



Red Hat 3scale API Management 2.1

Deployment Options

Deploy APIcast API Gateway on OpenShift, natively, or using Docker.

Red Hat 3scale API Management 2.1 Deployment Options

Deploy APIcast API Gateway on OpenShift, natively, or using Docker.

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide documents deployment options for Red Hat 3scale API Management 2.1.

Table of Contents

CHAPTER 1. APICAST OVERVIEW	4
1.1. PREREQUISITES	4
1.2. DEPLOYMENT OPTIONS	4
1.3. ENVIRONMENTS	4
1.4. API CONFIGURATION	4
1.5. CONFIGURE THE INTEGRATION SETTINGS	5
1.6. CONFIGURE YOUR SERVICE	5
1.7. MAPPING RULES	6
1.8. MAPPING RULES WORKFLOW	7
1.9. HOST HEADER	8
1.10. PRODUCTION DEPLOYMENT	8
1.11. PUBLIC BASE URL	8
1.12. PROTECTING YOUR API BACKEND	9
1.13. USING APICAST WITH PRIVATE APIS	9
CHAPTER 2. APICAST HOSTED	10
2.1. PREREQUISITES	10
2.2. STEP 1: DEPLOY YOUR API WITH APICAST HOSTED IN A STAGING ENVIRONMENT	10
2.3. STEP 2: DEPLOY YOUR API WITH THE APICAST HOSTED INTO PRODUCTION	11
2.3.1. Bear in mind	11
CHAPTER 3. APICAST SELF-MANAGED	12
3.1. PREREQUISITES	12
3.2. STEP 1: INSTALL OPENRESTY AND DEPENDENCIES	12
3.3. STEP 2: DEPLOY AND RUN APICAST	13
CHAPTER 4. APICAST ON THE DOCKER CONTAINERIZED ENVIRONMENT	14
4.1. PREREQUISITES	14
4.2. STEP 1: INSTALL THE DOCKER CONTAINERIZED ENVIRONMENT	14
4.3. STEP 2: RUN THE DOCKER CONTAINERIZED ENVIRONMENT GATEWAY	14
4.3.1. The Docker command options	15
4.4. STEP 3: TESTING APICAST	15
4.5. STEP 4: TROUBLESHOOTING APICAST ON THE DOCKER CONTAINERIZED ENVIRONMENT	15
4.5.1. Cannot connect to the Docker daemon error	15
4.5.2. Basic Docker command-line interface commands	16
4.6. STEP 5: CUSTOMISING THE GATEWAY	16
4.6.1. Custom Lua logic	16
4.6.2. Customising the configuration files	16
CHAPTER 5. RUNNING APICAST ON RED HAT OPENSIFT	17
5.1. PREREQUISITES	17
5.2. STEP 1: SET UP OPENSIFT	17
5.2.1. Install the Docker containerized environment	17
5.2.2. Start OpenShift cluster	18
5.2.2.1. Setting up OpenShift cluster on a remote server	18
5.3. STEP 2: DEPLOY APICAST USING THE OPENSIFT TEMPLATE	19
5.4. STEP 3: CREATE ROUTES IN OPENSIFT CONSOLE	20
CHAPTER 6. ADVANCED APICAST CONFIGURATION	23
6.1. DEFINE A SECRET TOKEN	23
6.2. CREDENTIALS	23
6.3. ERROR MESSAGES	24

6.4. CONFIGURATION HISTORY	25
6.5. DEBUGGING	26
6.6. EXTENDING THE GATEWAY	27
CHAPTER 7. CODE LIBRARIES	28
7.1. PYTHON	28
7.2. RUBY GEM	28
7.3. PERL	28
7.4. PHP	28
7.5. JAVA	28
7.6. .NET	29
7.7. NODE.JS	29
CHAPTER 8. PLUGIN SETUP	30
8.1. HOW PLUGINS WORK	30
8.2. STEP 1: SELECT YOUR LANGUAGE AND DOWNLOAD THE PLUGIN	30
8.3. STEP 2: CHECK THE METRICS YOU'VE SET FOR YOUR API	31
8.4. STEP 3: INSTALL THE PLUGIN CODE TO YOUR APPLICATION	31
8.5. STEP 4: ADD CALLS TO AUTHORIZE AS API TRAFFIC ARRIVES	33
8.5.1. Usage on authrep mode	33
8.6. STEP 5: DEPLOY AND RUN TRAFFIC	34

CHAPTER 1. APICAST OVERVIEW

APIcast is an NGINX based API gateway used to integrate your internal and external API services with the Red Hat 3scale Platform.

The latest released and supported version of APIcast is 2.0.

In this guide you will learn more about deployment options, environments provided, and how to get started.

1.1. PREREQUISITES

APIcast is not a standalone API gateway, it needs connection to 3scale API Manager.

- You will need a working 3scale [On-Premises](#) instance.

1.2. DEPLOYMENT OPTIONS

You can use APIcast hosted or self-managed, in both cases, it needs connection to the rest of the 3scale API management platform:

- **APIcast hosted:** 3scale hosts APIcast in the cloud. In this case, APIcast is already deployed for you and it is limited to 50,000 calls per day.
- **APIcast built-in:** Two APIcast (staging and production) come by default with the 3scale API Management Platform (AMP) installation. They come pre-configured and ready to use out-of-the-box.
- **APIcast self-managed:** You can deploy APIcast wherever you want. The self-managed mode is the intended mode of operation for production environments. Here are a few recommended options to deploy APIcast:
 - **Native deployment:** Install OpenResty and other dependencies on your own server and run APIcast using the code and configuration provided by 3scale.
 - **the Docker containerized environment:** Download a ready to use Docker-formatted container image, which includes all of the dependencies to run APIcast in a Docker-formatted container.
 - **OpenShift:** Run APIcast on a [supported version](#) of OpenShift. You can connect self-managed APIcasts both to a 3scale AMP installation or to a 3scale online account.

1.3. ENVIRONMENTS

By default, when you create a 3scale account, you get APIcast **hosted** in two different environments:

- **Staging:** Intended to be used only while configuring and testing your API integration. When you have confirmed that your setup is working as expected, then you can choose to deploy it to the production environment.
- **Production:** Limited to 50,000 calls per day and supports the following out-of-the-box authentication options: API key, and App ID and App key pair.

1.4. API CONFIGURATION

Follow the next steps to configure APIcast in no time.

1.5. CONFIGURE THE INTEGRATION SETTINGS

Go to the **Dashboard** → **API** tab and click on the Integration link. If you have more than one API in your account, you will need to select the API first.

On top of the Integration page you will see your integration options. By default, the deployment option is APIcast hosted, and the authentication mode is API key. You can change these settings by clicking on **edit integration settings** in the top right corner. Note that OAuth 2.0 authentication is only available for the Self-managed deployment.

1.6. CONFIGURE YOUR SERVICE

You will need to declare your API backend in the Private Base URL field, which is the endpoint host of your API backend. APIcast will redirect all traffic to your API backend after all authentication, authorization, rate limits and statistics have been processed.

Typically, the Private Base URL of your API will be something like <https://api-backend.yourdomain.com:443>, on the domain that you manage (**yourdomain.com**). For instance, if you were integrating with the Twitter API the Private Base URL would be <https://api.twitter.com/>. In this example will use the Echo API hosted by 3scale – a simple API that accepts any path and returns information about the request (path, request parameters, headers, etc.). Its Private Base URL is <https://echo-api.3scale.net:443>. Private Base URL

Staging: configure & test your integration [documentation](#)

[deployed](#) | [deployment history](#)

The screenshot shows a configuration page for an API. On the left, there is a green puzzle piece icon labeled 'API'. To the right, there is a search bar with a question mark icon. Below the search bar, the 'Private Base URL*' field is filled with the text 'https://echo-api.3scale.net:443'. A tooltip below the field reads 'Private address of your API that will be called by the API gateway.'

Test your private (unmanaged) API is working. For example, for the Echo API we can make the following call with **curl** command:

```
curl "https://echo-api.3scale.net:443"
```

You will get the following response:

```
{
  "method": "GET",
  "path": "/",
  "args": "",
  "body": "",
  "headers": {
    "HTTP_VERSION": "HTTP/1.1",
    "HTTP_HOST": "echo-api.3scale.net",
    "HTTP_ACCEPT": "**/*",
    "HTTP_USER_AGENT": "curl/7.51.0",
    "HTTP_X_FORWARDED_FOR": "2.139.235.79, 10.0.103.58",
    "HTTP_X_FORWARDED_HOST": "echo-api.3scale.net",
    "HTTP_X_FORWARDED_PORT": "443",
    "HTTP_X_FORWARDED_PROTO": "https",
    "HTTP_FORWARDED": "for=10.0.103.58;host=echo-api.3scale.net;proto=https"
```

```

    },
    "uuid": "ee626b70-e928-4cb1-a1a4-348b8e361733"
  }
}

```

Once you have confirmed that your API is working, you will need to configure the test call for the hosted staging environment. Enter a path existing in your API in the *API test GET request field* (for example, `/v1/word/good.json`).

Save the settings by clicking on the **Update & Test Staging Configuration** button in the bottom right part of the page. This will deploy the APIcast configuration to the 3scale hosted staging environment. If everything is configured correctly, the vertical line on the left should turn green.

If you are using one of the Self-managed deployment options save the configuration from the GUI and make sure it is pointing to your deployed API gateway by adding the correct host in the staging or production Public base URL field. Before making any calls to your production gateway, do not forget to click on the **Promote v.x to Production** button.

Find the sample **curl** at the bottom of the staging section and run it from the console:

```
curl "https://XXX.staging.apicast.io:443/v1/word/good.json?user_key=YOUR_USER_KEY"
```



NOTE

You should get the same response as above, however, this time the request will go through the 3scale hosted APIcast instance. Note: You should make sure you have an application with valid credentials for the service. If you are using the default API service created on sign up to 3scale, you should already have an application. Otherwise, if you see **USER_KEY** or **APP_ID** and **APP_KEY** values in the test curl, you need to create an application for this service first.

Now you have your API integrated with 3scale.

3scale hosted APIcast gateway does the validation of the credentials and applies the rate limits that you defined for the application plan of the application. If you try to make a call without credentials, or with invalid credentials, you will see an error message. The code and the text of the message can be configured, check out the [Advanced APIcast configuration](#) article for more information.

1.7. MAPPING RULES

By default we start with a very simple mapping rule,

Verb	Pattern		Metric or Method (Define)
GET	/	1	hits

[Add Mapping Rule](#)

This rule means that any **GET** request that starts with `/` will increment the metric **hits** by 1. This mapping rule will match any request to your API. Most likely you will change this rule since it is too generic.

The mapping rules define which metrics (and methods) you want to report depending on the requests to your API. For instance, below you can see the rules for the Echo API that serves us as an example:

▼ MAPPING RULES ?

Verb	Pattern	+	Metric or Method (Define)
GET	/hello	1	gethello ✎ 🗑️
GET	/goodbye	1	getgoodby ✎ 🗑️

➕ Add Mapping Rule

The matching of the rules is done by prefix and can be arbitrarily complex (the notation follows Swagger and ActiveDocs specification)

- You can do a match on the path over a literal string: **/hello**
- Mapping rules can contain named wildcards: **/{word}**

This rule will match anything in the placeholder **{word}**, making requests like **/morning** match the rule.

Wildcards can appear between slashes or between slash and dot.

- Mapping rules can also include parameters on the query string or in the body: **/{word}?value={value}**

APIcast will try to fetch the parameters from the query string when it is a GET and from the body when it is a POST, DELETE, PUT.

Parameters can also have named wildcards.

Note that all mapping rules are evaluated. There is no precedence (order does not matter). If we added a rule **/v1** to the example on the figure above, it would always be matched for the requests whose path starts with **/v1** regardless if it is **/v1/word** or **/v1/sentence**. Keep in mind that if two different rules increment the same metric by one, and the two rules are matched, the metric will be incremented by two.

1.8. MAPPING RULES WORKFLOW

The intended workflow to define mapping rules is as follows:

- You can add new rules by clicking the **Add Mapping Rule** button. Then you select an HTTP method, a pattern, a metric (or method) and finally its increment. When you are done, click **Update & Test Staging Configuration** to apply the changes.
- Mapping rules will be grayed out on the next reload to prevent accidental modifications.
- To edit an existing mapping rule you must enable it first by clicking the pencil icon on the right.
- To delete a rule click on the trash icon.
- Modifications and deletions will be saved when you hit the **Update & Test Staging Configuration** button.

For more advanced configuration options, you can check the [APIcast advanced configuration](#) tutorial.

1.9. HOST HEADER

This option is only needed for those API backends that reject traffic unless the **Host** header matches the expected one. In these cases, having a gateway in front of your API backend will cause problems since the **Host** will be the one of the gateway, e.g. **xxx-yyy.staging.apicast.io**

To avoid this issue you can define the host your API backend expects in the **Host Header** field in the Authentication Settings, and the hosted APIcast instance will rewrite the host.

▼ AUTHENTICATION SETTINGS

Host Header

Lets you define a custom `Host` request header. This is needed if your API backend only accepts traffic from a specific host.

1.10. PRODUCTION DEPLOYMENT

Once you have configured your API integration and verified it is working in the Staging environment, you can go ahead with one of the APIcast production deployments. See the [Deployment options](#) in the beginning of this article.

At the bottom of the Integration page you will find the *Production* section. You will find two fields here: the *Private Base URL*, which will be the same as you configured in the *Staging* section, and the *Public Base URL*.

1.11. PUBLIC BASE URL

The **Public Base URL** is the URL, which your developers will use to make requests to your API, protected by 3scale. This will be the URL of your APIcast instance.

In [APIcast hosted](#), the Public Base URL is set by 3scale and cannot be changed.

If you are using one of the Self-managed deployment options, you can choose your own Public Base URL for each one of the environments provided (staging and production), on a domain name you are managing. Note that this URL should be different from the one of your API backend, and could be something like <https://api.yourdomain.com:443>, where **yourdomain.com** is the domain that belongs to you. After setting the Public Base URL make sure you save the changes and, if necessary, promote the changes in staging to production.

Please note that APIcast v2 will only accept calls to the hostname which is specified in the Public Base URL. For example, if for the Echo API used as an example above, we specify <https://echo-api.3scale.net:443> as the Public Base URL, the correct call would be be:

```
curl "https://echo-api.3scale.net:443/hello?user_key=YOUR_USER_KEY"
```

In case you do not yet have a public domain for your API, you can also use the APIcast IP in the requests, but you still need to specify a value in the Public Base URL field (even if the domain is not real), and in this case make sure you provide the host in the Host header, for example:

```
curl "http://192.0.2.12:80/hello?user_key=YOUR_USER_KEY" -H "Host: echo-api.3scale.net"
```

If you are deploying on local machine, you can also just use "localhost" as the domain, so the Public Base URL will look like <http://localhost:80>, and then you can make requests like this:

```
curl "http://localhost:80/hello?user_key=YOUR_USER_KEY"
```

In case you have multiple API services, you will need to set this Public Base URL appropriately for each service. APIcast will route the requests based on the hostname.

1.12. PROTECTING YOUR API BACKEND

Once you have APIcast working in production, you might want to restrict direct access to your API backend without credentials. The easiest way to do this is by using the Secret Token set by APIcast. Please refer to the [Advanced APIcast configuration](#) for information on how to set it up.

1.13. USING APICAST WITH PRIVATE APIS

With APIcast it is possible to protect the APIs which are not publicly accessible on the Internet. The requirements that must be met are:

- APIcast self-managed must be used as the deployment option
- APIcast needs to be accessible from the public internet and be able to make outbound calls to the 3scale Service Management API
- the API backend should be accessible by APIcast

In this case you can set your internal domain name or the IP address of your API in the *Private Base URL* field and follow the rest of the steps as usual. Note, however, that you will not be able to take advantage of the Staging environment, and the test calls will not be successful, as the Staging APIcast instance is hosted by 3scale and will not have access to your private API backend). But once you deploy APIcast in your production environment, if the configuration is correct, APIcast will work as expected.

CHAPTER 2. APICAST HOSTED

Once you complete this tutorial, you'll have your API fully protected by a secure gateway in the cloud.

APIcast hosted is the best deployment option if you want to launch your API as fast as possible, or if you want to make the minimum infrastructure changes on your side.

2.1. PREREQUISITES

- You have reviewed the [deployment alternatives](#) and decided to use APIcast hosted to integrate your API with 3scale.
- Your API backend service is accessible over the public Internet (a secure communication will be established to prevent users from bypassing the access control gateway).
- You do not expect demand for your API to exceed the limit of 50,000 hits/day (beyond this, we recommend upgrading to the self-managed gateway).

2.2. STEP 1: DEPLOY YOUR API WITH APICAST HOSTED IN A STAGING ENVIRONMENT

The first step is to configure your API and test it in your staging environment. Define the private base URL and its endpoints, choose the placement of credentials and other configuration details that you can read about [here](#). Once you're done entering your configuration, go ahead and click on Update & Test Staging Environment button to run a test call that will go through the APIcast staging instance to your API.

Configuration: configure & test immediately in the staging environment [documentation](#)

The screenshot shows a configuration interface for APIcast. It is divided into three main sections: API, API GATEWAY, and CLIENT. Each section has a title, a field for configuration, and a help icon (question mark). The API section has a 'Private Base URL*' field with the value 'https://echo-api.3scale.net:443'. The API GATEWAY section has 'Public Base URL*' and 'Production Public Base URL*' fields, both with values starting with 'https://api-2445581460490.staging.apicast.io:443'. Below these are expandable sections for 'MAPPING RULES' and 'AUTHENTICATION SETTINGS'. The CLIENT section has an 'API test GET request' field with the value '/'. Below this field is a code block containing a curl command. At the bottom right is a blue button labeled 'Update & Test in Staging Environment'. At the bottom left is a note: 'Hit the test button to check the connections between client, gateway & API.'

API ?

Private Base URL*
Private address of your API that will be called by the API gateway.

API GATEWAY ?

Public Base URL*
Public address of your API gateway in the staging environment.

Production Public Base URL*
Public address of your API gateway in the production environment.

▶ **MAPPING RULES**

▶ **AUTHENTICATION SETTINGS**

CLIENT ?

API test GET request
Optional GET request to a API gateway endpoint. We will use this call to validate your API gateway setup using credentials of the first live application. You can try it yourself by copying the following command into your shell:

```
curl "https://api-2445581460490.staging.apicast.io:443/?user_key=063a01e356790b831f749b0b8b726e38"
```

Hit the test button to check the connections between client, gateway & API.

Update & Test in Staging Environment

[← Back to Integration & Configuration](#)

If everything was configured correctly, you should see a green confirmation message.

Before moving on to the next step, make sure that you have configured a secret token to be validated by your backend service. You can define the value for the secret token under **Authentication Settings**. This will ensure that nobody can bypass APIcast's access control.

2.3. STEP 2: DEPLOY YOUR API WITH THE APICAST HOSTED INTO PRODUCTION

At this point, you're ready to take your API configuration to a production environment. To deploy your 3scale-hosted APIcast instance, go back to the 'Integration and Configuration' page and click on the '**Promote to v.x to Production**' button. Repeat this step to promote further changes in your staging environment to your production environment.

The screenshot shows the APIcast configuration and environment management interface. At the top, there is a section for 'APICast Configuration' with a link to 'edit APICast configuration'. Below this, the configuration details are listed: Private Base URL: https://echo-api.3scale.net:443, Mapping rules: / => h1ta, Credential Location: query, and Secret Token: Shared_secret_sent_from_proxy_to_API_backend. Below the configuration section is the 'Environments' section, which contains a 'Staging Environment' and a 'Production Environment'. The 'Staging Environment' section shows the URL https://api-244581460490.staging.apicast.io:443 and a blue button labeled 'Promote v. 1 to Production'. The 'Production Environment' section shows the message 'no configuration has been saved for the production environment yet'. A link for 'Configuration history' is also visible in the top right corner of the Environments section.

It will take between 5 and 7 minutes for your configuration to deploy and propagate to all the cloud APIcast instances. During redeployment, your API will not experience any downtime. However, API calls may return different responses depending on which instance serves the call. You'll know it has been deployed once the box around your production environment has turned green.

Both the staging and production APIcast instances have base URLs on the apicast.io domain. You can easily tell them apart because the staging environment URLs have a staging subdomain. For example:

- staging: <https://api-2445581448324.staging.apicast.io:443>
- production: <https://api-2445581448324.apicast.io:443>

2.3.1. Bear in mind

- 50,000 hits/day is the maximum allowed for your API through the APIcast production cloud instance. You can check your API usage in the Analytics section of your Admin Portal.
- There is a hard throttle limit of 20 hits/second on any spike in API traffic.
- Above the throttle limit, APIcast returns a response code of **403**. This is the same as the default for an application over rate limits. If you want to differentiate the errors, please check the response body.

CHAPTER 3. APICAST SELF-MANAGED

This tutorial shows the necessary steps to deploy the latest version of APIcast on your own server to have it ready to be used as a 3scale API gateway.

For details on the previous version of self-managed APIcast (Nginx downloadable configuration files), please see [here](#)

3.1. PREREQUISITES

You will need to configure APIcast in your 3scale Admin Portal as per the [APIcast Overview](#), if you haven't done so already. Make sure *Self-managed Gateway* is selected as the deployment option in the integration settings.

You should also have a server where you'll deploy your API gateway(s). This tutorial covers how to install your self-managed APIcast instance on a server running Red Hat Enterprise Linux – the operating system supported by the Red Hat 3scale API management platform. The server can be located either in the cloud, or on premise.

3.2. STEP 1: INSTALL OPENRESTY AND DEPENDENCIES

APIcast requires some external modules for NGINX. Even though it's possible to compile NGINX with these modules from source, we strongly recommend using OpenResty – an excellent bundle that already includes all the necessary modules.

This guide covers the steps to set up the official pre-built packages that OpenResty provides for Red Hat Enterprise Linux (RHEL) version 7. The latest installation instructions can be found in the [OpenResty documentation](#).

For other operating systems please refer to the [OpenResty installation instructions](#).

First, add the *openresty* repository to your RHEL system by creating the file `/etc/yum.repos.d/OpenResty.repo` with the following content:

```
[openresty]
name=Official OpenResty Repository
baseurl=https://copr-be.cloud.fedoraproject.org/results/openresty/openresty/epel-7-$basearch/
skip_if_unavailable=True
gpgcheck=1
gpgkey=https://copr-be.cloud.fedoraproject.org/results/openresty/openresty/pubkey.gpg
enabled=1
enabled_metadata=1
```

Install OpenResty and the **resty** command-line utility with the following command:

```
sudo yum install openresty openresty-resty
```

You can learn more about these and other OpenResty packages in [OpenResty documentation](#).

APIcast uses LuaRocks for managing Lua dependencies. As it's not in the standard Yum repositories, you must first enable the [EPEL](#) (Extra Packages for Enterprise Linux) package repository. Please refer to the [following article](#) on Red Hat Customer Portal for more information on how to enable EPEL on RHEL.

For RHEL 7 you can run the following command:

```
sudo yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

Install LuaRocks:

```
sudo yum install luarocks
```

3.3. STEP 2: DEPLOY AND RUN APICAST

In the latest version of APICast the code that implements the gateway logic is separated from the configuration of your API. In order to deploy APICast to your server, you'll first need to get the APICast code.

To use the latest stable version of APICast, check the APICast [releases page](#).

Go to the **apicast** directory, that you checked out with git or extracted from the downloaded archive.

Run the following command to install all the Lua dependencies:

```
sudo luarocks make apicast/*.rockspec --tree /usr/local/openresty/luajit
```

You can start APICast using the **bin/apicast** executable included in the package.

```
THREESCALE_PORTAL_ENDPOINT=https://<access_token>@<domain>-admin.3scale.net  
bin/apicast
```

Here **<access_token>** is an [Access Token](#) for the 3scale Account Management API, and **<domain>-admin.3scale.net** is the URL of your 3scale Admin Portal.

This command will start APICast and download the latest APICast configuration from the 3scale Admin Portal.

bin/apicast executable accepts a number of options, you can check them out by running:

```
bin/apicast -h
```

Additional parameters can be specified using environment variables.

```
APICAST_LOG_FILE=logs/error.log bin/apicast -c config.json -d -v -v -v
```

The above command will run APICast as a daemon (**-d** option,) using the configuration file **config.json**, with the error logs at **debug** level (**-v -v -v**) and written to the **logs/error.log** file inside the **apicast** directory (the **prefix** directory).

And that's it! You should now be running APICast Self-Managed. If you're looking to install APICast on [OpenShift](#) or using the [Docker containerized environment](#), we recommend checking out those tutorials.

CHAPTER 4. APICAST ON THE DOCKER CONTAINERIZED ENVIRONMENT

This is a step-by-step guide to deploy APIcast inside a Docker-formatted container ready to be used as a 3scale API gateway.

4.1. PREREQUISITES

You must configure APIcast in your 3scale Admin Portal as per the [APIcast Overview](#).

4.2. STEP 1: INSTALL THE DOCKER CONTAINERIZED ENVIRONMENT

This guide covers the steps to set up the Docker containerized environment on Red Hat Enterprise Linux (RHEL) 7.

Docker-formatted containers provided by Red Hat are released as part of the Extras channel in RHEL. To enable additional repositories, you can use either the [Subscription Manager](#) or the yum config manager. For details, see the [RHEL product documentation](#).

To deploy RHEL 7 on an AWS EC2 instance, take the following steps:

1. List all repositories: **sudo yum repolist all**.
2. Find the ***-extras** repository.
3. Enable the **extras** repository: **sudo yum-config-manager --enable rhui-REGION-rhel-server-extras**.
4. Install the Docker containerized environment package: **sudo yum install docker**.

For other operating systems, refer to the following Docker documentation:

- [Installing the Docker containerized environment on Linux distributions](#)
- [Installing the Docker containerized environment on Mac](#)
- [Installing the Docker containerized environment on Windows](#)

4.3. STEP 2: RUN THE DOCKER CONTAINERIZED ENVIRONMENT GATEWAY

1. Start the Docker daemon: **sudo systemctl start docker.service**.
2. Check if the Docker daemon is running: **sudo systemctl status docker.service**. You can download a ready to use Docker-formatted container image from the Red Hat registry: **sudo docker pull registry.access.redhat.com/3scale-amp21/apicast-gateway:1.4**.
3. Run APIcast in a Docker-formatted container: **sudo docker run --name apicast --rm -p 8080:8080 -e THREESCALE_PORTAL_ENDPOINT=https://<access_token>@<domain>-admin.3scale.net registry.access.redhat.com/3scale-amp21/apicast-gateway:1.4**.

Here, **<access_token>** is the [Access Token](#) for the 3scale Account Management API. You can use the [Provider Key](#) instead of the access token. **<domain>-admin.3scale.net** is the URL of your 3scale admin portal.

This command runs a Docker-formatted container called "apicast" on port **8080** and fetches the JSON configuration file from your 3scale portal. For other configuration options, see the [APIcast Overview](#) guide.

4.3.1. The Docker command options

You can use the following options with the **docker run** command:

- **--rm**: Automatically removes the container when it exits.
- **-d** or **--detach**: Runs the container in the background and prints the container ID. When it is not specified, the container runs in the foreground mode and you can stop it using **CTRL + c**. When started in the detached mode, you can reattach to the container with the **docker attach** command, for example, **docker attach apicast**.
- **-p** or **--publish**: Publishes a container's port to the host. The value should have the format **<host port="">>:<container port=">**, so **-p 80:8080** will bind port **8080** of the container to port **80** of the host machine. For example, the [Management API](#) uses port **8090**, so you may want to publish this port by adding **-p 8090:8090** to the **docker run** command.
- **-e** or **--env**: Sets environment variables.
- **-v** or **--volume**: Mounts a volume. The value is typically represented as **<host path="">>:<container path=">[:<options>]**. **<options>** is an optional attribute; you can set it to **:ro** to specify that the volume will be read only (by default, it is mounted in read-write mode). Example: **-v /host/path:/container/path:ro**.

For more information on available options, see [Docker run reference](#).

4.4. STEP 3: TESTING APICAST

The preceding steps ensure that your Docker-formatted container is running with your own configuration file and the Docker-formatted image from the 3scale registry. You can test calls through APIcast on port **8080** and provide the correct authentication credentials, which you can get from your 3scale account.

Test calls will not only verify that APIcast is running correctly but also that authentication and reporting is being handled successfully.



NOTE

Ensure that the host you use for the calls is the same as the one configured in the **Public Base URL** field on the **Integration** page.

4.5. STEP 4: TROUBLESHOOTING APICAST ON THE DOCKER CONTAINERIZED ENVIRONMENT

4.5.1. Cannot connect to the Docker daemon error

The **docker: Cannot connect to the Docker daemon. Is the docker daemon running on this host?** error message may be because the Docker service hasn't started. You can check the status of the Docker daemon by running the **sudo systemctl status docker.service** command.

Ensure that you are run this command as the **root** user because the Docker containerized environment requires root permissions in RHEL by default. For more information, see [here](#)).

4.5.2. Basic Docker command-line interface commands

If you started the container in the detached mode (**-d** option) and want to check the logs for the running APIcast instance, you can use the **log** command: **sudo docker logs <container>**. Where **<container>** is the container name ("*apicast*" in the example above) or the container ID. You can get a list of the running containers and their IDs and names by using the **sudo docker ps** command.

To stop the container, run the **sudo docker stop <container>** command. You can also remove the container by running the **sudo docker rm <container>** command.

For more information on available commands, see [Docker commands reference](#).

4.6. STEP 5: CUSTOMISING THE GATEWAY

There are some customizations that cannot be managed through the admin portal and require writing custom logic to APIcast itself.

4.6.1. Custom Lua logic

The easiest way to customise APIcast logic is to rewrite the existing file with your own and attach it as a volume.

```
-v $(pwd)/path_to/file.lua:/opt/app-root/src/src/file.lua:ro
```

The **-v** flag attaches the local file to the stated path in the Docker-formatted container.

4.6.2. Customising the configuration files

The same steps apply to custom configuration files as the Lua scripting. If you want to add to the existing conf files rather than overwriting, ensure the name of your new file does not clash with pre-existing ones. This will automatically be included in the main *nginx.conf*.

To edit the *config.json* file, you can fetch this file from your admin portal with the following URL: <https://<account>-admin.3scale.net/admin/api/nginx/spec.json> and copy and paste the contents locally. You can pass this local JSON file with the following command to start APIcast:

```
docker run --name apicast --rm -p 8080:8080 -v  
$(pwd)/path_to/config.json:/opt/app/config.json:ro -e  
THREESCALE_CONFIG_FILE=/opt/app/config.json
```

You should now be able to run APIcast on the Docker containerized environment. For other deployment options, check out the related articles.

CHAPTER 5. RUNNING APICAST ON RED HAT OPENSIFT

This tutorial describes how to deploy the APICast API Gateway on Red Hat OpenShift.

5.1. PREREQUISITES

To follow the tutorial steps below, you will first need to configure APICast in your 3scale Admin Portal as per the [APICast Overview](#). Make sure *Self-managed Gateway* is selected as the deployment option in the integration settings. You should have both Staging and Production environment configured to proceed.

5.2. STEP 1: SET UP OPENSIFT

If you already have a running OpenShift cluster, you can skip this step. Otherwise, continue reading.

For production deployments you can follow the [instructions for OpenShift installation](#). In order to get started quickly in development environments, there are a couple of ways you can install OpenShift:

- Using **oc cluster up** command – https://github.com/openshift/origin/blob/master/docs/cluster_up_down.md (used in this tutorial, with detailed instructions for Mac and Windows in addition to Linux which we cover here)
- All-In-One Virtual Machine using Vagrant – <https://www.openshift.org/vm>

In this tutorial the OpenShift cluster will be installed using:

- Red Hat Enterprise Linux (RHEL) 7
- Docker containerized environment v1.10.3
- OpenShift Origin command line interface (CLI) - v1.3.1

5.2.1. Install the Docker containerized environment

Docker-formatted container images provided by Red Hat are released as part of the Extras channel in RHEL. To enable additional repositories, you can use either the [Subscription Manager](#), or yum config manager. See the [RHEL product documentation](#) for details.

For a RHEL 7 deployed on a AWS EC2 instance we'll use the following the instructions:

1. List all repositories:

```
sudo yum repolist all
```

Find the ***-extras** repository.

1. Enable *extras* repository:

```
sudo yum-config-manager --enable rhui-REGION-rhel-server-extras
```

2. Install Docker-formatted container images:

```
sudo yum install docker docker-registry
```

3. Add an insecure registry of **172.30.0.0/16** by adding or uncommenting the following line in **/etc/sysconfig/docker** file:

```
INSECURE_REGISTRY='--insecure-registry 172.30.0.0/16'
```

4. Start the Docker containerized environment:

```
sudo systemctl start docker
```

You can verify that the Docker containerized environment is running with the command:

```
sudo systemctl status docker
```

5.2.2. Start OpenShift cluster

Download the latest stable release of the client tools (**openshift-origin-client-tools-VERSION-linux-64bit.tar.gz**) from [OpenShift releases page](#), and place the Linux **oc** binary extracted from the archive in your **PATH**.



NOTE

- Please be aware that the **oc cluster** set of commands are only available in the 1.3+ or newer releases.
- the docker command runs as the **root** user, so you will need to run any **oc** or docker commands with root privileges.

Open a terminal with a user that has permission to run docker commands and run:

```
oc cluster up
```

At the bottom of the output you will find information about the deployed cluster:

```
-- Server Information ...
OpenShift server started.
The server is accessible via web console at:
https://172.30.0.112:8443

You are logged in as:
User: developer
Password: developer

To login as administrator:
oc login -u system:admin
```

Note the IP address that is assigned to your OpenShift server, we will refer to it in the tutorial as **OPENSIFT-SERVER-IP**.

5.2.2.1. Setting up OpenShift cluster on a remote server

In case you are deploying the OpenShift cluster on a remote server, you will need to explicitly specify a public hostname and a routing suffix on starting the cluster, in order to be able to access the OpenShift web console remotely.

For example, if you are deploying on an AWS EC2 instance, you should specify the following options:

```
oc cluster up --public-hostname=ec2-54-321-67-89.compute-1.amazonaws.com --routing-suffix=54.321.67.89.xip.io
```

where **ec2-54-321-67-89.compute-1.amazonaws.com** is the Public Domain, and **54.321.67.89** is the IP of the instance. You will then be able to access the OpenShift web console at <https://ec2-54-321-67-89.compute-1.amazonaws.com:8443>.

5.3. STEP 2: DEPLOY APICAST USING THE OPENSIFT TEMPLATE

1. By default you are logged in as *developer* and can proceed to the next step. Otherwise login into OpenShift using the **oc login** command from the OpenShift Client tools you downloaded and installed in the previous step. The default login credentials are *username = "developer"* and *password = "developer"*:

```
oc login https://OPENSIFT-SERVER-IP:8443
```

You should see **Login successful.** in the output.

2. Create your project. This example sets the display name as *gateway*

```
oc new-project "3scalegateway" --display-name="gateway" --description="3scale gateway demo"
```

The response should look like this:

```
Now using project "3scalegateway" on server "https://172.30.0.112:8443".
```

Ignore the suggested next steps in the text output at the command prompt and proceed to the next step below.

3. Create a new Secret to reference your project by replacing **<access_token>** and **<domain>** with yours.

```
oc secret new-basicauth apicast-configuration-url-secret --password=https://<access_token>@<domain>-admin.3scale.net
```

Here **<access_token>** is an [Access Token](#) (not a Service Token) for the 3scale Account Management API, and **<domain>-admin.3scale.net** is the URL of your 3scale Admin Portal.

The response should look like this:

```
secret/apicast-configuration-url-secret
```

4. Create an application for your APIcast Gateway from the template, and start the deployment:

```
oc new-app -f https://raw.githubusercontent.com/3scale/3scale-amp-openshift-templates/2.1.0-GA/apicast-gateway/apicast.yml
```

You should see the following messages at the bottom of the output:


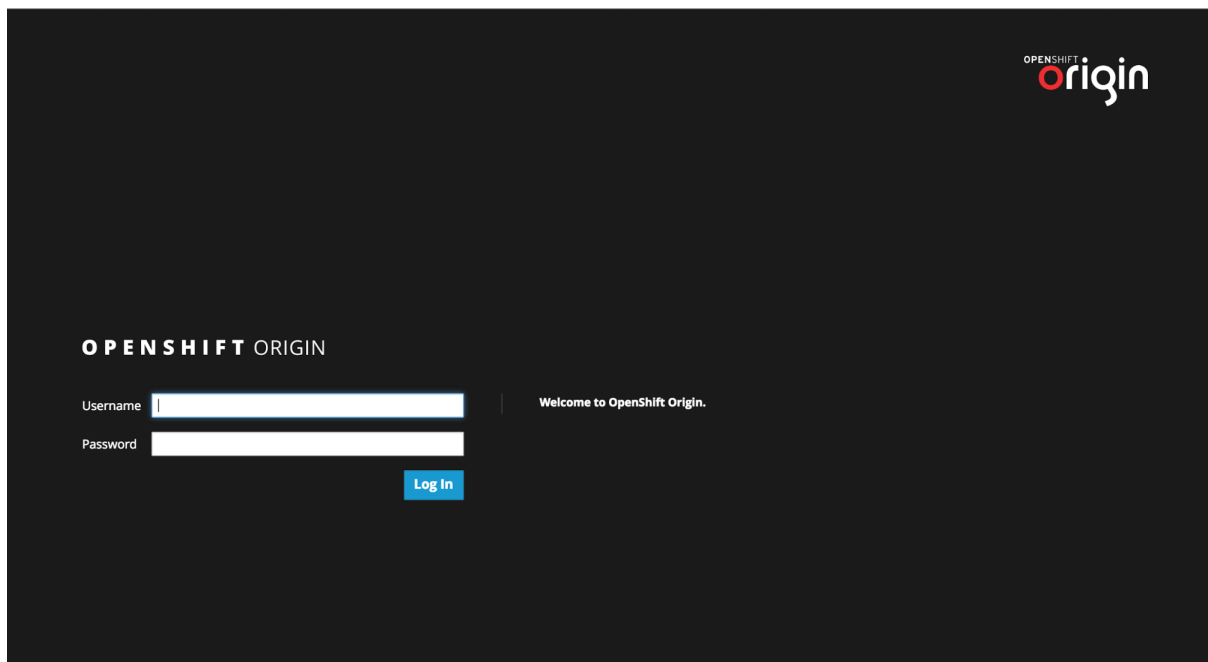
```
--> Creating resources with label app=3scale-gateway ...
    deploymentconfig "apicast" created
    service "apicast" created
--> Success
    Run 'oc status' to view your app.
```

5.4. STEP 3: CREATE ROUTES IN OPENSIFT CONSOLE

1. Open the web console for your OpenShift cluster in your browser: <https://OPENSIFT-SERVER-IP:8443/console/>

Use the value specified in **--public-hostname** instead of **OPENSIFT-SERVER-IP** if you started OpenShift cluster on a remote server.

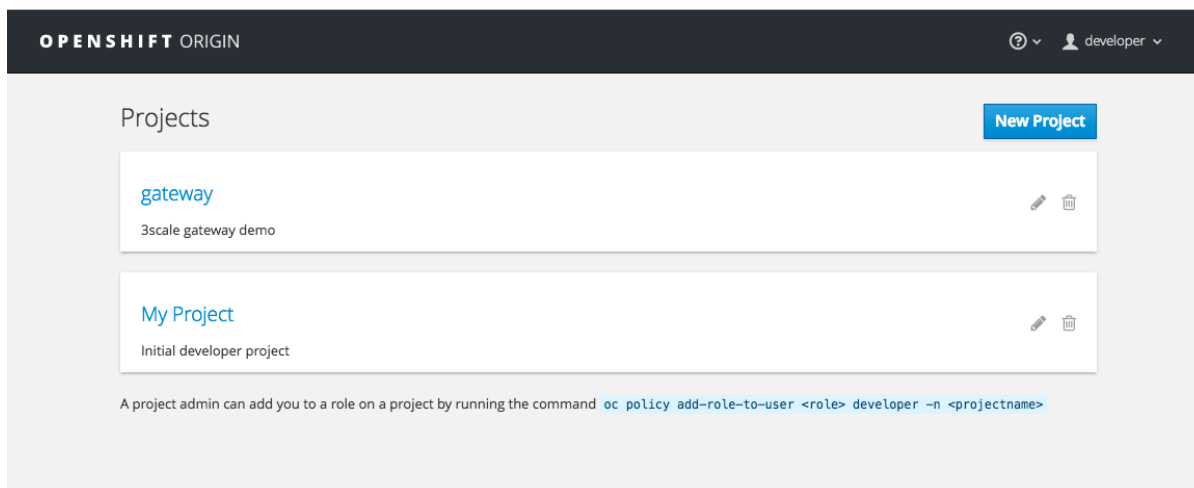
You should see the login screen:



NOTE

You may receive a warning about an untrusted web-site. This is expected, as we are trying to access the web console through secure protocol, without having configured a valid certificate. While you should avoid this in production environment, for this test setup you can go ahead and create an exception for this address.

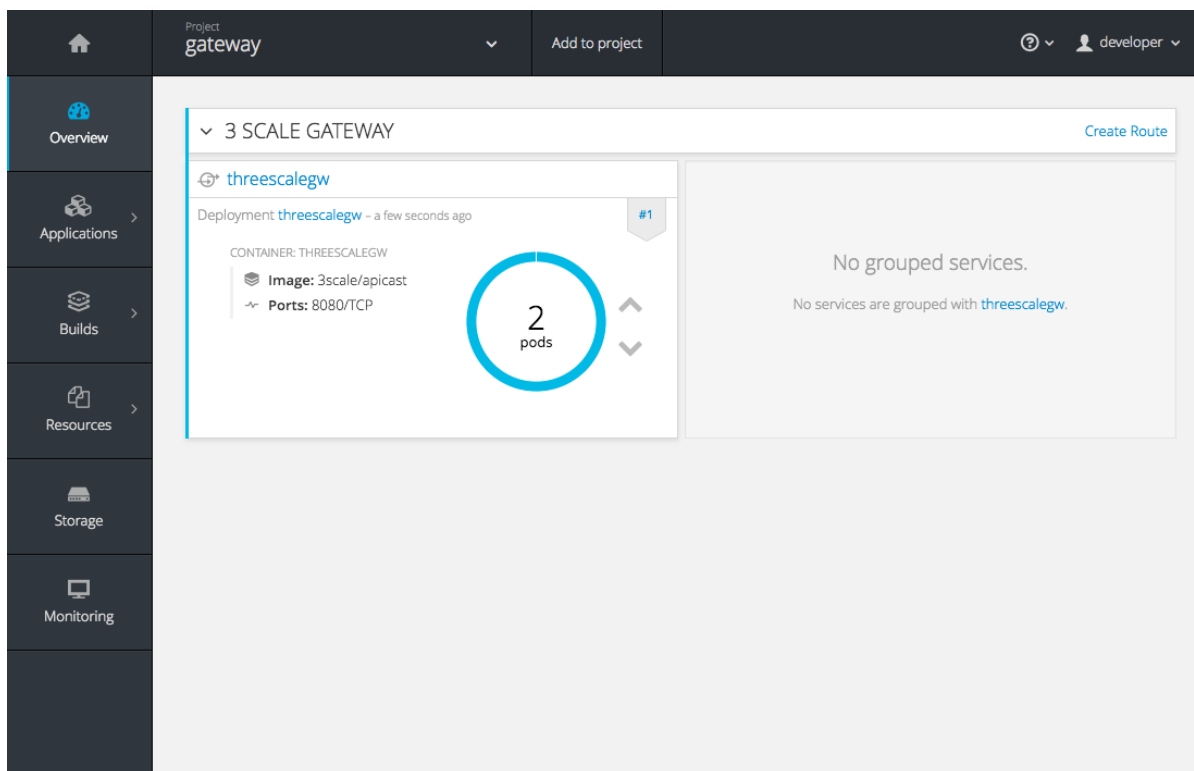
2. Log in using the *developer* credentials created or obtained in the *Setup OpenShift* section above. You will see a list of projects, including the "gateway" project you created from the command line above.



If you do not see your gateway project, you probably created it with a different user and will need to assign the policy role to this user.

3. Click on "gateway" and you will see the *Overview* tab. OpenShift downloaded the code for APIcast and started the deployment. You may see the message *Deployment #1 running* when the deployment is in progress.


When the build completes, the UI will refresh and show two instances of APIcast (2 pods) that have been started by OpenShift, as defined in the template.




Each APIcast instance, upon starting, downloads the required configuration from 3scale using the settings you provided on the **Integration** page of your 3scale Admin Portal.

OpenShift will maintain two APIcast instances and monitor the health of both; any unhealthy APIcast instance will automatically be replaced with a new one.

4. In order to allow your APIcast instances to receive traffic, you'll need to create a route. Start by clicking on **Create Route**.



Enter the same host you set in 3scale above in the section **Public Base URL** (without the `http://` and without the port) , e.g. **gateway.openshift.demo**, then click the **Create** button.



Create a new route for every 3scale service you define. Alternatively, you could avoid having to create a new route for every 3scale service you define by deploying a wildcard router.

CHAPTER 6. ADVANCED APICAST CONFIGURATION

This section covers the advanced settings option of 3scale's API gateway in the staging environment.

6.1. DEFINE A SECRET TOKEN

For security reasons, any request from 3scale's gateway to your API backend will contain a header called **X-3scale-proxy-secret-token**. You can set the value of this header in **Authentication Settings** on the Integration page.

▼ AUTHENTICATION SETTINGS

Host Header

Lets you define a custom `Host` request header. This is needed if your API backend only accepts traffic from a specific host.

Secret Token

Enables you to block any direct developer requests to your API backend; each 3scale API gateway call to your API backend contains a request header called `x-3scale-proxy-secret-token`. The value of this header can be set by you here. It's up to you ensure your backend only allows calls with this secret header.

Setting the secret token will act as a shared secret between the proxy and your API so that you can block all API requests that do not come from the gateway if you so wish. This gives an extra layer of security to protect your public endpoint while you're in the process of setting up your traffic management policies with the sandbox gateway.

Your API backend must have a public resolvable domain for the gateway to work, so anyone who might know your API backend could bypass the credentials checking. This should not be a problem because the API gateway in the staging environment is not meant for production use, but it's always better to have a fence available.

6.2. CREDENTIALS

The API credentials within 3scale are always **user_key** or **app_id/app_key** depending on the authentication mode you're using (OAuth is not available for the API gateway in the staging environment). However, you might want to use different credential names in your API. In this case, you'll need to set custom names for the **user_key** if you're using API key mode:

Auth user key

or for the **app_id** and **app_key**:

App ID parameter

Name of the parameter that acts of behalf of app id

App Key parameter

Name of the parameter that acts of behalf of app key

For instance, you could rename **app_id** to **key** if that fits your API better. The gateway will take the name **key** and convert it to **app_id** before doing the authorization call to 3scale's backend. Note that the new credential name has to be alphanumeric.

You can decide whether your API passes credentials in the query string (or body if not a GET) or in the headers.

**CREDENTIALS
LOCATION***

As HTTP Headers

As query parameters (GET) or body parameters (POST/PUT/DELETE)

6.3. ERROR MESSAGES

Another important element for a full-fledged configuration is to define your own custom error messages.

It's important to note that 3scale's API gateway in the staging environment will do a pass through of any error message generated by your API. However, since the management layer of your API is now carried out by the gateway, there are some errors that your API will never see since some requests will be terminated by the gateway.

AUTHENTICATION FAILED ERROR

Response Code*	403
Content-type	text/plain; charset=us-ascii
Response Body	Authentication failed

AUTHENTICATION MISSING ERROR

Response Code*	403
Content-type	text/plain; charset=us-ascii
Response Body	Authentication parameters missing

NO MATCH ERROR

Response Code*	404
Content-type	text/plain; charset=us-ascii
Response Body	No Mapping Rule matched

These errors are the following:

- Authentication missing: this error will be generated whenever an API request does not contain any credentials. This occurs when users forget to add their credentials to an API request.
- Authentication failed: this error will be generated whenever an API request does not contain valid credentials. This can be because the credentials are fake or because the application has been temporarily suspended.
- No match: this error means that the request did not match any mapping rule, therefore no metric is updated. This is not necessary an error, but it means that either the user is trying random paths or that your mapping rules do not cover legitimate cases.

6.4. CONFIGURATION HISTORY

Every time you click on the **Update & Test Staging Configuration** button, the current configuration will be saved in a JSON file. The staging gateway will pull the latest configuration with each new request. For each environment, staging or production, you can see a history of all the previous configuration files.

Note that it is not possible to automatically roll back to previous versions. Instead we provide a history of all your configuration versions with their associated JSON files. These files can be used to check what configuration you had deployed at any moment on time. If you want to, you can recreate any deployments manually.

6.5. DEBUGGING

Setting up the gateway configuration is easy, but still some errors can occur on the way. For those cases the gateway can return some useful debug information that will be helpful to track down what is going on.

To enable the debug mode on 3scale's API gateway in the staging environment you can add the following header with your provider key to a request to your gateway: **X-3scale-debug: YOUR_PROVIDER_KEY**

When the header is found and the provider key is valid, the gateway will add the following information to the response headers:

```
X-3scale-matched-rules: /v1/word/{word}.json, /v1
X-3scale-credentials: app_key=APP_KEY&app_id=APP_ID
X-3scale-usage: usage[version_1]=1&usage[word]=1
```

Basically, **X-3scale-matched-rules** tells you which mapping rules have been activated by the request. Note that it is a list. The header **X-3scale-credentials** returns the credentials that have been passed to 3scale's backend. Finally **X-3scale-usage** tells you the usage that will be reported to 3scale's backend.

You can check the logic for your mapping rules and usage reporting in the Lua file, in the function **extract_usage_x()** where **x** is your **service_id**.

```
...
local args = get_auth_params(nil, method)
local m = ngx.re.match(path,[=^/v1/word/([w_\.]+)\.json=])
if (m and method == "GET") then
  -- rule: /v1/word/{word}.json --
  params["word"] = m[1]
  table.insert(matched_rules, "/v1/word/{word}.json")
  usage_t["word"] = set_or_inc(usage_t, "word", 1)
  found = true
end
...
```

In this example, the comment **-- rule: /v1/word/{word}.json --** shows which particular rule the Lua code refers to. Each rule has a Lua snippet like the one above. Comments are delimited by **--**, **--[[,]]** in Lua, and with **#** in NGINX.

Unfortunately, there is no automatic rollback for Lua files if you make any changes. However, if your current configuration is not working but the previous one was OK, you can download the previous configuration files from the deployment history.

Staging – configure & test your integration [visit documentation](#) deployed [deployment history](#)

API ?

Private Base URL*

Private address of your API that will be called by the API gateway.

6.6. EXTENDING THE GATEWAY

3scale's hosted APIcast is quite flexible, but there are sometimes things that cannot be done because the console interface doesn't allow it or because of security reasons due to a multi-tenant gateway.

If you need to extend your API gateway, you can always run it locally on your servers ([see the APIcast self-managed section to deploy your own API gateway](#)) and point it to the proper service and configuration from your 3scale account.

When you're running the gateway on your servers, you can customize it to enable any feature you might need – NGINX with Lua is an extremely powerful, open-source piece of technology.

We have written a blog post explaining [how to augment APIs with NGINX and Lua](#) . Here are some examples that can be done:

- Basic DoS protection: white-lists, black-lists, rate-limiting at the second level.
- Define arbitrarily complex mapping rules.
- API rewrite rules – for example, you might want API requests starting with `/v1/*` to be rewritten to `/versions/1/*` when they hit your API backend.
- Content filtering: you can add checks on the content of the requests, either for security or to filter out undesired side effects.
- Content rewrites: you can transform the results of your API.
- Many, many more. Combining the power and flexibility of NGINX with Lua scripting is an awesome combination.

CHAPTER 7. CODE LIBRARIES

3scale plugins allow you to connect to the 3scale architecture in a variety of core programming languages. Plugins can be deployed anywhere to act as control agents for your API traffic. Download the latest version of the plugin and gems for the programming language of your choice along with their corresponding documentation below.

You can also find us on [GitHub](#). For other technical details, refer to the [technical overview](#).

7.1. PYTHON

Download the code [here](#).

Standard distutils installation – unpack the file and from the new directory run:

```
sudo python setup.py install
```

Or you can put ThreeScale.py in the same directory as your program. See the README for some usage examples.

7.2. RUBY GEM

You can download from [GitHub](#) or install it directly using:

```
gem install 3scale_client
```

7.3. PERL

Download the code bundle [here](#).

It requires:

- » LWP::UserAgent (installed by default on most systems)
- » XML::Parser (now standard in perl >= 5.6)

Run this to install into the default library location:

```
perl Makefile.PL && make install
```

Run this for further install options:

```
make help
```

7.4. PHP

Download the code [here](#) and drop into your code directory – see the README in the bundle.

7.5. JAVA

Download the JAR bundle [here](#).

Just unpack it into a directory and then run the maven build.

7.6. .NET

Download the code [here](#), and see the bundled README.txt for instructions.

7.7. NODE.JS

Download the code [here](#), or install it through the package manager npm:

```
npm install 3scale
```

CHAPTER 8. PLUGIN SETUP

By the time you complete this tutorial, you'll have configured your API to use the available 3scale code plugins to manage access traffic.

There are several different ways to add 3scale API Management to your API – including using [APIcast \(API gateway\)](#), [Amazon API Gateway with Lambda](#), [3scale Service Management API](#), or code plugins. This tutorial drills down into how to use the code plugin method to get you set up.

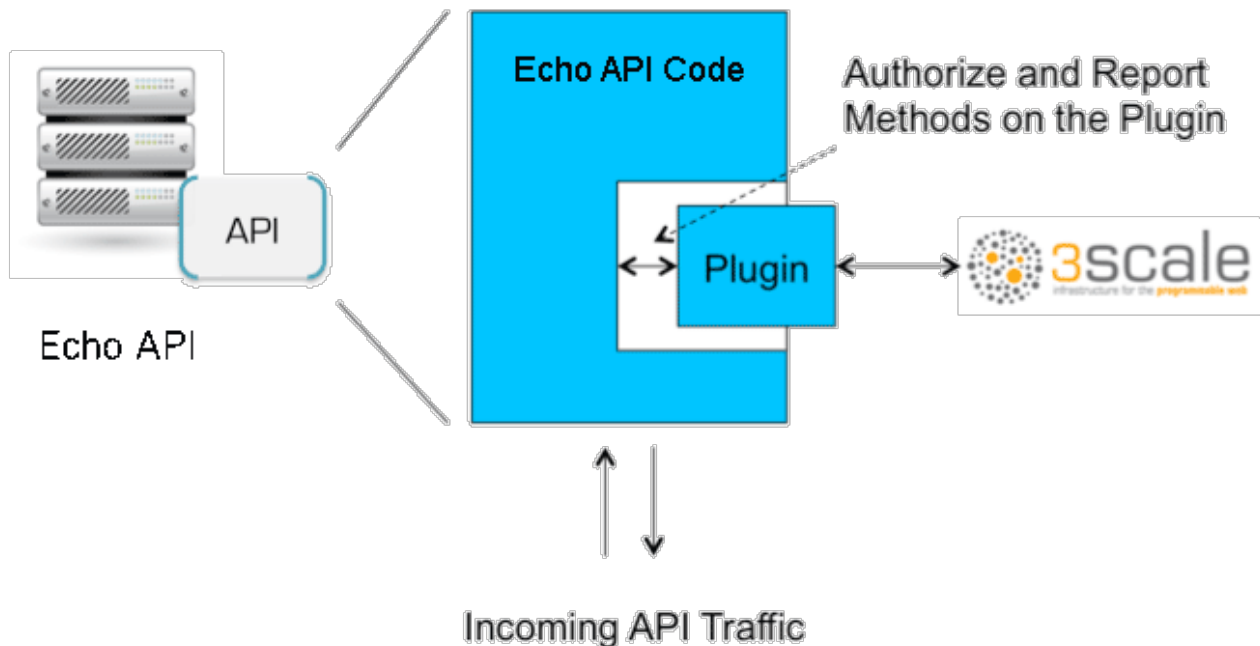
8.1. HOW PLUGINS WORK

3scale API plugins are available for a variety of implementation languages including Java, Ruby, PHP, .NET, and others – the full list can be found in the [code libraries section](#). The plugins provide a wrapper for the 3scale Service Management API. Find the Service Management API in the 3scale ActiveDocs, available in your Admin Portal, under the **Documentation → 3scale API Docs** section, to enable:

- API access control and security
- API traffic reporting
- API monetization

This wrapper connects back into the 3scale system to set and manage policies, keys, rate limits, and other controls that you can put in place through the interface. Check out the [Getting Started](#) guide to see how to configure these elements.

Plugins are deployed within your API code to insert a traffic filter on all calls as shown in the figure.



8.2. STEP 1: SELECT YOUR LANGUAGE AND DOWNLOAD THE PLUGIN

Once you've created your 3scale account ([signup here](#)), navigate to the [code libraries](#) section on this site and choose a plugin in the language you plan to work with. Click through to the code repository to get the bundle in the form that you need.

If your language is not supported or listed, [let us know](#) and we'll inform you of any ongoing support efforts for your language. Alternatively, you can connect directly to the [3scale Service Management API](#).

8.3. STEP 2: CHECK THE METRICS YOU'VE SET FOR YOUR API

As described in the [API definition tutorial](#), you can configure multiple metrics and methods for your API on the API control panel. Each metric and method has a system name that will be required when configuring your plugin. You can find the metrics in **API > Definition** area of your Admin Portal.

The screenshot shows the 3scale Admin Portal interface for API Definition. The top navigation bar includes Dashboard, Developers, Applications, Billing, Analytics, API (highlighted), Developer Portal, and Settings. The left sidebar has Overview, Active Docs, Definition (highlighted), Integration, Application Plans, Settings, and Alerts. The main content area is titled 'Definition' and shows the API Name: API and System Name: api. Below this, there are sections for 'Methods' and 'Metrics'. The 'Methods' section includes a table with columns: Method, System Name, Unit, Description, and Mapped. The 'Metrics' section includes a table with columns: Metric, System Name, Unit, Description, and Mapped. Red arrows point to 'Create new method' and 'Create new metric' buttons.

Methods

Method	System Name	Unit	Description	Mapped	
transactions/create_single	transactions/create_single	hit		✓	
transactions/create_multiple	transactions/create_multiple	hit		✓	
transactions/confirm	transactions/confirm	hit		✓	
transactions/destroy	transactions/destroy	hit		Add a mapping rule	

Metrics

Metric	System Name	Unit	Description	Mapped	
Hits	hits	hit	Number of API hits	✓	
Number of transactions	transactions	transaction		Add a mapping rule	

For more advanced information on metrics, methods, and rate limits see the specific [tutorial on rate limits](#).

8.4. STEP 3: INSTALL THE PLUGIN CODE TO YOUR APPLICATION

Armed with this information, return to the code and add the downloaded code bundle to your application. This step varies for each type of plugin, and the form that it takes depends on the way each language framework uses libraries. For the purposes of this tutorial, this example will proceed with Ruby plugin instructions. Other plugins integration details are included in the README documentation of each repository.

This library is distributed as a gem, for which Ruby 2.1 or JRuby 9.1.1.0 are minimum requirements:

```
gem install 3scale_client
```

Or alternatively, download the source code from GitHub.

If you are using Bundler, please add this to your Gemfile:

```
gem '3scale_client'
```

and do a bundle install.

If you are using Rails' config.gems, put this into your config/environment.rb

```
config.gem '3scale_client'
```

Otherwise, require the gem in whatever way is natural to your framework of choice.

Then, create an instance of the client:

```
client = ThreeScale::Client.new(service_tokens: true)
```



NOTE

unless you specify `service_tokens: true` you will be expected to specify a `provider_key` parameter, which is deprecated in favor of Service Tokens:

```
client = ThreeScale::Client.new(provider_key: 'your_provider_key')
```

This will communicate with the 3scale platform SaaS default server.

If you want to create a Client with a given host and port when connecting to an on-premise instance of the 3scale platform, you can specify them when creating the instance:

```
client = ThreeScale::Client.new(service_tokens: true, host: 'service_management_api.example.com',  
port: 80)
```

or

```
client = ThreeScale::Client.new(provider_key: 'your_provider_key', host:  
'service_management_api.example.com', port: 80)
```

Because the object is stateless, you can create just one and store it globally. Then you can perform calls in the client:

```
client.authorize(service_token: 'token', service_id: '123', usage: usage)  
client.report(service_token: 'token', service_id: '123', usage: usage)
```

If you had configured a (deprecated) provider key, you would instead use:

```
client.authrep(service_id: '123', usage: usage)
```



NOTE

`service_id` is mandatory since November 2016, both when using service tokens and when using provider keys



NOTE

You might use the option `warn_deprecated: false` to avoid deprecation warnings. This is enabled by default.

SSL and Persistence

Starting with version 2.4.0 you can use two more options: **secure** and **persistent** like:

```
client = ThreeScale::Client.new(provider_key: '...', secure: true, persistent: true)
```

secure

Enabling secure will force all traffic going through HTTPS. Because establishing SSL/TLS for every call is expensive, there is persistent.

persistent

Enabling persistent will use HTTP Keep-Alive to keep open connection to our servers. This option requires installing gem net-http-persistent.

8.5. STEP 4: ADD CALLS TO AUTHORIZE AS API TRAFFIC ARRIVES

The plugin supports the 3 main calls to the 3scale backend:

- **authrep** grants access to your API and reports the traffic on it in one call.
- **authorize** grants access to your API.
- **report** reports traffic on your API.

8.5.1. Usage on authrep mode

You can make request to this backend operation using **service_token** and **service_id**, and an authentication pattern like **user_key**, or **app_id** with an optional key, like this:

```
response = client.authrep(service_token: 'token', service_id: 'service_id', app_id: 'app_id', app_key: 'app_key')
```

Then call the **success?** method on the returned object to see if the authorization was successful.

```
if response.success?
  # All fine, the usage will be reported automatically. Proceed.
else
  # Something's wrong with this application.
end
```

The example is using the **app_id** authentication pattern, but you can also use other patterns such as **user_key**.

Example:

```
class ApplicationController < ActionController
  # Call the authenticate method on each request to the API
  before_filter :authenticate

  # You only need to instantiate a new Client once and store it as a global variable
  # If you used a provider key it is advisable to fetch it from the environment, as
  # it is secret.
  def create_client
    @@threescale_client ||= ThreeScale::Client.new(service_tokens: true)
  end
end
```

```

# To record usage, create a new metric in your application plan. You will use the
# "system name" that you specified on the metric/method to pass in as the key to the usage hash.
# The key needs to be a symbol.
# A way to pass the metric is to add a parameter that will pass the name of the metric/method along
#
# Note that you don't always want to retrieve the service token and service id from
# the parameters - this will depend on your application.
def authenticate
  response = create_client.authrep(service_token: params['service_token']
                                service_id: params['service_id'],
                                app_id: params['app_id'],
                                app_key: params['app_key'],
                                usage: { params['metric'].to_sym => 1 })


  if response.success?
    return true
    # All fine, the usage will be reported automatically. Proceed.
  else
    # Something's wrong with this application.
    puts "#{response.error_message}"
    # raise error
  end
end
end
end

```

8.6. STEP 5: DEPLOY AND RUN TRAFFIC

Once the required calls are added to your code, you can deploy (ideally to a staging/testing environment) and make API calls to your endpoints. As the traffic reaches the API, it will be filtered by the plugin and keys checked against issues API credentials. Refer to the [Getting Started guide](#) for how to generate valid keys – a set will have also been created as sample data in your 3scale account.

To see if traffic is flowing, log in to your API Admin Portal and navigate to the Analytics tab – there you will see traffic reported via the plugin.



Once the application is making calls to the API, they will become visible on the statistics dashboard and the Statistics > Usage page for more detailed view.

If you're receiving errors in the connection, you can also check the "Errors" menu item under Monitoring.