this point we've installed chalice and the AWS CLI and we have a basic app created locally. Next we'll run the `package` command:

```
$ chalice package --pkg-format terraform /tmp/packaged-app/  
Creating deployment package.  
$ ls -la /tmp/packaged-app/  
-rw-r--r--  1 j          wheel  3355270 May 25 14:20 deployment.zip  
-rw-r--r--  1 j          wheel    3068 May 25 14:20 chalice.tf.json  
  
$ unzip -l /tmp/packaged-app/deployment.zip | tail -n 5  
   17292  05-25-17 14:19  chalice/app.py  
     283  05-25-17 14:19  chalice/__init__.py  
     796  05-25-17 14:20  app.py  
-----  
 9826899                      723 files
```

As you can see in the above example, the `package --pkg-format` command created a directory that contained two files, a `deployment.zip` file, which is the Lambda deployment package, and a `chalice.tf.json` file, which is the Terraform template that can be deployed using Terraform. Next we're going to use the Terraform CLI to deploy our app.

Note terraform will deploy run against all terraform files in this directory, so we can add additional resources for our application by adding terraform additional files here. The Chalice terraform template includes two static data values (*app* and *stage* names) that we can optionally use when constructing these additional resources, ie. `${data.null_data_source.chalice.outputs.app}`

First let's run Terraform init to install the AWS Terraform Provider:

```
$ cd /tmp/packaged-app  
$ terraform init
```

Now we can deploy our app using the `terraform apply` command:

```
$ terraform apply  
data.aws_region.chalice: Refreshing state...  
data.aws_caller_identity.chalice: Refreshing state...  
  
An execution plan has been generated and is shown below.  
Resource actions are indicated with the following symbols:  
+ create  
  
... (omit plan output)  
  
Plan: 14 to add, 0 to change, 0 to destroy.  
  
Do you want to perform these actions?  
  Terraform will perform the actions described above.  
  Only 'yes' will be accepted to approve.  
Enter a value: yes  
  
... (omit apply output)  
  
Apply complete! Resources: 14 added, 0 changed, 0 destroyed.  
  
Outputs:  
  
EndpointURL = https://7bnxriulj5.execute-api.us-east-1.amazonaws.com/dev
```

This will take a minute to complete, but once it's done, the endpoint url will be available as an output which we can then curl:

```
$ http "$(terraform output EndpointURL)"
HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 18
Content-Type: application/json
...
{
  "hello": "world"
}
```

14.13 Authorization

Chalice supports multiple mechanisms for authorization. This topic covers how you can integrate authorization into your Chalice applications.

In Chalice, all the authorizers are configured per-route and specified using the `authorizer` kwarg to an `@app.route()` call. You control which type of authorizer to use based on what's passed as the `authorizer` kwarg. You can use the same authorizer instance for multiple routes.

The first set of authorizers chalice supports cover the scenario where you have some existing authorization mechanism that you just want your Chalice app to use.

Chalice also supports built-in authorizers, which allows Chalice to manage your custom authorizers as part of `chalice deploy`. This is covered in the [Built-in Authorizers](#) section.

14.13.1 AWS IAM Authorizer

The IAM Authorizer allows you to control access to API Gateway with [IAM permissions](#)

To associate an IAM authorizer with a route in chalice, you use the `IAMAuthorizer` class:

```
from chalice import IAMAuthorizer

authorizer = IAMAuthorizer()

@app.route('/iam-auth', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"success": True}
```

See the [API Gateway documentation](#) for more information on controlling access to API Gateway with IAM permissions.

14.13.2 Amazon Cognito User Pools

In addition to using IAM roles and policies with the `IAMAuthorizer` you can also use a [Cognito user pools](#) to control who can access your Chalice app. A cognito user pool serves as your own identity provider to maintain a user directory.

To integrate Cognito user pools with Chalice, you'll need to have an existing cognito user pool configured.

```
from chalice import CognitoUserPoolAuthorizer

authorizer = CognitoUserPoolAuthorizer(
    'MyPool', provider_arns=['arn:aws:cognito:...:userpool/name'])
```

```
@app.route('/user-pools', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"success": True}
```

For more information about using Cognito user pools with API Gateway, see the [Use Amazon Cognito User Pools documentation](#).

14.13.3 Custom Authorizers

API Gateway also lets you write custom authorizers using a Lambda function. You can configure a Chalice route to use a pre-existing Lambda function as a custom authorizer. If you also want to write and manage your Lambda authorizer using Chalice, see the next section, Built-in Authorizers.

To connect an existing Lambda function as a custom authorizer in chalice, you use the `CustomAuthorizer` class:

```
from chalice import CustomAuthorizer

authorizer = CustomAuthorizer(
    'MyCustomAuth', header='Authorization',
    authorizer_uri=('arn:aws:apigateway:region:lambda:path/2015-03-31'
                   '/functions/arn:aws:lambda:region:account-id:'
                   'function:FunctionName/invocations'))

@app.route('/custom-auth', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"success": True}
```

14.13.4 Built-in Authorizers

The `IAMAuthorizer`, `CognitoUserPoolAuthorizer`, and the `CustomAuthorizer` classes are all for cases where you have existing resources for managing authorization and you want to wire them together with your Chalice app. A Built-in authorizer is used when you'd like to write your custom authorizer in Chalice, and have the additional Lambda functions managed when you run `chalice deploy/delete`. This section will cover how to use the built-in authorizers in chalice.

Creating an authorizer in chalice requires you use the `@app.authorizer` decorator to a function. The function must accept a single arg, which will be an instance of `AuthRequest`. The function must return a `AuthResponse`. As an example, we'll port the example from the [API Gateway documentation](#). First, we'll show the code and then walk through it:

```
from chalice import Chalice, AuthResponse

app = Chalice(app_name='demoauth1')

@app.authorizer()
def demo_auth(auth_request):
    token = auth_request.token
    # This is just for demo purposes as shown in the API Gateway docs.
    # Normally you'd call an oauth provider, validate the
    # jwt token, etc.
    # In this example, the token is treated as the status for demo
    # purposes.
    if token == 'allow':
        return AuthResponse(routes=['/'], principal_id='user')
```

```

else:
    # By specifying an empty list of routes,
    # we're saying this user is not authorized
    # for any URLs, which will result in an
    # Unauthorized response.
    return AuthResponse(routes=[], principal_id='user')

@app.route('/', authorizer=demo_auth)
def index():
    return {'context': app.current_request.context}

```

In the example above we define a built-in authorizer by decorating the `demo_auth` function with the `@app.authorizer()` decorator. Note you must use `@app.authorizer()` and not `@app.authorizer`. A built-in authorizer function has this type signature:

```
def auth_handler(auth_request: AuthRequest) -> AuthResponse: ...
```

Within the auth handler you must determine if the request is authorized or not. The `AuthResponse` contains the allowed URLs as well as the principal id of the user. You can optionally return a dictionary of key value pairs (as the `context` kwarg). This dictionary will be passed through on subsequent requests. In our example above we're not using the context dictionary. API Gateway will convert all the values in the `context` dictionary to string values.

Now let's deploy our app. As usual, we just need to run `chalice deploy` and `chalice` will automatically deploy all the necessary Lambda functions for us.

Now when we try to make a request, we'll get an Unauthorized error:

```

$ http https://api.us-west-2.amazonaws.com/api/
HTTP/1.1 401 Unauthorized

{
  "message": "Unauthorized"
}

```

If we add the appropriate authorization header, we'll see the call succeed:

```

$ http https://api.us-west-2.amazonaws.com/api/ 'Authorization: allow'
HTTP/1.1 200 OK

{
  "context": {
    "accountId": "12345",
    "apiId": "api",
    "authorizer": {
      "principalId": "user"
    },
    "httpMethod": "GET",
    "identity": {
      "accessKey": null,
      "accountId": null,
      "apiKey": "",
      "caller": null,
      "cognitoAuthenticationProvider": null,
      "cognitoAuthenticationType": null,
      "cognitoIdentityId": null,
      "cognitoIdentityPoolId": null,
      "sourceIp": "1.1.1.1",
      "user": null,

```

```
        "userAgent": "HTTPie/0.9.9",
        "userArn": null
    },
    "path": "/api/",
    "requestId": "d35d2063-56be-11e7-9ce1-dd61c24a3668",
    "resourceId": "id",
    "resourcePath": "/",
    "stage": "dev"
}
}
```

The low level API for API Gateway’s custom authorizer feature requires that an IAM policy must be returned. The `AuthResponse` class we’re using is a wrapper over building the IAM policy yourself. If you want low level control and would prefer to construct the IAM policy yourself you can return a dictionary of the IAM policy instead of an instance of `AuthResponse`. If you do that, the dictionary is returned without modification back to API Gateway.

For more information on custom authorizers, see the [Use API Gateway Custom Authorizers](#) page in the API Gateway user guide.

14.14 Lambda Event Sources

14.14.1 Scheduled Events

Chalice has support for *scheduled events*. This feature allows you to periodically invoke a lambda function based on some regular schedule. You can specify a fixed rate or a cron expression.

To create a scheduled event in chalice, you use the `@app.schedule()` decorator. Let’s look at an example.

```
app = chalice.Chalice(app_name='foo')

@app.schedule('rate(1 hour)')
def every_hour(event):
    print(event.to_dict())
```

In this example, we have a single lambda function that we want automatically invoked every hour. When you run `chalice deploy` Chalice will create a lambda function as well as the necessary CloudWatch events/rules such that the `every_hour` function is invoked every hour.

The `Chalice.schedule()` method accepts either a string or an instance of `Rate` or `Cron`. For example:

```
app = chalice.Chalice(app_name='foo')

@app.schedule(Rate(1, unit=Rate.HOURS))
def every_hour(event):
    print(event.to_dict())
```

The function you decorate must accept a single argument, which will be of type `CloudWatchEvent`.

You can use the `schedule()` decorator multiple times in your chalice app. Each `schedule()` decorator will result in a new lambda function and associated CloudWatch event rule. For example:

```
app = chalice.Chalice(app_name='foo')

@app.schedule(Rate(1, unit=Rate.HOURS))
def every_hour(event):
    print(event.to_dict())
```



```
@app.schedule(Rate(2, unit=Rate.HOURS))
def every_two_hours(event):
    print(event.to_dict())
```

In the app above, chalice will create two lambda functions, and configure `every_hour` to be invoked once an hour, and `every_two_hours` to be invoked once every two hours.

14.14.2 CloudWatch Events

You can configure a lambda function to subscribe to any *CloudWatch Event* <<https://amzn.to/2SCgWA6>>.

To subscribe to a CloudWatch Event in chalice, you use the `@app.on_cw_event()` decorator. Let's look at an example.

```
app = chalice.Chalice(app_name='foo')

@app.on_cw_event({"source": ["aws.codecommit"]})
def on_code_commit_changes(event):
    print(event.to_dict())
```

In this example, we have a single lambda function that we subscribe to all events from the AWS Code Commit service. The first parameter to the decorator is the event pattern that will be used to filter the events sent to the function.

See the CloudWatch Event pattern docs for additional syntax and examples <https://amzn.to/2OlqZso>

The function you decorate must accept a single argument, which will be of type *CloudWatchEvent*.

14.14.3 S3 Events

You can configure a lambda function to be invoked whenever certain events happen in an S3 bucket. This uses the *event notifications* feature provided by Amazon S3.

To configure this, you just tell Chalice the name of an existing S3 bucket, along with what events should trigger the lambda function. This is done with the *Chalice.on_s3_event()* decorator.

Here's an example:

```
from chalice import Chalice

app = chalice.Chalice(app_name='s3eventdemo')
app.debug = True

@app.on_s3_event(bucket='mybucket-name',
                 events=['s3:ObjectCreated:*'])
def handle_s3_event(event):
    app.log.debug("Received event for bucket: %s, key: %s",
                  event.bucket, event.key)
```

In this example above, Chalice connects the S3 bucket to the `handle_s3_event` Lambda function such that when an object is uploaded to the `mybucket-name` bucket, the Lambda function will be invoked. This example also uses the `.bucket` and `.key` attribute from the `event` parameter, which is of type *S3Event*.

It will automatically create the appropriate S3 notification configuration as needed. Chalice will also leave any existing notification configuration on the `mybucket-name` untouched. It will only merge in the additional configuration needed for the `handle_s3_event` Lambda function.

Warning: This feature only works when using *chalice deploy*. Because you configure the lambda function with the name of an existing S3 bucket, it is not possible to describe this using a CloudFormation/SAM template. The *chalice* package command will fail. You will eventually be able to request that chalice create a bucket for you, which will support the *chalice* package command.

The function you decorate must accept a single argument, which will be of type *S3Event*.

14.14.4 SNS Events

You can configure a lambda function to be automatically invoked whenever something publishes to an SNS topic. Chalice will automatically handle creating the lambda function, subscribing the lambda function to the SNS topic, and modifying the lambda function policy to allow SNS to invoke the function.

To configure this, you just need the name of an existing SNS topic you'd like to subscribe to. The SNS topic must already exist.

Below is an example of how to set this up. The example uses *boto3* to create the SNS topic. If you don't have *boto3* installed in your virtual environment, be sure to install it with:

```
$ pip install boto3
```

First, we'll create an SNS topic using *boto3*.

```
$ python
>>> import boto3
>>> sns = boto3.client('sns')
>>> sns.create_topic(Name='my-demo-topic')
{'TopicArn': 'arn:aws:sns:us-west-2:12345:my-demo-topic',
 'ResponseMetadata': {}}
```

Next, we'll create our chalice app:

```
$ chalice new-project chalice-demo-sns
$ cd chalice-demo-sns/
```

We'll update the *app.py* file to use the *on_sns_message* decorator:

```
from chalice import Chalice

app = chalice.Chalice(app_name='chalice-sns-demo')
app.debug = True

@app.on_sns_message(topic='my-demo-topic')
def handle_sns_message(event):
    app.log.debug("Received message with subject: %s, message: %s",
                  event.subject, event.message)
```

We can now deploy our chalice app:

```
$ chalice deploy
Creating deployment package.
Creating IAM role: chalice-demo-sns-dev
Creating lambda function: chalice-demo-sns-dev-handle_sns_message
Subscribing chalice-demo-sns-dev-handle_sns_message to SNS topic my-demo-topic
Resources deployed:
- Lambda ARN: arn:aws:lambda:us-west-2:123:function:...
```

And now we can test our app by publishing a few SNS messages to our topic. We'll do this using `boto3`. In the example below, we're using `list_topics()` to find the ARN associated with our topic name before calling the `publish()` method.

```
$ python
>>> import boto3
>>> sns = boto3.client('sns')
>>> topic_arn = [t['TopicArn'] for t in sns.list_topics()['Topics']
...               if t['TopicArn'].endswith(':my-demo-topic')][0]
>>> sns.publish(Message='TestMessage1', Subject='TestSubject1',
...             TopicArn=topic_arn)
{'MessageId': '12345', 'ResponseMetadata': {}}
>>> sns.publish(Message='TestMessage2', Subject='TestSubject2',
...             TopicArn=topic_arn)
{'MessageId': '54321', 'ResponseMetadata': {}}
```

To verify our function was called correctly, we can use the `chalice logs` command:

```
$ chalice logs -n handle_sns_message
2018-06-28 17:49:30.513000 547e0f chalice-demo-sns - DEBUG - Received message with subject: TestSubject
2018-06-28 17:49:40.391000 547e0f chalice-demo-sns - DEBUG - Received message with subject: TestSubject
```

In this example we used the SNS topic name to register our handler, but you can also use the topic arn. This can be useful if your topic is in another region or account.

14.14.5 SQS Events

You can configure a lambda function to be invoked whenever messages are available on an SQS queue. To configure this, use the `Chalice.on_sqs_message()` decorator and provide the name of the SQS queue and an optional batch size.

The message visibility timeout of your SQS queue must be greater than or equal to the lambda timeout. The default message visibility timeout when you create an SQS queue is 30 seconds, and the default timeout for a Lambda function is 60 seconds, so you'll need to modify one of these values in order to successfully connect an SQS queue to a Lambda function.

You can check the visibility timeout of your queue using the `GetQueueAttributes` API call. Using the [AWS CLI](#), you can run this command to check the value:

```
$ aws sqs get-queue-attributes \
  --queue-url https://us-west-2.queue.amazonaws.com/1/testq \
  --attribute-names VisibilityTimeout
{
  "Attributes": {
    "VisibilityTimeout": "30"
  }
}
```

You can set the visibility timeout of your SQS queue using the `SetQueueAttributes` API call. Again using the [AWS CLI](#) you can run this command:

```
$ aws sqs set-queue-attributes \
  --queue-url https://us-west-2.queue.amazonaws.com/1/testq \
  --attributes VisibilityTimeout=60
```

If you would prefer to change the timeout of your lambda function instead, you can specify this timeout value using the `lambda_timeout` config key if your `.chalice/config.json` file. See [Lambda Specific Configuration](#) for

a list of all supported lambda configuration values in chalice. In this example below, we're setting the timeout of our `handle_sqs_message` lambda function to 30 seconds:

```
$ cat .chalice/config.json
{
  "stages": {
    "dev": {
      "lambda_functions": {
        "handle_sqs_message": {
          "lambda_timeout": 30
        }
      }
    }
  },
  "version": "2.0",
  "app_name": "chalice-sqs-demo"
}
```

In this example below, we're connecting the `handle_sqs_message` lambda function to the `my-queue` SQS queue. Note that we are specifying the queue name, not the queue URL or queue ARN. If you are connecting your lambda function to a FIFO queue, make sure you specify the `.fifo` suffix, e.g. `my-queue.fifo`.

```
from chalice import Chalice

app = chalice.Chalice(app_name='chalice-sqs-demo')
app.debug = True

@app.on_sqs_message(queue='my-queue', batch_size=1)
def handle_sqs_message(event):
    for record in event:
        app.log.debug("Received message with contents: %s", record.body)
```

Whenever a message is sent to the SQS queue our function will be automatically invoked. The function argument is an `SQSEvent` object, and each `record` in the example above is of type `SQSRecord`. Lambda takes care of automatically scaling your function as needed. See [Understanding Scaling Behavior](#) for more information on how Lambda scaling works.

If your lambda functions completes without raising an exception, then Lambda will automatically delete all the messages associated with the `SQSEvent`. You don't need to manually call `sqs.delete_message()` in your lambda function. If your lambda function raises an exception, then Lambda won't delete any messages, and once the visibility timeout has been reached, the messages will be available again in the SQS queue. Note that if you are using a batch size of more than one, the entire batch succeeds or fails. This means that it is possible for your lambda function to see a message multiple times, even if it's successfully processed the message previously. There are a few options available to mitigate this:

- Use a batch size of 1 (the default value).
- Use a separate data store to check if you've already processed an SQS message. You can use services such as Amazon DynamoDB or Amazon ElastiCache.
- Manually call `sqs.delete_message()` in your Lambda function once you've successfully processed a message.

For more information on Lambda and SQS, see the [AWS documentation](#).

14.15 Pure Lambda Functions

Chalice provides abstractions over AWS Lambda functions, including:

- An API handler the coordinates with API Gateway for creating rest APIs.
- A custom authorizer that allows you to integrate custom auth logic in your rest API.
- A scheduled event that includes managing the CloudWatch Event rules, targets, and permissions.

However, chalice also supports managing pure Lambda functions that don't have any abstractions built on top. This is useful if you want to create a Lambda function for something that's not supported by chalice or if you just want to create Lambda functions but don't want to manage handling dependencies and deployments yourself.

In order to do this, you can use the `Chalice.lambda_function()` decorator to denote that this python function is a pure lambda function that should be invoked as is, without any input or output mapping. When you use this function, you must provide a function that maps to the same function signature expected by AWS Lambda as [defined here](#).

Let's look at an example.

```
app = chalice.Chalice(app_name='foo')

@app.route('/')
def index():
    return {'hello': 'world'}

@app.lambda_function()
def custom_lambda_function(event, context):
    # Anything you want here.
    return {}

@app.lambda_function(name='MyFunction')
def other_lambda_function(event, context):
    # Anything you want here.
    return {}
```

In this example, we've updated the starter hello world app with two extra Lambda functions. When you run `chalice deploy` Chalice will create three Lambda functions. The first lambda function is for the API handler used by API gateway. The second and third lambda function will be pure lambda functions. These two additional lambda functions won't be hooked up to anything. You'll need to manage connecting them to any additional AWS Resources on your own.

14.16 Blueprints

Chalice blueprints are used to organize your application into logical components. Using a blueprint, you define your resources and decorators in modules outside of your `app.py`. You then register a blueprint in your main `app.py` file. Blueprints support any decorator available on an application object.

Note: The Chalice blueprints are conceptually similar to [Blueprints](#) in Flask. Flask blueprints allow you to define a set of URL routes separately from the main `Flask` object. This concept is extended to all resources in Chalice. A Chalice blueprint can have Lambda functions, event handlers, built-in authorizers, etc. in addition to a collection of routes.

14.16.1 Example

In this example, we'll create a blueprint with part of our routes defined in a separate file. First, let's create an application:

```
$ chalice new-project blueprint-demo
$ cd blueprint-demo
$ mkdir chalicelib
$ touch chalicelib/__init__.py
$ touch chalicelib/blueprints.py
```

Next, we'll open the `chalicelib/blueprints.py` file:

```
from chalice import Blueprint

extra_routes = Blueprint(__name__)

@extra_routes.route('/foo')
def foo():
    return {'foo': 'bar'}
```

The `__name__` is used to denote the import path of the blueprint. This name must match the import name of the module so the function can be properly imported when running in Lambda. We'll now import this module in our `app.py` and register this blueprint. We'll also add a route in our `app.py` directly:

```
from chalice import Chalice
from chalicelib.blueprints import extra_routes

app = Chalice(app_name='blueprint-demo')
app.register_blueprint(extra_routes)

@app.route('/')
def index():
    return {'hello': 'world'}
```

At this point, we've defined two routes. One route, `/`, is directly defined in our `app.py` file. The other route, `/foo` is defined in `chalicelib/blueprints.py`. It was added to our Chalice app when we registered it via `app.register_blueprint(extra_routes)`.

We can deploy our application to verify this works as expected:

```
$ chalice deploy
Creating deployment package.
Creating IAM role: blueprint-demo-dev
Creating lambda function: blueprint-demo-dev
Creating Rest API
Resources deployed:
- Lambda ARN: arn:aws:lambda:us-west-2:1234:function:blueprint-demo-dev
- Rest API URL: https://rest-api.execute-api.us-west-2.amazonaws.com/api/
```

We should now be able to request the `/` and `/foo` routes:

```
$ http https://rest-api.execute-api.us-west-2.amazonaws.com/api/
HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 17
Content-Type: application/json
Date: Sat, 22 Dec 2018 01:05:48 GMT
Via: 1.1 5ab5dc09da67e3ea794ec8a82992cc89.cloudfront.net (CloudFront)
X-Amz-Cf-Id: Cdsow9--fnTH5EdjkjWBMWINCCMD4nGmi4S_3iMYMK0rpc8Mpiymgw==
X-Amzn-Trace-Id: Root=1-5c1d8dec-f1ef3ee83c7c654ca7fb3a70;Sampled=0
```

```

X-Cache: Miss from cloudfront
x-amz-apigw-id: SSMc6H_yvHcFcEw=
x-amzn-RequestId: b7bd0c87-0585-11e9-90cf-59b71c1a1de1

{
  "hello": "world"
}

$ http https://rest-api.execute-api.us-west-2.amazonaws.com/api/foo
HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 13
Content-Type: application/json
Date: Sat, 22 Dec 2018 01:05:51 GMT
Via: 1.1 95b0ac620fa3a80ee590ecf1cdalc698.cloudfront.net (CloudFront)
X-Amz-Cf-Id: HX4l1BNdWvYDRXan17PFZyalvaomoJel4rP7d8_stdw2qT50v7Iybg==
X-Amzn-Trace-Id: Root=1-5c1d8def-214e7f681ff82c00fd81f37a;Sampled=0
X-Cache: Miss from cloudfront
x-amz-apigw-id: SSMdXF40vHcF-mg=
x-amzn-RequestId: b96f77bf-0585-11e9-b229-01305cd40040

{
  "foo": "bar"
}

```

14.16.2 Blueprint Registration

The `app.register_blueprint` function accepts two optional arguments, `name_prefix` and `url_prefix`. This allows you to register the resources in your blueprint at a certain url and name prefix. If you specify `url_prefix`, any routes defined in your blueprint will have the `url_prefix` prepended to it. If you specify the `name_prefix`, any Lambda functions created will have the `name_prefix` prepended to the resource name.

Note: The `name_prefix` parameter does not apply to the Lambda function associated with API Gateway, which is anything decorated with `@app.route()`.

14.16.3 Advanced Example

Let's create a more advanced example. If this application, let's say we want to organize our application into separate modules for our API and our event sources. We can create an app with these files:

```

$ ls -la chaliceelib/
__init__.py
api.py
events.py

```

The contents of `api.py` are:

```

from chalice import Blueprint

myapi = Blueprint(__name__)

@myapi.route('/')

```

```
def index():
    return {'hello': 'world'}

@myapi.route('/foo')
def index():
    return {'foo': 'bar'}
```

The contents of `events.py` are:

```
from chalice import Blueprint

myevents = Blueprint(__name__)

@myevents.schedule('rate(5 minutes)')
def cron(event):
    pass

@myevents.on_sns_message('MyTopic')
def handle_sns_message(event):
    pass
```

In our `app.py` we'll register these blueprints:

```
from chalice import Chalice
from chalicelib.events import myevents
from chalicelib.api import myapi

app = Chalice(app_name='blueprint-demo')
app.register_blueprint(myevents)
app.register_blueprint(myapi)
```

Now our `app.py` only registers the necessary blueprints, and all our resources are defined in blueprints.

14.17 Websockets

Warning: Websockets are considered an experimental API. You'll need to opt-in to this feature using the WEBSOCKETS feature flag:

```
app = Chalice('myapp')
app.experimental_feature_flags.extend([
    'WEBSOCKETS'
])
```

See *Experimental APIs* for more information.

Chalice supports websockets through integration with an API Gateway Websocket API. If any of the decorators are present in a Chalice app, then an API Gateway Websocket API will be deployed and wired to Lambda Functions.

14.17.1 Responding to websocket events

In a Chalice app the websocket API is accessed through the three decorators `on_ws_connect`, `on_ws_message`, `on_ws_disconnect`. These handle a new websocket connection, an incoming message on an existing connection, and a connection being cleaned up respectively.

A decorated websocket handler function takes one argument `event` with the type `WebsocketEvent`. This class allows easy access to information about the API Gateway Websocket API, and information about the particular socket the handler is being invoked to serve.

Below is a simple working example application that prints to CloudWatch Logs for each of the events.

```
from boto3.session import Session
from chalice import Chalice

app = Chalice(app_name='test-websockets')
app.experimental_feature_flags.update([
    'WEBSOCKETS',
])
app.websocket_api.session = Session()

@app.on_ws_connect()
def connect(event):
    print('New connection: %s' % event.connection_id)

@app.on_ws_message()
def message(event):
    print('%s: %s' % (event.connection_id, event.body))

@app.on_ws_disconnect()
def disconnect(event):
    print('%s disconnected' % event.connection_id)
```

14.17.2 Sending a message over a websocket

To send a message to a websocket client Chalice, use the `app.websocket_api.send()` method. This method will work in any of the decorated functions outlined in the above section.

Two pieces of information are needed to send a message. The identifier of the websocket, and the contents for the message. Below is a simple example that when it receives a message, it sends back the message "I got your message!" over the same socket.

```
from boto3.session import Session
from chalice import Chalice

app = Chalice(app_name='test-websockets')
app.experimental_feature_flags.update([
    'WEBSOCKETS',
])
app.websocket_api.session = Session()

@app.on_ws_message()
def message(event):
    app.websocket_api.send(event.connection_id, 'I got your message!')
```

See [Websocket Tutorials](#) for completely worked example applications.

14.18 Continuous Deployment (CD)

Chalice can be used to set up a basic Continuous Deployment pipeline. The `chalice deploy` command is good for getting up and running quickly with Chalice, but in a team environment properly managing permissions and sharing and updating the `deployed.json` file will get messy.

One way to scale up your chalice app is to create a continuous deployment pipeline. The pipeline can run tests on code changes and, if they pass, promote the new build to a testing stage. More checks can be put in place to manually promote a build to production, or you can do so automatically. This model greatly simplifies managing what resources belong to your Chalice app as they are all stored in the Continuous Deployment pipeline.

Chalice can generate a CloudFormation template that will create a starter CD pipeline. It contains a CodeCommit repo, a CodeBuild stage for packaging your chalice app, and a CodePipeline stage to deploy your application using CloudFormation.

14.18.1 Usage example

Setting up the deployment pipeline is a two step process. First use the `chalice generate-pipeline` command to generate a base CloudFormation template. Second use the AWS CLI to deploy the CloudFormation template using the `aws cloudformation deploy` command. Below is an example.

```
$ chalice generate-pipeline pipeline.json
$ aws cloudformation deploy --stack-name mystack
  --template-file pipeline.json --capabilities CAPABILITY_IAM
Waiting for changeset to be created..
Waiting for stack create/update to complete
Successfully created/updated stack - mystack
```

Once the CloudFormation template has finished creating the stack, you will have several new AWS resources that make up a bare bones CD pipeline.

- **CodeCommit Repository** - The [CodeCommit](#) repository is the entrypoint into the pipeline. Any code you want to deploy should be pushed to this remote.
- **CodePipeline Pipeline** - The [CodePipeline](#) is what coordinates the build process, and pushes the released code out.
- **CodeBuild Project** - The [CodeBuild](#) project is where the code bundle is built that will be pushed to Lambda. The default CloudFormation template will create a CodeBuild stage that builds a package using `chalice package` and then uploads those artifacts for CodePipeline to deploy.
- **S3 Buckets** - Two S3 buckets are created on your behalf.
 - **artifactbucketstore** - This bucket stores artifacts that are built by the CodeBuild project. The only artifact by default is the `transformed.yaml` created by the `aws cloudformation package` command.
 - **applicationbucket** - Stores the application bundle after the Chalice application has been packaged in the CodeBuild stage.
- Each resource is created with all the required IAM roles and policies.

14.18.2 CodeCommit repository

The CodeCommit repository can be added as a git remote for deployment. This makes it easy to kick off deployments. The developer doing the deployment only needs to push the release code up to the CodeCommit repository master branch. All the developer needs is keys that allow for push access to the CodeCommit repository. This is a lot easier than managing a set of `deployed.json` resources across a repository and manually doing `chalice deploy` whenever a change needs to be deployed.

The default CodeCommit repository that is created is empty, you will have to populate it with the Chalice application code. Permissions will also need to be set up, you can find the documentation on how to do that [here](#).

14.18.3 CodePipeline

CodePipeline is the main coordinator between all the other resources. It watches for changes on the CodeCommit repository, and triggers builds in the CodeBuild project. If the build succeeds then it will start a CloudFormation deployment of the built artifacts to a beta stage. This should be treated as a starting point, not a fully featured CD system.

14.18.4 CodeBuild build script

By default Chalice will create the CodeBuild project with a default builds spec that does the following.

```
version: 0.1
phases:
  install:
    commands:
      - sudo pip install --upgrade awscli
      - aws --version
      - sudo pip install chalice
      - sudo pip install -r requirements.txt
      - chalice package /tmp/packaged
      - aws cloudformation package --template-file
        tmp/packaged/sam.json --s3-bucket ${APP_S3_BUCKET}
        --output-template-file transformed.yaml
artifacts:
  type: zip
  files:
    - transformed.yaml
```

The CodeBuild stage installs both the AWS CLI and Chalice, then creates a package out of your chalice project, pushing the package to the application S3 bucket that was created for you. The transformed CloudFormation template is the only artifact, and can be run by CodePipeline after the build has succeeded.

14.18.5 Deploying to beta stage

Once the CodeBuild stage has finished building the Chalice package and creating the `transformed.yaml`, CodePipeline will take these artifacts and use them to create or update the beta stage. The `transformed.yaml` is a CloudFormation template that CodePipeline will execute, all the code it references has been uploaded to the application bucket by the AWS CLI in the CodeBuild stage, so this is the only artifact we need.

Once the CodePipeline beta build stage is finished, the beta version of the app is deployed and ready for testing.

14.18.6 Extending

It is recommended to use this pipeline as a starting point. The default template does not run any tests on the Chalice app before deploying to beta. There is also no mechanism provided by Chalice for a production stage. Ideally the CodeBuild stage would be used to run unit and functional tests before deploying to beta. After the beta stage is up, integration tests can be run against that endpoint, and if they all pass the beta stage could be promoted to a production stage using the CodePipeline manual approval feature.

14.19 Experimental APIs

Chalice maintains backwards compatibility for all features that appear in this documentation. Any Chalice application using version 1.x will continue to work for all future versions of 1.x.

We also believe that Chalice has a lot of potential for new ideas and APIs, many of which will take several iterations to get right. We may implement a new idea and need to make changes based on customer usage and feedback. This may include backwards incompatible changes all the way up to the removal of a feature.

To accommodate these new features, Chalice has support for experimental APIs, which are features that are added to Chalice on a provisional basis. Because these features may include backwards incompatible changes, you must explicitly opt-in to using these features. This makes it clear that you are using an experimental feature that may change.

14.19.1 Opting-in to Experimental APIs

Each experimental feature in chalice has a name associated with it. To opt-in to an experimental API, you must have the feature name to the `experimental_feature_flags` attribute on your app object. This attribute's type is a set of strings.

```
from chalice import Chalice

app = Chalice('myapp')
app.experimental_feature_flags.update([
    'MYFEATURE1',
    'MYFEATURE2',
])
```

If you use an experimental API without opting-in, you will receive a message whenever you run a Chalice CLI command. The error message tells you which feature flags you need to add:

```
$ chalice deploy
You are using experimental features without explicitly opting in.
Experimental features do not guarantee backwards compatibility and may be removed in the future.
If you still like to use these experimental features, you can opt-in by adding this to your app.py f

app.experimental_feature_flags.update([
    'FEATURE_FLAG_NAME'
])

See https://chalice.readthedocs.io/en/latest/topics/experimental.rst for more details.
```

The feature flag only happens when running CLI commands. There are no runtime checks for experimental features once your application is deployed.

14.19.2 List of Experimental APIs

In the table below, the “Feature Flag Name” column is the value you must add to the `app.experimental_feature_flags` attribute. The status of an experimental API can be:

- **Trial** - You must explicitly opt-in to use this feature.
- **Accepted** - This feature has graduated from an experimental feature to a fully supported, backwards compatible feature in Chalice. Accepted features still appear in the table for auditing purposes.
- **Rejected** - This feature has been removed.

Table 14.1: Experimental APIs

Feature	Feature Flag Name	Version Added	Version Finalized	Status	GitHub Issue(s)
<i>Blueprints</i>	BLUEPRINTS	1.7.0	1.15.0	Accepted	#1023 , #651
<i>Websockets</i>	WEBSOCKETS	1.10.0	n/a	Trial	#1041 , #1017

See the [original discussion](#) for more background information and alternative proposals.

API Reference

15.1 Chalice

class `Chalice` (*app_name*)

This class represents a chalice application. It provides:

- The ability to register routes using the `route()` method.
- Within a view function, the ability to introspect the current request using the `current_request` attribute which is an instance of the `Request` class.

current_request

An object of type `Request`. This value is only set when a view function is being called. This attribute can be used to introspect the current HTTP request.

api

An object of type `APIGateway`. This attribute can be used to control how apigateway interprets Content-Type headers in both requests and responses.

lambda_context

A Lambda context object that is passed to the invoked view by AWS Lambda. You can find out more about this object by reading the [lambda context object documentation](#).

Note: This is only set on `@app.route` handlers. For other handlers it will be `None`. Instead the event parameter will have a `context` property. For example `S3Event.context`.

debug

A boolean value that enables debugging. By default, this value is `False`. If debugging is true, then internal errors are returned back to the client. Additionally, debug log messages generated by the framework will show up in the cloudwatch logs. Example usage:

```
from chalice import Chalice

app = Chalice(app_name="appname")
app.debug = True
```

websocket_api

An object of type `WebsocketAPI`. This attribute can be used to send messages to websocket clients connected through API Gateway.

route (*path*, ***options*)

Register a view function for a particular URI path. This method is intended to be used as a decorator for a

view function. For example:

```
from chalice import Chalice

app = Chalice(app_name="appname")

@app.route('/resource/{value}', methods=['PUT'])
def viewfunction(value):
    pass
```

Parameters

- **path** (*str*) – The path to associate with the view function. The path should only contain `[a-zA-Z0-9._-]` chars and curly braces for parts of the URL you would like to capture. The path should not end in a trailing slash, otherwise a validation error will be raised during deployment.
- **methods** (*list*) – Optional parameter that indicates which HTTP methods this view function should accept. By default, only GET requests are supported. If you only wanted to support POST requests, you would specify `methods=['POST']`. If you support multiple HTTP methods in a single view function (`methods=['GET', 'POST']`), you can check the `app.current_request.method` attribute to see which HTTP method was used when making the request.
- **name** (*str*) – Optional parameter to specify the name of the view function. You generally do not need to set this value. The name of the view function is used as the default value for the view name.
- **authorizer** (*Authorizer*) – Specify an authorizer to use for this view. Can be an instance of *CognitoUserPoolAuthorizer*, *CustomAuthorizer* or *IAMAuthorizer*.
- **content_types** (*str*) – A list of content types to accept for this view. By default `application/json` is accepted. If this value is specified, then chalice will reject any incoming request that does not match the provided list of content types with a 415 Unsupported Media Type response.
- **api_key_required** (*boolean*) – Optional parameter to specify whether the method required a valid API key.
- **cors** – Specify if CORS is supported for this view. This can either be a boolean value, `None`, or an instance of *CORSConfig*. Setting this value is set to `True` gives similar behavior to enabling CORS in the AWS Console. This includes injecting the `Access-Control-Allow-Origin` header to have a value of `*` as well as adding an `OPTIONS` method to support preflighting requests. If you would like more control over how CORS is configured, you can provide an instance of *CORSConfig*.

authorizer (*name*, ***options*)
Register a built-in authorizer.

```
from chalice import Chalice, AuthResponse

app = Chalice(app_name="appname")

@app.authorizer(ttl_seconds=30)
def my_auth(auth_request):
    # Validate auth_request.token, and then:
    return AuthResponse(routes=['/'], principal_id='username')
```



```
@app.route('/', authorizer=my_auth)
def viewfunction(value):
    pass
```

Parameters

- **ttl_seconds** – The number of seconds to cache this response. Subsequent requests that require this authorizer will use a cached response if available. The default is 300 seconds.
- **execution_role** – An optional IAM role to specify when invoking the Lambda function associated with the built-in authorizer.

schedule (*expression*, *name=None*)

Register a scheduled event that's invoked on a regular schedule. This will create a lambda function associated with the decorated function. It will also schedule the lambda function to be invoked with a scheduled CloudWatch Event.

See *Scheduled Events* for more information.

```
@app.schedule('cron(15 10 ? * 6L 2002-2005)')
def cron_handler(event):
    pass

@app.schedule('rate(5 minutes)')
def rate_handler(event):
    pass

@app.schedule(Rate(5, unit=Rate.MINUTES))
def rate_obj_handler(event):
    pass

@app.schedule(Cron(15, 10, '?', '*', '6L', '2002-2005'))
def cron_obj_handler(event):
    pass
```

Parameters

- **expression** – The schedule expression to use for the CloudWatch event rule. This value can either be a string value or an instance of type `ScheduleExpression`, which is either a `Cron` or `Rate` object. If a string value is provided, it will be provided directly as the `ScheduleExpression` value in the `PutRule` API call.
- **name** – The name of the function to use. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

on_cw_event (*pattern*, *name=None*)

Create a lambda function and configure it to be invoked whenever an event that matches the given pattern flows through CloudWatch Events or Event Bridge.

Parameters

- **pattern** – The event pattern to use to filter subscribed events. See the CloudWatch Events docs for examples <https://amzn.to/2OlqZso>
- **name** – The name of the function to create. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

on_s3_event (*bucket, events=None, prefix=None, suffix=None, name=None*)

Create a lambda function and configure it to be automatically invoked whenever an event happens on an S3 bucket.

Warning: You can't use the `chalice` package command when using the `on_s3_event` decorator. This is because CFN does not support configuring an existing S3 bucket.

See [S3 Events](#) for more information.

This example shows how you could implement an image resizer that's triggered whenever an object is uploaded to the `images/` prefix of an S3 bucket (e.g. `s3://mybucket/images/house.jpg`).

```
@app.on_s3_event('mybucket', events=['s3:ObjectCreated:Put'],
                 prefix='images/', suffix='.jpg')
def resize_image(event):
    with tempfile.NamedTemporaryFile('w') as f:
        s3.download_file(event.bucket, event.key, f.name)
        resize_image(f.name)
        s3.upload_file(event.bucket, 'resized/%s' % event.key, f.name)
```

Parameters

- **bucket** – The name of the S3 bucket. This bucket must already exist.
- **events** – A list of strings indicating the events that should trigger the lambda function. See [Supported Event Types](#) for the full list of strings you can provide. If this option is not provided, a default of `['s3:ObjectCreated:*']` is used, which will configure the lambda function to be invoked whenever a new object is created in the S3 bucket.
- **prefix** – An optional key prefix. This specifies that the lambda function should only be invoked if the key starts with this prefix (e.g. `prefix='images/'`). Note that this value is not a glob (e.g. `images/*`), it is a literal string match for the start of the key.
- **suffix** – An optional key suffix. This specifies that the lambda function should only be invoked if the key name ends with this suffix (e.g. `suffix='.jpg'`). Note that this value is not a glob (e.g. `*.txt`), it is a literal string match for the end of the key.
- **name** – The name of the function to use. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

on_sns_message (*topic, name=None*)

Create a lambda function and configure it to be automatically invoked whenever an SNS message is published to the specified topic.

See [SNS Events](#) for more information.

This example prints the subject and the contents of the message whenever something publishes to the sns topic of `mytopic`. In this example, the input parameter is of type `SNSEvent`.

```
app.debug = True

@app.on_sns_message(topic='mytopic')
def handler(event):
    app.log.info("SNS subject: %s", event.subject)
    app.log.info("SNS message: %s", event.message)
```

Parameters

- **topic** – The name or ARN of the SNS topic you want to subscribe to.
- **name** – The name of the function to use. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

on_sqs_message (*queue, batch_size=1, name=None*)

Create a lambda function and configure it to be automatically invoked whenever a message is published to the specified SQS queue.

The lambda function must accept a single parameter which is of type *SQSEvent*.

If the decorated function returns without raising any exceptions then Lambda will automatically delete the SQS messages associated with the *SQSEvent*. You don't need to manually delete messages. If any exception is raised, Lambda won't delete any messages, and the messages will become available once the visibility timeout has been reached. Note that for batch sizes of more than one, either the entire batch succeeds and all the messages in the batch are deleted by Lambda, or the entire batch fails. The default batch size is 1. See the [Using AWS Lambda with Amazon SQS](#) for more information on how Lambda integrates with SQS.

See the [SQS Events](#) topic guide for more information on using SQS in Chalice.

```
app.debug = True

@app.on_sqs_message(queue='myqueue')
def handler(event):
    app.log.info("Event: %s", event.to_dict())
    for record in event:
        app.log.info("Message body: %s", record.body)
```

Parameters

- **queue** – The name of the SQS queue you want to subscribe to. This is the name of the queue, not the ARN or Queue URL.
- **batch_size** – The maximum number of messages to retrieve when polling for SQS messages. The event parameter can have multiple SQS messages associated with it. This is why the event parameter passed to the lambda function is iterable. The batch size controls how many messages can be in a single event.
- **name** – The name of the function to use. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

lambda_function (*name=None*)

Create a pure lambda function that's not connected to anything.

See [Pure Lambda Functions](#) for more information.

Parameters name – The name of the function to use. This name is combined with the chalice app name as well as the stage name to create the entire lambda function name. This parameter is optional. If it is not provided, the name of the python function will be used.

register_blueprint (*blueprint, name_prefix=None, url_prefix=None*)

Register a [Blueprint](#) to a Chalice app. See [Blueprints](#) for more information.

Parameters

- **blueprint** – The [Blueprint](#) to register to the app.

- **name_prefix** – An optional name prefix that's added to all the resources specified in the blueprint.
- **url_prefix** – An optional url prefix that's added to all the routes defined the Blueprint. This allows you to set the root mount point for all URLs in a Blueprint.

on_ws_connect (*event*)

Create a Websocket API connect event handler.

Parameters event – The *WebsocketEvent* received to indicate a new connection has been registered with API Gateway. The identifier of this connection is under the *WebsocketEvent.connection_id* attribute.

see *Websockets* for more information.

on_ws_message (*event*)

Create a Websocket API message event handler.

Parameters event – The *WebsocketEvent* received to indicate API Gateway received a message from a connected client. The identifier of the client that sent the message is under the *WebsocketEvent.connection_id* attribute. The content of the message is available in the *WebsocketEvent.body* attribute.

see *Websockets* for more information.

on_ws_disconnect (*event*)

Create a Websocket API disconnect event handler.

Parameters event – The *WebsocketEvent* received to indicate an existing connection has been disconnected from API Gateway. The identifier of this connection is under the *WebsocketEvent.connection_id* attribute.

see *Websockets* for more information.

15.2 Request

class Request

A class that represents the current request. This is mapped to the `app.current_request` object.

```
@app.route('/objects/{key}', methods=['GET', 'PUT'])
def myobject(key):
    request = app.current_request
    if request.method == 'PUT':
        # handle PUT request
        pass
    elif request.method == 'GET':
        # handle GET request
        pass
```

query_params

A dict of the query params for the request. This value is `None` if no query params were provided in the request.

headers

A dict of the request headers.

uri_params

A dict of the captured URI params. This value is `None` if no URI params were provided in the request.

method

The HTTP method as a string.

json_body

The parsed JSON body (`json.loads(raw_body)`). This value will only be non-None if the Content-Type header is `application/json`, which is the default content type value in chalice.

raw_body

The raw HTTP body as bytes. This is useful if you need to calculate a checksum of the HTTP body.

context

A dict of additional context information.

stage_vars

A dict of configuration for the API Gateway stage.

to_dict()

Convert the `Request` object to a dictionary. This is useful for debugging purposes. This dictionary is guaranteed to be JSON serializable so you can return this value from a chalice view.

15.3 Response

class Response (*body, headers=None, status_code=200*)

A class that represents the response for the view function. You can optionally return an instance of this class from a view function if you want complete control over the returned HTTP response.

```
from chalice import Chalice, Response

app = Chalice(app_name='custom-response')

@app.route('/')
def index():
    return Response(body='hello world!',
                    status_code=200,
                    headers={'Content-Type': 'text/plain'})
```

New in version 0.6.0.

body

The HTTP response body to send back. This value must be a string.

headers

An optional dictionary of HTTP headers to send back. This is a dictionary of header name to header value, e.g `{ 'Content-Type': 'text/plain' }`

status_code

The integer HTTP status code to send back in the HTTP response.

15.4 Authorization

Each of these classes below can be provided using the `authorizer` argument for an `@app.route(authorizer=...)` call:

```
authorizer = CognitoUserPoolAuthorizer(
    'MyPool', header='Authorization',
    provider_arns=['arn:aws:cognito:...:userpool/name'])

@app.route('/user-pools', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

class CognitoUserPoolAuthorizer (*name, provider_arns, header='Authorization'*)

New in version 0.8.1.

name

The name of the authorizer.

provider_arns

The Cognito User Pool arns to use.

header

The header where the auth token will be specified.

class IAMAuthorizer

New in version 0.8.3.

class CustomAuthorizer (*name, authorizer_uri, ttl_seconds, header='Authorization'*)

New in version 0.8.1.

name

The name of the authorizer.

authorizer_uri

The URI of the lambda function to use for the custom authorizer. This usually has the form `arn:aws:apigateway:{region}:lambda:path/2015-03-31/functions/{lambda_arn}/invocations`

ttl_seconds

The number of seconds to cache the returned policy from a custom authorizer.

header

The header where the auth token will be specified.

15.4.1 Built-in Authorizers

These classes are used when defining built-in authorizers in Chalice.

class AuthRequest (*auth_type, token, method_arn*)

An instance of this class is passed as the first argument to an authorizer defined via `@app.authorizer()`. You generally do not instantiate this class directly.

auth_type

The type of authentication

token

The authorization token. This is usually the value of the `Authorization` header.

method_arn

The ARN of the API gateway being authorized.

class AuthResponse (*routes, principal_id, context=None*)

routes

A list of authorized routes. Each element in the list can either be a string route such as `"/foo/bar"` or an

instance of `AuthRoute`. If you specify the URL as a string, then all supported HTTP methods will be authorized. If you want to specify which HTTP methods are allowed, you can use `AuthRoute`. If you want to specify that all routes and HTTP methods are supported you can use the wildcard value of `"*"`:

```
AuthResponse(routes=['*'], ...)
```

principal_id

The principal id of the user.

context

An optional dictionary of key value pairs. This dictionary will be accessible in the `app.current_request.context` in all subsequent authorized requests for this user.

class AuthRoute (*path, methods*)

This class be used in the `routes` attribute of a `AuthResponse` instance to get fine grained control over which HTTP methods are allowed for a given route.

path

The allowed route specified as a string

methods

A list of allowed HTTP methods.

15.5 APIGateway

class APIGateway

This class is used to control how API Gateway interprets Content-Type headers in both requests and responses.

There is a single instance of this class attached to each `Chalice` object under the `api` attribute.

cors

Global cors configuration. If a route-level cors configuration is not provided, or is `None` then this configuration will be used. By default it is set to `False`. This can either be `True`, `False`, or an instance of the `CORSConfig` class. This makes it easy to enable CORS for your entire application by setting `app.api.cors = True`.

New in version 1.12.1.

default_binary_types

The value of `default_binary_types` are the Content-Types that are considered binary by default. This value should not be changed, instead you should modify the `binary_types` list to change the behavior of a content type. Its value is: `application/octet-stream`, `application/x-tar`, `application/zip`, `audio/basic`, `audio/ogg`, `audio/mp4`, `audio/mpeg`, `audio/wav`, `audio/webm`, `image/png`, `image/jpg`, `image/jpeg`, `image/gif`, `video/ogg`, `video/mpeg`, `video/webm`.

binary_types

The value of `binary_types` controls how API Gateway interprets requests and responses as detailed below.

If an incoming request has a Content-Type header value that is present in the `binary_types` list it will be assumed that its body is a sequence of raw bytes. You can access these bytes by accessing the `app.current_request.raw_body` property.

If an outgoing response from Chalice has a header Content-Type that matches one of the `binary_types` its body must be a `bytes` type object. It is important to note that originating request must have the `Accept` header for the same type as the Content-Type on the response. Otherwise a 400 error will be returned.

This value can be modified to change what types API Gateway treats as binary. The easiest way to do this is to simply append new types to the list.

```
app.api.binary_types.append('application/my-binary-data')
```

Keep in mind that there can only be a total of 25 binary types at a time and Chalice by default has a list of 16 types. It is recommended if you are going to make extensive use of binary types to reset the list to the exact set of content types you will be using. This can easily be done by reassigning the whole list.

```
app.api.binary_types = [
    'application/octet-stream',
    'application/my-binary-data',
]
```

Implementation Note: API Gateway and Lambda communicate through a JSON event which is encoded using UTF-8. The raw bytes are temporarily encoded using base64 when being passed between API Gateway and Lambda. In the worst case this encoding can cause the binary body to be inflated up to 4/3 its original size. Lambda only accepts an event up to 6mb, which means even if your binary data was not quite at that limit, with the base64 encoding it may exceed that limit. This will manifest as a 502 Bad Gateway error.

15.6 WebsocketAPI

class WebsocketAPI

This class is used to send messages to websocket clients connected to an API Gateway Websocket API.

session

A boto3 Session that will be used to send websocket messages to clients. Any custom configuration can be set through a botocore session. This **must** be manually set before websocket features can be used.

```
import botocore
from boto3.session import Session
from chalice import Chalice

app = Chalice('example')
session = botocore.session.Session()
session.set_config_variable('retries', {'max_attempts': 0})
app.websocket_api.session = Session(botocore_session=session)
```

configure (domain_name, stage)

Configure prepares the *WebsocketAPI* to call the *send()* method. Without first calling this method calls to *send()* will fail with the message *WebsocketAPI needs to be configured before sending messages..* This is because a boto3 apigatewaymanagementapi client must be created from the *session* with a custom endpoint in order to properly communicate with our API Gateway WebsocketAPI. This method is called on your behalf before each of the websocket handlers: *on_ws_connect*, *on_ws_message*, *on_ws_disconnect*. This ensures that the *send()* method is available in each of those handlers.

send (connection_id, message)

requires boto3>=1.9.91

Method to send a message to a client. The *connection_id* is the unique identifier of the socket to send the message to. The message must be a utf-8 string.

If the socket is disconnected it raises a *WebsocketDisconnectedError* error.

close (connection_id)

requires boto3>=1.9.221

Method to close a WebSocket connection. The `connection_id` is the unique identifier of the socket to close.

If the socket is already disconnected it raises a `WebsocketDisconnectedError` error.

info (*connection_id*)
requires boto3>=1.9.221

Method to get info about a WebSocket. The `connection_id` is the unique identifier of the socket to get info about.

The following is an example of the format this method returns:

```
{
    'ConnectedAt': datetime(2015, 1, 1),
    'Identity': {
        'SourceIp': 'string',
        'UserAgent': 'string'
    },
    'LastActiveAt': datetime(2015, 1, 1)
}
```

If the socket is disconnected it raises a `WebsocketDisconnectedError` error.

class WebsocketDisconnectedError

An exception raised when a message is sent to a websocket that has disconnected.

connection_id
 The unique identifier of the websocket that was disconnected.

15.7 CORS

class CORSConfig (*allow_origin='*', allow_headers=None, expose_headers=None, max_age=None, allow_credentials=None*)

CORS configuration to attach to a route, or globally on `app.api.cors`.

```
from chalice import CORSConfig
cors_config = CORSConfig(
    allow_origin='https://foo.example.com',
    allow_headers=['X-Special-Header'],
    max_age=600,
    expose_headers=['X-Special-Header'],
    allow_credentials=True
)

@app.route('/custom_cors', methods=['GET'], cors=cors_config)
def supports_custom_cors():
    return {'cors': True}
```

New in version 0.8.1.

allow_origin

The value of the Access-Control-Allow-Origin to send in the response. Keep in mind that even though the Access-Control-Allow-Origin header can be set to a string that is a space separated list of origins, this behavior does not work on all clients that implement CORS. You should only supply a single origin to the `CORSConfig` object. If you need to supply multiple origins you will need to define a custom handler for it that accepts `OPTIONS` requests and matches the `Origin` header against a whitelist of origins. If the match is successful then return just their `Origin` back to them in the Access-Control-Allow-Origin header.

allow_headers

The list of additional allowed headers. This list is added to list of built in allowed headers: Content-Type, X-Amz-Date, Authorization, X-API-Key, X-Amz-Security-Token.

expose_headers

A list of values to return for the Access-Control-Expose-Headers:

max_age

The value for the Access-Control-Max-Age

allow_credentials

A boolean value that sets the value of Access-Control-Allow-Credentials.

15.8 Event Sources

New in version 1.0.0b1.

class Rate (*value, unit*)

An instance of this class can be used as the expression value in the `Chalice.schedule()` method:

```
@app.schedule(Rate(5, unit=Rate.MINUTES))
def handler(event):
    pass
```

Examples:

```
# Run every minute.
Rate(1, unit=Rate.MINUTES)

# Run every 2 hours.
Rate(2, unit=Rate.HOURS)
```

value

An integer value that presents the amount of time to wait between invocations of the scheduled event.

unit

The unit of the provided value attribute. This can be either `Rate.MINUTES`, `Rate.HOURS`, or `Rate.DAYS`.

MINUTES, HOURS, DAYS

These values should be used for the `unit` attribute.

class Cron (*minutes, hours, day_of_month, month, day_of_week, year*)

An instance of this class can be used as the expression value in the `Chalice.schedule()` method.

```
@app.schedule(Cron(15, 10, '?', '*', '6L', '2002-2005'))
def handler(event):
    pass
```

It provides more capabilities than the `Rate` class. There are a few limits:

- You can't specify `day_of_month` and `day_of_week` fields in the same Cron expression. If you specify a value in one of the fields, you must use a `?` in the other.
- Cron expressions that lead to rates faster than 1 minute are not supported.

For more information, see the [API docs page](#).

Examples:

```

# Run at 10:00am (UTC) every day.
Cron(0, 10, '*', '*', '?', '*')

# Run at 12:15pm (UTC) every day.
Cron(15, 12, '*', '*', '?', '*')

# Run at 06:00pm (UTC) every Monday through Friday.
Cron(0, 18, '?', '*', 'MON-FRI', '*')

# Run at 08:00am (UTC) every 1st day of the month.
Cron(0, 8, 1, '*', '?', '*')

# Run every 15 minutes.
Cron('0/15', '*', '*', '*', '?', '*')

# Run every 10 minutes Monday through Friday.
Cron('0/10', '*', '?', '*', 'MON-FRI', '*')

# Run every 5 minutes Monday through Friday between
# 08:00am and 5:55pm (UTC).
Cron('0/5', '8-17', '?', '*', 'MON-FRI', '*')

```

class CloudWatchEvent

This is the input argument for a scheduled or CloudWatch events.

```

@app.schedule('rate(1 hour)')
def every_hour(event: CloudWatchEvent):
    pass

```

In the code example above, the event argument is of type `CloudWatchEvent`, which will have the following attributes.

version

By default, this is set to 0 (zero) in all events.

account

The 12-digit number identifying an AWS account.

region

Identifies the AWS region where the event originated.

detail

For CloudWatch events this will be the event payload. For scheduled events, this will be an empty dictionary.

detail_type

For scheduled events, this value will be "Scheduled Event".

source

Identifies the service that sourced the event. All events sourced from within AWS will begin with "aws." Customer-generated events can have any value here as long as it doesn't begin with "aws." We recommend the use of java package-name style reverse domain-name strings.

For scheduled events, this will be `aws.events`.

time

The event timestamp, which can be specified by the service originating the event. If the event spans a time interval, the service might choose to report the start time, so this value can be noticeably before the time the event is actually received.

event_id

A unique value is generated for every event. This can be helpful in tracing events as they move through rules to targets, and are processed.

resources

This JSON array contains ARNs that identify resources that are involved in the event. Inclusion of these ARNs is at the discretion of the service.

For scheduled events, this will include the ARN of the CloudWatch rule that triggered this event.

context

A [Lambda context object](#) that is passed to the handler by AWS Lambda. This is useful if you need the AWS request ID for tracing, or any other data in the context object.

to_dict()

Return the original event dictionary provided from Lambda. This is useful if you need direct access to the lambda event, for example if a new key is added to the lambda event that has not been mapped as an attribute to the `CloudWatchEvent` object. Example:

```
{'account': '123457940291',
 'detail': {},
 'detail-type': 'Scheduled Event',
 'id': '12345678-b9f1-4667-9c5e-39f98e9a6113',
 'region': 'us-west-2',
 'resources': ['arn:aws:events:us-west-2:123457940291:rule/testevents-dev-every_minute'],
 'source': 'aws.events',
 'time': '2017-06-30T23:28:38Z',
 'version': '0'}
```

class S3Event

This is the input argument for an S3 event.

```
@app.on_s3_event(bucket='mybucket')
def event_handler(event: S3Event):
    app.log.info("Event received for bucket: %s, key %s",
                 event.bucket, event.key)
```

In the code example above, the `event` argument is of type `S3Event`, which will have the following attributes.

bucket

The S3 bucket associated with the event.

key

The S3 key name associated with the event. The original key name in the S3 event payload is URL encoded. However, this `key` attribute automatically URL decodes the key name for you. If you need access to the original URL encoded key name, you can access it through the `to_dict()` method.

context

A [Lambda context object](#) that is passed to the handler by AWS Lambda. This is useful if you need the AWS request ID for tracing, or any other data in the context object.

to_dict()

Return the original event dictionary provided from Lambda. This is useful if you need direct access to the lambda event, for example if a new key is added to the lambda event that has not been mapped as an attribute to the `S3Event` object. Note that this event is not modified in any way. This means that the key name of the S3 object is URL encoded, which is the way that S3 sends this value to Lambda.

class SNSEvent

This is the input argument for an SNS event handler.

```
@app.on_sns_message(topic='mytopic')
def event_handler(event: SNSEvent):
    app.log.info("Message received with subject: %s, message: %s",
                 event.subject, event.message)
```

In the code example above, the `event` argument is of type `SNSEvent`, which will have the following attributes.

subject

The subject of the SNS message that was published.

message

The string value of the SNS message that was published.

context

A [Lambda context object](#) that is passed to the handler by AWS Lambda. This is useful if you need the AWS request ID for tracing, or any other data in the context object.

to_dict()

Return the original event dictionary provided from Lambda. This is useful if you need direct access to the lambda event, for example if a new key is added to the lambda event that has not been mapped as an attribute to the `SNSEvent` object.

class SQSEvent

This is the input argument for an SQS event handler.

```
@app.on_sqs_message(queue='myqueue')
def event_handler(event: SQSEvent):
    app.log.info("Event: %s", event.to_dict())
```

In the code example above, the `event` argument is of type `SQSEvent`. An `SQSEvent` can have multiple sqs messages associated with it. To access the multiple messages, you can iterate over the `SQSEvent`.

__iter__()

Iterate over individual SQS messages associated with the event. Each element in the iterable is of type [SQSRecord](#).

context

A [Lambda context object](#) that is passed to the handler by AWS Lambda. This is useful if you need the AWS request ID for tracing, or any other data in the context object.

to_dict()

Return the original event dictionary provided from Lambda. This is useful if you need direct access to the lambda event, for example if a new key is added to the lambda event that has not been mapped as an attribute to the `SQSEvent` object.

class SQSRecord

Represents a single SQS record within an [SQSEvent](#).

body

The body of the SQS message.

receipt_handle

The receipt handle associated with the message. This is useful if you need to manually delete an SQS message to account for partial failures.

context

A [Lambda context object](#) that is passed to the handler by AWS Lambda.

`to_dict()`

Return the original dictionary associated with the given message. This is useful if you need direct access to the lambda event.

15.9 Blueprints

class `Blueprint` (*import_name*)

An object used for grouping related handlers together. This is primarily used as a mechanism for organizing your lambda handlers. Any decorator methods defined in the *Chalice* object are also defined on a *Blueprint* object. You can register a blueprint to a Chalice app using the *Chalice.register_blueprint()* method.

The *import_name* is the module in which the *Blueprint* is defined. It is used to construct the appropriate handler string when creating the Lambda functions associated with a *Blueprint*. This is typically the `__name__` attribute: `mybp = Blueprint(__name__)`.

See *Blueprints* for more information.

```
# In ./app.py

from chalice import Chalice
from chalicelib import myblueprint

app = Chalice(app_name='blueprints')
app.register_blueprint(myblueprint)

# In chalicelib/myblueprint.py

from chalice import Blueprint

myblueprint = Blueprint(__name__)

@myblueprint.route('/')
def index():
    return {'hello': 'world'}
```

15.10 Websockets

class `WebsocketEvent`

Event object event that is passed as the sole argument to any handler function decorated with one of the three websocket related handlers: `on_ws_connect`, `on_ws_disconnect`, `on_ws_message`.

domain_name

The domain name of the endpoint for the API Gateway Websocket API.

stage

The API Gateway stage of the Websocket API.

connection_id

A handle that uniquely identifies a connection with API Gateway.

body

The message body received. This is only populated on the `on_ws_message` otherwise it will be set to `None`.

json_body

The parsed JSON body (`json.loads(body)`) of the message. If the body is not JSON parsable then using this attribute will raise a `ValueError`.

See [Websockets](#) for more information.

16.1 Websocket Tutorials

16.1.1 Echo Server Example

An echo server is a simple server that echos any message it receives back to the client that sent it.

First install a copy of Chalice in a fresh environment, create a new project and cd into the directory:

```
$ pip install -U chalice
$ chalice new-project echo-server
$ cd echo-server
```

Our Chalice application will need boto3 as a dependency for both API Gateway to send websocket messages. Let's add a boto3 to the `requirements.txt` file:

```
$ echo "boto3>=1.9.91" > requirements.txt
```

Now that the requirement has been added. Let's install it locally since our next script will need it as well:

```
$ pip install -r requirements.txt
```

Next replace the contents of the `app.py` file with the code below.

Listing 16.1: `app.py`

```
1 from boto3.session import Session
2
3 from chalice import Chalice
4 from chalice import WebsocketDisconnectedError
5
6 app = Chalice(app_name="echo-server")
7 app.websocket_api.session = Session()
8 app.experimental_feature_flags.update([
9     'WEBSOCKETS'
10 ])
11
12
13 @app.on_ws_message()
14 def message(event):
15     try:
16         app.websocket_api.send(
17             connection_id=event.connection_id,
```

```

18         message=event.body,
19     )
20     except WebsocketDisconnectedError as e:
21         pass # Disconnected so we can't send the message back.

```

Stepping through this app line by line, the first thing to note is that we need to import and instantiate a boto3 session. This session is manually assigned to `app.websocket_api.session`. This is needed because in order to send websocket responses to API Gateway we need to construct a boto3 client. Chalice does not take a direct dependency on boto3 or botocore, so we need to provide the Session ourselves.

```

from boto3.session import Session
app.websocket_api.session = Session()

```

Next we enable the experimental feature `WEBSOCKETS`. Websockets are an experimental feature and are subject to API changes. This includes all aspects of the Websocket API exposed in Chalice. Including any public members of `app.websocket_api`, and the three decorators `on_ws_connect`, `on_ws_message`, and `on_ws_disconnect`.

```

app.experimental_feature_flags.update([
    'WEBSOCKETS'
])

```

To register a websocket handler, and cause Chalice to deploy an API Gateway Websocket API we use the `app.on_ws_message()` decorator. The event parameter here is a wrapper object with some convenience parameters attached. The most useful are `event.connection_id` and `event.body`. The `connection_id` is an API Gateway specific identifier that allows you to refer to the connection that sent the message. The `body` is the content of the message.

```

@app.on_ws_message()
def message(event):

```

Since this is an echo server, the message handler simply reads the content it received on the socket, and rewrites it back to the same socket. To send a message to a socket we call `app.websocket_api.send(connection_id, message)`. In this case, we just use the same `connection_id` we got the message from, and use the `body` we got from the event as the message to send.

```

app.websocket_api.send(
    connection_id=event.connection_id,
    message=event.body,
)

```

Finally, we catch the exception `WebsocketDisconnectedError` which is raised by `app.websocket_api.send` if the provided `connection_id` is not connected anymore. In our case this doesn't really matter since we don't have anything tracking our connections. The error has a `connection_id` property that contains the offending connection id.

```

except WebsocketDisconnectedError as e:
    pass # Disconnected so we can't send the message back.

```

Now that we understand the code, let's deploy it with `chalice deploy`:

```

$ chalice deploy
Creating deployment package.
Creating IAM role: echo-server-dev
Creating lambda function: echo-server-dev-websocket_message
Creating websocket api: echo-server-dev-websocket-api
Resources deployed:
  - Lambda ARN: arn:aws:lambda:region:0123456789:function:echo-server-dev-websocket_message
  - Websocket API URL: wss://{websocket_api_id}.execute-api.region.amazonaws.com/api/

```

To test out the echo server we will use the `websocket-client` package. You install it from PyPI:

```
$ pip install websocket-client
```

After deploying the Chalice app the output will contain a URL for connecting to the websocket API labeled: `- Websocket API URL:..` The `websocket-client` package installs a command line tool called `wsdump.py` which can be used to test websocket echo server:

```
$ wsdump.py wss://{websocket_api_id}.execute-api.region.amazonaws.com/api/
Press Ctrl+C to quit
> foo
< foo
> bar
< bar
> foo bar baz
< foo bar baz
>
```

Every message sent to the server (lines that start with `>`) result in a message sent to us (lines that start with `<`) with the same content.

If something goes wrong, you can check the chalice error logs using the following command:

```
$ chalice logs -n websocket_message
```

Note: If you encounter an Internal Server Error here it is likely that you forgot to include `boto3>=1.9.91` in the `requirements.txt` file.

To tear down the example. Just run:

```
$ chalice delete
Deleting Websocket API: {websocket_api_id}
Deleting function: arn:aws:lambda:us-west-2:0123456789:function:echo-server-dev-websocket_message
Deleting IAM role: echo-server-dev
```

16.1.2 Chat Server Example

Note: This example is for illustration purposes and does not represent best practices.

A simple chat server example application. This example will walk through deploying a chat application with separate chat rooms and nicknames. It uses a DynamoDB table to store state like connection IDs between websocket messages.

First install a copy of Chalice in a fresh environment, create a new project and `cd` into the directory:

```
$ pip install -U chalice
$ chalice new-project chalice-chat-example
$ cd chalice-chat-example
```

Our Chalice application will need `boto3` as a dependency for both DynamoDB access and in order to communicate back with API Gateway to send websocket messages. Let's add a `boto3` to the `requirements.txt` file:

```
$ echo "boto3>=1.9.91" > requirements.txt
```

Now that the requirement has been added. Let's install it locally since our next script will need it as well:

```
$ pip install -r requirements.txt
```

Unlike our previous example where we used `chalice deploy`, we will use `chalice package` to create a CloudFormation template. The AWS CLI will be used to deploy the template. To install the AWS CLI run the command:

```
$ pip install -U awscli
```

Starting in Chalice 1.10, the `package` command has a `--merge-template` argument that allows us to merge in a custom JSON file to the generated CloudFormation template. Since Chalice does not have any built-in support for DynamoDB currently, we will make a `resources.json` file with the DynamoDB definition. The template file will set the environment variable `TABLE` in all our Lambda functions as a CloudFormation reference to the DynamoDB table. Finally, the template will also override our IAM policy with a custom one to allow all the DynamoDB operations our application will need.

Below is the JSON file that contains all of our custom Cloudformation.

Listing 16.2: `resources.json`

```
{
  "Resources": {
    "ChaliceChatTable": {
      "Type" : "AWS::DynamoDB::Table",
      "Properties" : {
        "AttributeDefinitions" : [
          {
            "AttributeName": "PK",
            "AttributeType": "S"
          },
          {
            "AttributeName": "SK",
            "AttributeType": "S"
          }
        ],
        "KeySchema" : [
          {
            "AttributeName" : "PK",
            "KeyType" : "HASH"
          },
          {
            "AttributeName" : "SK",
            "KeyType" : "RANGE"
          }
        ],
        "GlobalSecondaryIndexes": [
          {
            "IndexName" : "ReverseLookup",
            "KeySchema" : [
              {
                "AttributeName" : "SK",
                "KeyType" : "HASH"
              },
              {
                "AttributeName" : "PK",
                "KeyType" : "RANGE"
              }
            ],
            "Projection" : {
              "ProjectionType": "ALL"
            }
          }
        ]
      }
    }
  }
}
```

```

    },
    "ProvisionedThroughput": {
        "ReadCapacityUnits": 1,
        "WriteCapacityUnits": 1
    }
},
"ProvisionedThroughput" : {
    "ReadCapacityUnits": 1,
    "WriteCapacityUnits": 1
},
"TableName": "ChaliceChat"
}
},
"WebsocketConnect": {
    "Properties": {
        "Environment": {
            "Variables": {
                "TABLE": {"Ref": "ChaliceChatTable"}
            }
        }
    }
},
"WebsocketMessage": {
    "Properties": {
        "Environment": {
            "Variables": {
                "TABLE": {"Ref": "ChaliceChatTable"}
            }
        }
    }
},
"WebsocketDisconnect": {
    "Properties": {
        "Environment": {
            "Variables": {
                "TABLE": {"Ref": "ChaliceChatTable"}
            }
        }
    }
},
"DefaultRole": {
    "Type": "AWS::IAM::Role",
    "Properties": {
        "AssumeRolePolicyDocument": {
            "Version": "2012-10-17",
            "Statement": [
                {
                    "Sid": "",
                    "Effect": "Allow",
                    "Principal": {
                        "Service": "lambda.amazonaws.com"
                    },
                    "Action": "sts:AssumeRole"
                }
            ]
        }
    }
},
"Policies": [

```

```

{
  "PolicyName": "DefaultRolePolicy",
  "PolicyDocument": {
    "Version": "2012-10-17",
    "Statement": [
      {
        "Effect": "Allow",
        "Action": [
          "logs:CreateLogGroup",
          "logs:CreateLogStream",
          "logs:PutLogEvents"
        ],
        "Resource": "arn:aws:logs:*:*:*"
      },
      {
        "Effect": "Allow",
        "Action": [
          "execute-api:ManageConnections"
        ],
        "Resource": "arn:aws:execute-api:*:*:*/*connections/*"
      },
      {
        "Effect": "Allow",
        "Action": [
          "dynamodb:DeleteItem",
          "dynamodb:PutItem",
          "dynamodb:GetItem",
          "dynamodb:UpdateItem",
          "dynamodb:Query",
          "dynamodb:Scan"
        ],
        "Resource": [
          {
            "Fn::Sub": "arn:aws:dynamodb:${AWS::Region}:${AWS::AccountId}:table/${ChaliceCl
          },
          {
            "Fn::Sub": "arn:aws:dynamodb:${AWS::Region}:${AWS::AccountId}:table/${ChaliceCl
          }
        ]
      }
    ]
  }
}

```

The current directory layout should now look like this:

```

$ tree -a .
.
-- .chalice
|   -- config.json
-- .gitignore
-- app.py
-- resources.json

```

```
-- requirements.txt

1 directory, 5 files
```

Next let's fill out the `app.py` file since it is pretty simple. Most of this example code is contained in the `chalicelib/` directory.

Listing 16.3: `chalice-chat-example/app.py`

```
from boto3.session import Session

from chalice import Chalice

from chalicelib import Storage
from chalicelib import Sender
from chalicelib import Handler

app = Chalice(app_name="chalice-chat-example")
app.websocket_api.session = Session()
app.experimental_feature_flags.update([
    'WEBSOCKETS'
])

STORAGE = Storage.from_env()
SENDER = Sender(app, STORAGE)
HANDLER = Handler(STORAGE, SENDER)

@app.on_ws_connect()
def connect(event):
    STORAGE.create_connection(event.connection_id)

@app.on_ws_disconnect()
def disconnect(event):
    STORAGE.delete_connection(event.connection_id)

@app.on_ws_message()
def message(event):
    HANDLER.handle(event.connection_id, event.body)
```

Similar to the previous example. We need to use `boto3` to construct a `Session` and pass it to `app.websocket_api.session`. We opt into the usage of the `WEBSOCKET` experimental feature. Most of the actual work is done in some classes that we import from `chalicelib/`. These classes are detailed below, and the various parts are explained in comments and doc strings. In addition to the previous example, we register a handler for `on_ws_connect` and `on_ws_disconnect` to handle events from API gateway when a new socket is trying to connect, or an existing socket is disconnected.

Finally before being able to deploy and test the app out, we need to fill out the `chalicelib` directory. This is the bulk of the app and it is explained inline in comments. Create a new directory called `chalicelib` and inside that directory create an `__init__.py` file and fill it out with the following file.

Listing 16.4: `chalice-chat-example/chalicelib/__init__.py`

```
import os
```

```

import boto3
from boto3.dynamodb.conditions import Key

from chalice import WebsocketDisconnectedError

class Storage(object):
    """An abstraction to interact with the DynamoDB Table."""
    def __init__(self, table):
        """Initialize Storage object

        :param table: A boto3 dynamodb Table resource object.
        """
        self._table = table

    @classmethod
    def from_env(cls):
        """Create table from the environment.

        The environment variable TABLE is present for a deployed application
        since it is set in all of the Lambda functions by a CloudFormation
        reference. We default to '', which will happen when we run
        ``chalice package`` since it loads the application, and no
        environment variable has been set. For local testing, a value should
        be manually set in the environment if '' will not suffice.
        """
        table_name = os.environ.get('TABLE', '')
        table = boto3.resource('dynamodb').Table(table_name)
        return cls(table)

    def create_connection(self, connection_id):
        """Create a new connection object in the database.

        When a new connection is created, we create a stub for
        it in the table. The stub uses a primary key of the
        connection_id and a sort key of username_. This translates
        to a connection with an unset username. The first message
        sent over the wire from the connection is to be used as the
        username, and this entry will be re-written.

        :param connection_id: The connection id to write to
            the table.
        """
        self._table.put_item(
            Item={
                'PK': connection_id,
                'SK': 'username_',
            },
        )

    def set_username(self, connection_id, old_name, username):
        """Set the username.

        The SK entry that goes with this connection id that starts
        with username_ is taken to be the username. The previous
        entry needs to be deleted, and a new entry needs to be
        written.

```



```

        :param connection_id: Connection id of the user trying to
            change their name.

        :param old_name: The original username. Since this is part of
            the key, it needs to be deleted and re-created rather than
            updated.

        :param username: The new username the user wants.
    """
    self._table.delete_item(
        Key={
            'PK': connection_id,
            'SK': 'username_%s' % old_name,
        },
    )
    self._table.put_item(
        Item={
            'PK': connection_id,
            'SK': 'username_%s' % username,
        },
    )

def list_rooms(self):
    """Get a list of all rooms that exist.

    Scan through the table looking for SKs that start with room_
    which indicates a room that a user is in. Collect a unique set
    of those and return them.
    """
    r = self._table.scan()
    rooms = set([item['SK'].split('_', 1)[1] for item in r['Items']
                  if item['SK'].startswith('room_')])
    return rooms

def set_room(self, connection_id, room):
    """Set the room a user is currently in.

    The room a user is in is in the form of an SK that starts with
    room_ prefix.

    :param connection_id: The connection id to move to a room.

    :param room: The room name to join.
    """
    self._table.put_item(
        Item={
            'PK': connection_id,
            'SK': 'room_%s' % room,
        },
    )

def remove_room(self, connection_id, room):
    """Remove a user from a room.

    The room a user is in is in the form of an SK that starts with
    room_ prefix. To leave a room we need to delete this entry.

    :param connection_id: The connection id to move to a room.

```

```

        :param room: The room name to join.
        """
        self._table.delete_item(
            Key={
                'PK': connection_id,
                'SK': 'room_%s' % room,
            },
        )

def get_connection_ids_by_room(self, room):
    """Find all connection ids that go to a room.

    This is needed whenever we broadcast to a room. We collect all
    their connection ids so we can send messages to them. We use a
    ReverseLookup table here which inverts the PK, SK relationship
    creating a partition called room_{room}. Everything in that
    partition is a connection in the room.

    :param room: Room name to get all connection ids from.
    """
    r = self._table.query(
        IndexName='ReverseLookup',
        KeyConditionExpression=(
            Key('SK').eq('room_%s' % room)
        ),
        Select='ALL_ATTRIBUTES',
    )
    return [item['PK'] for item in r['Items']]

def delete_connection(self, connection_id):
    """Delete a connection.

    Called when a connection is disconnected and all its entries need
    to be deleted.

    :param connection_id: The connection partition to delete from
        the table.
    """
    try:
        r = self._table.query(
            KeyConditionExpression=(
                Key('PK').eq(connection_id)
            ),
            Select='ALL_ATTRIBUTES',
        )
        for item in r['Items']:
            self._table.delete_item(
                Key={
                    'PK': connection_id,
                    'SK': item['SK'],
                },
            )
    except Exception as e:
        print(e)

def get_record_by_connection(self, connection_id):
    """Get all the properties associated with a connection.

```

Each connection_id creates a partition in the table with multiple SK entries. Each SK entry is in the format {property}_{value}. This method reads all those records from the database and puts them all into dictionary and returns it.

:param connection_id: The connection to get properties for.

"""

```
r = self._table.query(
    KeyConditionExpression=(
        Key('PK').eq(connection_id)
    ),
    Select='ALL_ATTRIBUTES',
)
r = {
    entry['SK'].split('_', 1)[0]: entry['SK'].split('_', 1)[1]
    for entry in r['Items']}
return r
```

```
class Sender(object):
```

"""Class to send messages over websockets."""

```
def __init__(self, app, storage):
```

"""Initialize a sender object.

:param app: A Chalice application object.

:param storage: A Storage object.

"""

```
self._app = app
```

```
self._storage = storage
```

```
def send(self, connection_id, message):
```

"""Send a message over a websocket.

:param connection_id: API Gateway Connection ID to send a message to.

:param message: The message to send to the connection.

"""

```
try:
```

Call the chalice websocket api send method

```
self._app.websocket_api.send(connection_id, message)
```

```
except WebsocketDisconnectedError as e:
```

*# If the websocket has been closed, we delete the connection
from our database.*

```
self._storage.delete_connection(e.connection_id)
```

```
def broadcast(self, connection_ids, message):
```

"""Send a message to multiple connections.

:param connection_id: A list of API Gateway Connection IDs to send the message to.

:param message: The message to send to the connections.

"""

```
for cid in connection_ids:
```

```
    self.send(cid, message)
```

```

class Handler(object):
    """Handler object that handles messages received from a websocket.

    This class implements the bulk of our app behavior.
    """
    def __init__(self, storage, sender):
        """Initialize a Handler object.

        :param storage: Storage object to interact with database.

        :param sender: Sender object to send messages to websockets.
        """
        self._storage = storage
        self._sender = sender
        # Command table to translate a string command name into a
        # method to call.
        self._command_table = {
            'help': self._help,
            'nick': self._nick,
            'join': self._join,
            'room': self._room,
            'quit': self._quit,
            'ls': self._list,
        }

    def handle(self, connection_id, message):
        """Entry point for our application.

        :param connection_id: Connection id that the message came from.

        :param message: Message we got from the connection.
        """
        # First look the user up in the database and get a record for it.
        record = self._storage.get_record_by_connection(connection_id)
        if record['username'] == '':
            # If the user does not have a username, we assume that the message
            # is the username they want and we call _handle_login_message.
            self._handle_login_message(connection_id, message)
        else:
            # Otherwise we assume the user is logged in. So we call
            # a method to handle the message. We pass along the
            # record we loaded from the database so we don't need to
            # again.
            self._handle_message(connection_id, message, record)

    def _handle_login_message(self, connection_id, message):
        """Handle a login message.

        The message is the username to give the user. Re-write the
        database entry for this user to reset their username from ''
        to {message}. Once that is done send a message back to the user
        to confirm the name choice. Also send a /help prompt.
        """
        self._storage.set_username(connection_id, '', message)
        self._sender.send(
            connection_id,
            'Using nickname: %s\nType /help for list of commands.' % message

```

```

)

def _handle_message(self, connection_id, message, record):
    """Handle a message from a connected and logged in user.

    If the message starts with a / it's a command. Otherwise its a
    text message to send to all rooms in the room.

    :param connection_id: Connection id that the message came from.

    :param message: Message we got from the connection.

    :param record: A data record about the sender.
    """
    if message.startswith('/'):
        self._handle_command(connection_id, message[1:], record)
    else:
        self._handle_text(connection_id, message, record)

def _handle_command(self, connection_id, message, record):
    """Handle a command message.

    Check the command name and look it up in our command table.
    If there is an entry, we call that method and pass along
    the connection_id, arguments, and the loaded record.

    :param connection_id: Connection id that the message came from.

    :param message: Message we got from the connection.

    :param record: A data record about the sender.
    """
    args = message.split(' ')
    command_name = args.pop(0).lower()
    command = self._command_table.get(command_name)
    if command:
        command(connection_id, args, record)
    else:
        # If no command method is found, send an error message
        # back to the user.
        self._sender(
            connection_id, 'Unknown command: %s' % command_name)

def _handle_text(self, connection_id, message, record):
    """Handle a raw text message.

    :param connection_id: Connection id that the message came from.

    :param message: Message we got from the connection.

    :param record: A data record about the sender.
    """
    if 'room' not in record:
        # If the user is not in a room send them an error message
        # and return early.
        self._sender.send(
            connection_id, 'Cannot send message if not in chatroom.')
    return

```

```

    # Collect a list of connection_ids in the same room as the message
    # sender.
    connection_ids = self._storage.get_connection_ids_by_room(
        record['room'])
    # Prefix the message with the sender's name.
    message = '%s: %s' % (record['username'], message)
    # Broadcast the new message to everyone in the room.
    self._sender.broadcast(connection_ids, message)

def _help(self, connection_id, _message, _record):
    """Send the help message.

    Build a help message and send back to the same connection.

    :param connection_id: Connection id that the message came from.
    """
    self._sender.send(
        connection_id,
        '\n'.join([
            'Commands available:',
            '    /help',
            '        Display this message.',
            '    /join {chat_room_name}',
            '        Join a chatroom named {chat_room_name}.',
            '    /nick {nickname}',
            '        Change your name to {nickname}. If no {nickname}',
            '        is provided then your current name will be printed',
            '    /room',
            '        Print out the name of the room you are currently ',
            '        in.',
            '    /ls',
            '        If you are in a room, list all users also in the ',
            '        room. Otherwise, list all rooms.',
            '    /quit',
            '        Leave current room.',
            '',
            'If you are in a room, raw text messages that do not start ',
            'with a / will be sent to everyone else in the room.',
        ]),
    )

def _nick(self, connection_id, args, record):
    """Change or check nickname (username).

    :param connection_id: Connection id that the message came from.

    :param args: Argument list that came after the command.

    :param record: A data record about the sender.
    """
    if not args:
        # If a nickname argument was not provided, we just want to
        # report the current nickname to the user.
        self._sender.send(
            connection_id, 'Current nickname: %s' % record['username'])
        return
    # The first argument is assumed to be the new desired nickname.
    nick = args[0]

```

```

# Change the username from record['username'] to nick in the storage
# layer.
self._storage.set_username(connection_id, record['username'], nick)
# Send a message to the requestor to confirm the nickname change.
self._sender.send(connection_id, 'Nickname is: %s' % nick)
# Get the room the user is in.
room = record.get('room')
if room:
    # If the user was in a room, announce to the room they have
    # changed their name. Don't send this message to the user since
    # they already got a name change message.
    room_connections = self._storage.get_connection_ids_by_room(room)
    room_connections.remove(connection_id)
    self._sender.broadcast(
        room_connections,
        '%s is now known as %s.' % (record['username'], nick))

def _join(self, connection_id, args, record):
    """Join a chat room.

    :param connection_id: Connection id that the message came from.

    :param args: Argument list. The first argument should be the
        name of the room to join.

    :param record: A data record about the sender.
    """
    # Get the room name to join.
    room = args[0]
    # Call quit to leave the current room we are in if there is any.
    self._quit(connection_id, '', record)
    # Get a list of connections in the target chat room.
    room_connections = self._storage.get_connection_ids_by_room(room)
    # Join the target chat room.
    self._storage.set_room(connection_id, room)
    # Send a message to the requestor that they have joined the room.
    # At the same time send an announcement to everyone who was already
    # in the room to alert them of the new user.
    self._sender.send(
        connection_id, 'Joined chat room "%s"' % room)
    message = '%s joined room.' % record['username']
    self._sender.broadcast(room_connections, message)

def _room(self, connection_id, _args, record):
    """Report the name of the current room.

    :param connection_id: Connection id that the message came from.

    :param record: A data record about the sender.
    """
    if 'room' in record:
        # If the user is in a room send them the name back.
        self._sender.send(connection_id, record['room'])
    else:
        # If the user is not in a room. Tell them so, and how to
        # join a room.
        self._sender.send(
            connection_id,

```

```

        'Not currently in a room. Type /join {room_name} to do so.'
    )

def _quit(self, connection_id, _args, record):
    """Quit from a room.

    :param connection_id: Connection id that the message came from.

    :param record: A data record about the sender.
    """
    if 'room' not in record:
        # If the user is not in a room there is nothing to do.
        return
    # Find the current room name, and delete that entry from
    # the database.
    room_name = record['room']
    self._storage.remove_room(connection_id, room_name)
    # Send a message to the user to inform them they left the room.
    self._sender.send(
        connection_id, 'Left chat room "%s"' % room_name)
    # Tell everyone in the room that the user has left.
    self._sender.broadcast(
        self._storage.get_connection_ids_by_room(room_name),
        '%s left room.' % record['username'],
    )

def _list(self, connection_id, _args, record):
    """Show a context dependent listing.

    :param connection_id: Connection id that the message came from.

    :param record: A data record about the sender.
    """
    room = record.get('room')
    if room:
        # If the user is in a room, get a listing of everyone
        # in the room.
        result = [
            self._storage.get_record_by_connection(c_id) ['username']
            for c_id in self._storage.get_connection_ids_by_room(room)
        ]
    else:
        # If they are not in a room. Get a listing of all rooms
        # currently open.
        result = self._storage.list_rooms()
    # Send the result list back to the requestor.
    self._sender.send(connection_id, '\n'.join(result))

```

The final directory layout should be

```

$ tree -a .
.
-- .chalice
|   -- config.json
-- .gitignore
-- app.py
-- chalicelib
|   -- __init__.py

```



```
-- resources.json
-- requirements.txt

2 directories, 6 files
```

Deploying our app with CloudFormation requires 3 steps. First we use Chalice to package our app into a JSON CloudFormation template:

```
$ chalice package --merge-template resources.json out
```

This will result in a new directory called `out` being created, inside which there is a `sam.json` file. This template contains our Chalice app as a CloudFormation template, merged with our `resources.json` template.

Next we use the AWS CLI to package this template, and prepare it for deployment. In order for this to work you will need to replace `$BUCKET` with the name of a bucket you control:

```
$ aws cloudformation package --template-file out/sam.json --s3-bucket $BUCKET --output-template-file
```

Once this is complete, a new template should be located at `out/template.yml` this is the final CloudFormation template which is ready for deployment. Deploying it with the AWS CLI can be done with the following command:

```
$ aws cloudformation deploy --template-file out/template.yml --stack-name ChaliceChat --capabilities
```

This command should wait awhile, and once it exits the app should be ready. To get the websocket connection URL, we can use the AWS CLI again to check the stack output `WebsocketConnectEndpointURL`:

```
$ aws cloudformation describe-stacks --stack-name ChaliceChat --query "Stacks[0].Outputs[?OutputKey==wss://{id}.execute-api.{region}.amazonaws.com/api/
```

Once deployed we can take the result of the previous command and connect to it using `wsdump.py`. Below is a sample of two running clients, the first message sent to the server is used as the client's username.

Listing 16.5: client-1

```
$ wsdump.py wss://{id}.execute-api.{region}.amazonaws.com/api/
Press Ctrl+C to quit
> John
< Using nickname: John
Type /help for list of commands.
> /help
< Commands available:
    /help
        Display this message.
    /join {chat_room_name}
        Join a chatroom named {chat_room_name}.
    /nick {nickname}
        Change your name to {nickname}. If no {nickname}
        is provided then your current name will be printed
    /room
        Print out the name of the room you are currently
        in.
    /ls
        If you are in a room, list all users also in the
        room. Otherwise, list all rooms.
    /quit
        Leave current room.
```

If you are in a room, raw text messages that **do** not start with a `/` will be sent to everyone **else** in the room.

```
> /join chalice
< Joined chat room "chalice"
< Jenny joined room.
> Hi
< John: Hi
< Jenny is now known as JennyJones.
> /quit
< Left chat room "chalice"
> /ls
< chalice
> Ctrl-D
```

Listing 16.6: client-2

```
$ wsdump.py wss://{id}.execute-api.{region}.amazonaws.com/api/
Press Ctrl+C to quit
> Jenny
< Using nickname: Jenny
Type /help for list of commands.
> /help
< Commands available:
    /help
        Display this message.
    /join {chat_room_name}
        Join a chatroom named {chat_room_name}.
    /nick {nickname}
        Change your name to {nickname}. If no {nickname}
        is provided then your current name will be printed
    /room
        Print out the name of the room you are currently
        in.
    /ls
        If you are in a room, list all users also in the
        room. Otherwise, list all rooms.
    /quit
        Leave current room.

If you are in a room, raw text messages that do not start
with a / will be sent to everyone else in the room.
> /join chalice
< Joined chat room "chalice"
> /ls
< John
Jenny
< John: Hi
> /nick JennyJones
< Nickname is: JennyJones
< John left room.
> /ls
< JennyJones
> /room
< chalice
> /nick
< Current nickname: JennyJones
> Ctrl-D
```

To delete the resources you can run use the AWS CLI to delete the stack:

```
$ aws cloudformation delete-stack --stack-name ChaliceChat
```

Upgrade Notes

17.1 Upgrade Notes

This document provides additional documentation on upgrading your version of chalice. If you're just interested in the high level changes, see the [CHANGELOG.rst](#) file.

17.1.1 1.2.0

This release features a rewrite of the Chalice deployer (#604). This is a backwards compatible change, and should not have any noticeable changes with deployments with the exception of fixing deployer bugs (e.g. <https://github.com/aws/chalice/issues/604>). This code path affects the `chalice deploy`, `chalice delete`, and `chalice package` commands.

While this release is backwards compatible, you will notice several changes when you upgrade to version 1.2.0.

The output of `chalice deploy` has changed in order to give more details about the resources it creates along with a more detailed summary at the end:

```
$ chalice deploy
Creating deployment package.
Creating IAM role: myapp-dev
Creating lambda function: myapp-dev-foo
Creating lambda function: myapp-dev
Creating Rest API
Resources deployed:
- Lambda ARN: arn:aws:lambda:us-west-2:12345:function:myapp-dev-foo
- Lambda ARN: arn:aws:lambda:us-west-2:12345:function:myapp-dev
- Rest API URL: https://abcd.execute-api.us-west-2.amazonaws.com/api/
```

Also, the files used to store deployed values has changed. These files are used internally by the `chalice deploy/delete` commands and you typically do not interact with these files directly. It's mentioned here in case you notice new files in your `.chalice` directory. Note that these files are *not* part of the public interface of Chalice and are documented here for completeness and to help with debugging issues.

In versions < 1.2.0, the value of deployed resources was stored in `.chalice/deployed.json` and looked like this:

```
{
  "dev": {
    "region": "us-west-2",
    "api_handler_name": "demoauth4-dev",
    "api_handler_arn": "arn:aws:lambda:us-west-2:123:function:myapp-dev",
```

```
"rest_api_id": "abcd",
"lambda_functions": {
  "myapp-dev-foo": {
    "type": "pure_lambda",
    "arn": "arn:aws:lambda:us-west-2:123:function:myapp-dev-foo"
  }
},
"chalice_version": "1.1.1",
"api_gateway_stage": "api",
"backend": "api"
},
"prod": {...}
}
```

In version 1.2.0, the deployed resources are split into multiple files, one file per chalice stage. These files are in the `.chalice/deployed/<stage.json>`, so if you had a dev and a prod chalice stage you'd have `.chalice/deployed/dev.json` and `.chalice/deployed/prod.json`. The schema has also changed and looks like this:

```
$ cat .chalice/deployed/dev.json
{
  "schema_version": "2.0",
  "resources": [
    {
      "role_name": "myapp-dev",
      "role_arn": "arn:aws:iam::123:role/myapp-dev",
      "name": "default-role",
      "resource_type": "iam_role"
    },
    {
      "lambda_arn": "arn:aws:lambda:us-west-2:123:function:myapp-dev-foo",
      "name": "foo",
      "resource_type": "lambda_function"
    },
    {
      "lambda_arn": "arn:aws:lambda:us-west-2:123:function:myapp-dev",
      "name": "api_handler",
      "resource_type": "lambda_function"
    },
    {
      "name": "rest_api",
      "rest_api_id": "abcd",
      "rest_api_url": "https://abcd.execute-api.us-west-2.amazonaws.com/api",
      "resource_type": "rest_api"
    }
  ],
  "backend": "api"
}
```

When you run `chalice deploy` for the first time after upgrading to version 1.2.0, chalice will automatically converted `.chalice/deployed.json` over to the format as you deploy a given stage.

Warning: Once you upgrade to 1.2.0, chalice will only update the new <code>.chalice/deployed/<stage>.json</code> . This means you cannot downgrade to earlier versions of chalice unless you manually update <code>.chalice/deployed.json</code> as well.
--

The `chalice package` command has also been updated to use the `deployer`. This results in several changes

compared to the previous version:

- Pure lambdas are supported
- Scheduled events are supported
- Parity between the behavior of `chalice deploy` and `chalice package`

As part of this change, the CFN resource names have been updated to use `CamelCase` names. Previously, chalice converted your python function names to CFN resource names by removing all non alphanumeric characters and appending an md5 checksum, e.g `my_function` -> `myfunction3bfc`. With this new packager update, the resource name would be converted as `my_function` -> `MyFunction`. Note, the `Outputs` section renames unchanged in order to preserve backwards compatibility. In order to fix parity issues with `chalice deploy` and `chalice package`, we now explicitly create an IAM role resource as part of the default configuration.

17.1.2 1.0.0b2

The url parameter names and the function argument names must match. Previously, the routing code would use positional args `handler(*args)` to invoke a view function. In this version, kwargs are now used instead: `handler(**view_args)`. For example, this code will no longer work:

```
@app.route("/{a}/{b}")
def myview(first, second)
    return {}
```

The example above must be updated to:

```
@app.route("/{a}/{b}")
def myview(a, b)
    return {}
```

Now that functions are invoked with kwargs, the order doesn't matter. You may also write the above view function as:

```
@app.route("/{a}/{b}")
def myview(b, a)
    return {}
```

This was done to have consistent behavior with other web frameworks such as Flask.

17.1.3 1.0.0b1

The `Chalice.define_authorizer` method has been removed. This has been deprecated since v0.8.1. See [Authorization](#) for updated information on configuring authorizers in Chalice as well as the original deprecation notice in the [0.8.1](#) upgrade notes.

The optional deprecated positional parameter in the `chalice deploy` command for specifying the API Gateway stage has been removed. If you want to specify the API Gateway stage, you can use the `--api-gateway-stage` option in the `chalice deploy` command:

```
# Deprecated and removed in 1.0.0b1
$ chalice deploy prod

# Equivalent and updated way to specify an API Gateway stage:
$ chalice deploy --api-gateway-stage prod
```

17.1.4 0.9.0

The 0.9.0 release changed the type of `app.current_request.raw_body` to always be of type `bytes()`. This only affects users that were using `python3`. Previously you would get a type `str()`, but with the introduction of [binary content type support](#), the `raw_body` attribute was made to consistently be of type `bytes()`.

17.1.5 0.8.1

The 0.8.1 changed the preferred way of specifying authorizers for view functions. You now specify either an instance of `chalice.CognitoUserPoolAuthorizer` or `chalice.CustomAuthorizer` to an `@app.route()` function using the `authorizer` argument.

Deprecated:

```
@app.route('/user-pools', methods=['GET'], authorizer_name='MyPool')
def authenticated():
    return {"secure": True}

app.define_authorizer(
    name='MyPool',
    header='Authorization',
    auth_type='cognito_user_pools',
    provider_arns=['arn:aws:cognito:...:userpool/name']
)
```

Equivalent, and preferred way

```
from chalice import CognitoUserPoolAuthorizer

authorizer = CognitoUserPoolAuthorizer(
    'MyPool', header='Authorization',
    provider_arns=['arn:aws:cognito:...:userpool/name'])

@app.route('/user-pools', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

The `define_authorizer` is still available, but is now deprecated and will be removed in future versions of chalice. You can also use the new `authorizer` argument to provide a `CustomAuthorizer`:

```
from chalice import CustomAuthorizer

authorizer = CustomAuthorizer(
    'MyCustomAuth', header='Authorization',
    authorizer_uri=('arn:aws:apigateway:region:lambda:path/2015-03-01'
                   '/functions/arn:aws:lambda:region:account-id:'
                   'function:FunctionName/invocations'))

@app.route('/custom-auth', methods=['GET'], authorizer=authorizer)
def authenticated():
    return {"secure": True}
```

17.1.6 0.7.0

The 0.7.0 release adds several major features to chalice. While the majority of these features are introduced in a backwards compatible way, there are a few backwards incompatible changes that were made in order to support these new major features.

Separate Stages

Prior to this version, chalice had a notion of a “stage” that corresponded to an API gateway stage. You can create and deploy a new API gateway stage by running `chalice deploy <stage-name>`. In 0.7.0, stage support was been reworked such that a chalice stage is a completely separate set of AWS resources. This means that if you have two chalice stages, say `dev` and `prod`, then you will have two separate sets of AWS resources, one set per stage:

- Two API Gateway Rest APIs
- Two separate Lambda functions
- Two separate IAM roles

The [Chalice Stages](#) doc has more details on the new chalice stages feature. This section highlights the key differences between the old stage behavior and the new chalice stage functionality in 0.7.0. In order to ease transition to this new model, the following changes were made:

- A new `--stage` argument was added to the `deploy`, `logs`, `url`, `generate-sdk`, and `package` commands. If this value is specified and the stage does not exist, a new chalice stage with that name will be created for you.
- The existing form `chalice deploy <stage-name>` has been deprecated. The command will still work in version 0.7.0, but a deprecation warning will be printed to `stderr`.
- If you want the pre-existing behavior of creating a new API gateway stage (while using the same Lambda function), you can use the `--api-gateway-stage` argument. This is the replacement for the deprecated form `chalice deploy <stage-name>`.
- The default stage if no `--stage` option is provided is `dev`. By defaulting to a `dev` stage, the pre-existing behavior of not specifying a stage name, e.g `chalice deploy`, `chalice url`, etc. will still work exactly the same.
- A new `stages` key is supported in the `.chalice/config.json`. This allows you to specify configuration specific to a chalice stage. See the [Configuration File](#) doc for more information about stage specific configuration.
- Setting `autogen_policy` to `false` will result in chalice looking for a IAM policy file named `.chalice/policy-<stage-name>.json`. Previously it would look for a file named `.chalice/policy.json`. You can also explicitly set this value to `In order to ease transition, chalice will check for a .chalice/policy.json file when deploying to the dev stage. Support for .chalice/policy.json will be removed in future versions of chalice and users are encouraged to switch to the stage specific .chalice/policy-<stage-name>.json files.`

See the [Chalice Stages](#) doc for more details on the new chalice stages feature.

Note, the AWS resource names it creates now have the form “<app-name>-<stage-name>”, e.g. “myapp-dev“, “myapp-prod“.

We recommend using the new stage specific resource names. However, If you would like to use the existing resource names for a specific stage, you can create a `.chalice/deployed.json` file that specifies the existing values:

```
{
  "dev": {
    "backend": "api",
    "api_handler_arn": "lambda-function-arn",
    "api_handler_name": "lambda-function-name",
    "rest_api_id": "your-rest-api-id",
    "api_gateway_stage": "dev",
    "region": "your region (e.g us-west-2)",
    "chalice_version": "0.7.0",
```

```
}  
}
```

This file is discussed in the next section.

Deployed Values

In version 0.7.0, the way deployed values are stored and retrieved has changed. In prior versions, only the `lambda_arn` was saved, and its value was written to the `.chalice/config.json` file. Any of other deployed values that were needed (for example the API Gateway rest API id) was dynamically queried by assuming the resource names matches the app name. In this version of chalice, a separate `.chalice/deployed.json` file is written on every deployment which contains all the resources that have been created. While this should be a transparent change, you may noticed issues if you run commands such as `chalice url` and `chalice logs` without first deploying. To fix this issue, run `chalice deploy` and version 0.7.0 of chalice so a `.chalice/deployed.json` will be created for you.

Authorizer Changes

The “`authorizer_id`” and “`authorization_type`” args are no longer supported in “`@app.route(...)`” calls.

They have been replaced with an `authorizer_name` parameter and an `app.define_authorizer` method.

This version changed the internals of how an API gateway REST API is created. Prior to 0.7.0, the AWS SDK for Python was used to make the appropriate service API calls to API gateway include `create_rest_api` and `put_method / put_method_response` for each route. In version 0.7.0, this internal mechanism was changed to instead generate a swagger document. The rest api is then created or updated by calling `import_rest_api` or `put_rest_api` and providing the swagger document. This simplifies the internals and also unifies the code base for the newly added `chalice package` command (which uses a swagger document internally). One consequence of this change is that the entire REST API must be defined in the swagger document. With the previous `authorizer_id` parameter, you would create/deploy a rest api, create your authorizer, and then provide that `authorizer_id` in your `@app.route` calls. Now they must be defined all at once in the `app.py` file:

```
app = chalice.Chalice(app_name='demo')  
  
@app.route('/auth-required', authorizer_name='MyUserPool')  
def foo():  
    return {}  
  
app.define_authorizer(  
    name='MyUserPool',  
    header='Authorization',  
    auth_type='cognito_user_pools',  
    provider_arns=['arn:aws:cognito:...:userpool/name']  
)
```

17.1.7 0.6.0

This version changed how the internals of how API gateway resources are created by chalice. The integration type changed from `AWS` to `AWS_PROXY`. This was to enable additional functionality, notable to allows users to provide non-JSON HTTP responses and inject arbitrary headers to the HTTP responses. While this change to the internals is primarily internal, there are several user-visible changes.

- Uncaught exceptions with `app.debug = False` (the default value) will result in a more generic `InternalServerError` error. The previous behavior was to return a `ChaliceViewError`.

- When you enabled debug mode via `app.debug = True`, the HTTP response will contain the python stack trace as the entire request body. This is to improve the readability of stack traces. For example:

```
$ http https://endpoint/dev/
HTTP/1.1 500 Internal Server Error
Content-Length: 358
Content-Type: text/plain

Traceback (most recent call last):
  File "/var/task/chalice/app.py", line 286, in __call__
    response = view_function(*function_args)
  File "/var/task/app.py", line 12, in index
    return a()
  File "/var/task/app.py", line 16, in a
    return b()
  File "/var/task/app.py", line 19, in b
    raise ValueError("Hello, error!")
ValueError: Hello, error!
```

- Content type validation now has error responses that match the same error response format used for other chalice built in responses. Chalice was previously relying on API gateway to perform the content type validation. As a result of the AWS_PROXY work, this logic has moved into the chalice handler and now has a consistent error response:

```
$ http https://endpoint/dev/ 'Content-Type: text/plain'
HTTP/1.1 415 Unsupported Media Type
Content-Type: application/json

{
  "Code": "UnsupportedMediaType",
  "Message": "Unsupported media type: text/plain"
}
```

- The keys in the `app.current_request.to_dict()` now match the casing used by the AWS_PROXY lambda integration, which are `lowerCamelCased`. This method is primarily intended for introspection purposes.

17.2 Indices and tables

- `genindex`
- `search`

Symbols

`__iter__()` (SQSEvent method), 97

A

account (CloudWatchEvent attribute), 95
allow_credentials (CORSConfig attribute), 94
allow_headers (CORSConfig attribute), 93
allow_origin (CORSConfig attribute), 93
api (Chalice attribute), 83
APIGateway (built-in class), 91
auth_type (AuthRequest attribute), 90
authorizer() (Chalice method), 84
authorizer_uri (CustomAuthorizer attribute), 90
AuthRequest (built-in class), 90
AuthResponse (built-in class), 90
AuthRoute (built-in class), 91

B

binary_types (APIGateway attribute), 91
Blueprint (built-in class), 98
body (Response attribute), 89
body (SQSRecord attribute), 97
body (WebsocketEvent attribute), 98
bucket (S3Event attribute), 96

C

Chalice (built-in class), 83
close(), 92
CloudWatchEvent (built-in class), 95
CognitoUserPoolAuthorizer (built-in class), 90
configure() (WebsocketAPI method), 92
connection_id (WebsocketDisconnectedError attribute), 93
connection_id (WebsocketEvent attribute), 98
context (AuthResponse attribute), 91
context (CloudWatchEvent attribute), 96
context (Request attribute), 89
context (S3Event attribute), 96
context (SNSEvent attribute), 97
context (SQSEvent attribute), 97

context (SQSRecord attribute), 97
cors (APIGateway attribute), 91
CORSConfig (built-in class), 93
Cron (built-in class), 94
current_request (Chalice attribute), 83
CustomAuthorizer (built-in class), 90

D

debug (Chalice attribute), 83
default_binary_types (APIGateway attribute), 91
detail (CloudWatchEvent attribute), 95
detail_type (CloudWatchEvent attribute), 95
domain_name (WebsocketEvent attribute), 98

E

event_id (CloudWatchEvent attribute), 95
expose_headers (CORSConfig attribute), 94

H

header (CognitoUserPoolAuthorizer attribute), 90
header (CustomAuthorizer attribute), 90
headers (Request attribute), 88
headers (Response attribute), 89

I

IAMAuthorizer (built-in class), 90
info(), 93

J

json_body (Request attribute), 89
json_body (WebsocketEvent attribute), 98

K

key (S3Event attribute), 96

L

lambda_context (Chalice attribute), 83
lambda_function() (Chalice method), 87

M

max_age (CORSConfig attribute), 94
message (SNSEvent attribute), 97
method (Request attribute), 88
method_arn (AuthRequest attribute), 90
methods (AuthRoute attribute), 91

N

name (CognitoUserPoolAuthorizer attribute), 90
name (CustomAuthorizer attribute), 90

O

on_cw_event() (Chalice method), 85
on_s3_event() (Chalice method), 85
on_sns_message() (Chalice method), 86
on_sqs_message() (Chalice method), 87
on_ws_connect() (Chalice method), 88
on_ws_disconnect() (Chalice method), 88
on_ws_message() (Chalice method), 88

P

path (AuthRoute attribute), 91
principal_id (AuthResponse attribute), 91
provider_arns (CognitoUserPoolAuthorizer attribute), 90

Q

query_params (Request attribute), 88

R

Rate (built-in class), 94
raw_body (Request attribute), 89
receipt_handle (SQSRecord attribute), 97
region (CloudWatchEvent attribute), 95
register_blueprint() (Chalice method), 87
Request (built-in class), 88
resources (CloudWatchEvent attribute), 96
Response (built-in class), 89
route() (Chalice method), 83
routes (AuthResponse attribute), 90

S

S3Event (built-in class), 96
schedule() (Chalice method), 85
send(), 92
session (WebsocketAPI attribute), 92
SNSEvent (built-in class), 96
source (CloudWatchEvent attribute), 95
SQSEvent (built-in class), 97
SQSRecord (built-in class), 97
stage (WebsocketEvent attribute), 98
stage_vars (Request attribute), 89
status_code (Response attribute), 89
subject (SNSEvent attribute), 97

T

time (CloudWatchEvent attribute), 95
to_dict() (CloudWatchEvent method), 96
to_dict() (Request method), 89
to_dict() (S3Event method), 96
to_dict() (SNSEvent method), 97
to_dict() (SQSEvent method), 97
to_dict() (SQSRecord method), 97
token (AuthRequest attribute), 90
ttl_seconds (CustomAuthorizer attribute), 90

U

unit (Rate attribute), 94
uri_params (Request attribute), 88

V

value (Rate attribute), 94
version (CloudWatchEvent attribute), 95

W

websocket_api (Chalice attribute), 83
WebsocketAPI (built-in class), 92
WebsocketDisconnectedError (built-in class), 93
WebsocketEvent (built-in class), 98