



Red Hat 3scale API Management 2.5

Operating 3scale

How to automate the deployment of your API Gateway using Chef, production tips,
and more.

Red Hat 3scale API Management 2.5 Operating 3scale

How to automate the deployment of your API Gateway using Chef, production tips, and more.

Legal Notice

Copyright © 2020 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, the Red Hat 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 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 development operations with Red Hat 3scale API Management 2.5.

Table of Contents

CHAPTER 1. 3SCALE OPERATIONS AND SCALING GUIDE	4
1.1. INTRODUCTION	4
1.1.1. Prerequisites	4
1.1.2. Further reading	4
1.2. RE-DEPLOYING APICAST	4
1.3. APICAST BUILT-IN WILDCARD ROUTING	4
1.3.1. Modify wildcards	5
1.4. SCALING UP 3SCALE ON-PREMISE	5
1.4.1. Scaling up storage	5
1.4.1.1. Method 1: Backup and swap persistent volumes	5
1.4.1.2. Method 2: Back up and redeploy 3scale	6
1.4.2. Scaling up performance	6
1.4.2.1. Configuring 3scale on-premise deployments	6
1.4.2.2. Vertical and horizontal hardware scaling	7
1.4.2.3. Scaling up routers	7
1.4.2.4. Further reading	7
1.5. OPERATIONS TROUBLESHOOTING	7
1.5.1. Accessing your logs	8
1.5.2. Checking job queues	8
1.5.3. Preventing monotonic growth	8
CHAPTER 2. USING THE 3SCALE TOOLBOX	10
2.1. INSTALLING THE TOOLBOX	10
2.2. IMPORTING SERVICES	10
2.3. COPYING SERVICES	11
2.4. COPYING SERVICE SETTINGS ONLY	11
2.5. IMPORTING OPENAPI DEFINITIONS	12
2.5.1. Optional flags	12
2.6. MANAGING REMOTE WEB ADDRESSES	12
2.6.1. Listing remote web addresses	13
2.6.2. Adding remote web addresses	13
2.6.3. Removing remote web addresses	13
2.6.4. Renaming remote web addresses	13
2.7. TROUBLESHOOTING SSL ISSUES	13
CHAPTER 3. PROVISIONING 3SCALE SERVICES AND CONFIGURATIONS VIA THE OPERATOR (CAPABILITIES)	15
3.1. INTRODUCTION	15
3.1.1. Prerequisites	15
3.2. DEPLOYING CAPABILITIES RELATED CUSTOM RESOURCES	15
3.3. DEPLOYING OPTIONAL TENANTS CUSTOM RESOURCE	18
3.4. DELETING CREATED CUSTOM RESOURCES	19
CHAPTER 4. TROUBLESHOOTING	20
4.1. COMMON ISSUES	20
4.1.1. Integration issues	20
4.1.1.1. APICast Hosted	20
4.1.1.2. APICast self-managed	21
4.1.2. Production issues	22
4.1.2.1. Availability issues	22
4.1.3. Post-deploy issues	23
4.2. TROUBLESHOOTING 101	24

4.2.1.1. Can we connect?	24
4.2.2. 2. Is it me or is it them?	25
4.2.3. 3. Is it a DNS issue?	25
4.2.4. 4. Is it an SSL issue?	25
4.3. TROUBLESHOOTING CHECKLISTS	27
4.3.1. API	28
4.3.2. API Gateway > API	28
4.3.3. API gateway	28
4.3.3.1. 1. Is the API gateway up and running?	28
4.3.3.2. 2. Are there any errors in the gateway logs?	28
4.3.4. API gateway > 3scale	28
4.3.4.1. 1. Can the API gateway reach 3scale?	29
4.3.4.2. 2. Is the API gateway resolving 3scale addresses correctly?	29
4.3.4.3. 3. Is the API gateway calling 3scale correctly?	30
4.3.5. 3scale	31
4.3.5.1. 1. Is 3scale available?	31
4.3.5.2. 2. Is 3scale returning an error?	31
4.3.5.3. 3. Use the 3scale debug headers	31
4.3.5.4. 4. Check the integration errors	31
4.3.6. Client API gateway	31
4.3.6.1. 1. Is the API gateway reachable from the public internet?	32
4.3.6.2. 2. Is the API gateway reachable by the client?	32
4.3.7. Client	32
4.3.7.1. 1. Test the same call using a different client	32
4.3.7.2. 2. Inspect the traffic sent by client	32
4.4. OTHER ISSUES	32
4.4.1. ActiveDocs issues	32
4.4.1.1. 1. Use petstore.swagger.io	32
4.4.1.2. 2. Check that firewall allows connections from ActiveDocs proxy	32
4.4.1.3. 3. Call the API with incorrect credentials	32
4.4.1.4. 4. Compare calls	33
4.5. APPENDIX	33
4.5.1. Logging in NGINX	33
4.5.1.1. Enabling debugging log	33
4.5.2. 3scale error codes	33

CHAPTER 1. 3SCALE OPERATIONS AND SCALING GUIDE

1.1. INTRODUCTION

This document describes operations and scaling tasks of a Red Hat 3scale AMP 2.5 On-Premises installation.

1.1.1. Prerequisites

An installed and initially configured AMP On-Premises instance on a [supported OpenShift version](#).

This document is not intended for local installations on laptops or similar end user equipment.

1.1.2. Further reading

- [OpenShift Documentation](#)
- [Health and Liveness Monitoring](#)

1.2. RE-DEPLOYING APICAST

After you have deployed AMP On-Premises and your chosen APICast deployment method, you can test and promote system changes through 3scale Admin Portal. By default, APICast deployments on OpenShift, both built-in and on other OpenShift clusters, are configured to allow you to publish changes to your staging and production gateways through the AMP UI.

Redeploy APICast on OpenShift:

1. Make system changes.
2. In the UI, deploy to staging and test.
3. In the UI, promote to production.

By default, APICast retrieves and publishes the promoted update once every 5 minutes.

If you are using APICast on the Docker containerized environment or a native installation, you must configure your staging and production gateways, and configure how often your gateway retrieves published changes. After you have configured your APICast gateways, you can redeploy APICast through the AMP UI.

To redeploy APICast on the Docker containerized environment or a native installations:

1. Configure your APICast gateway and connect it to AMP On-Premises.
2. Make system changes.
3. In the UI, deploy to staging and test.
4. In the UI, promote to production.

APICast retrieves and publishes the promoted update at the configured frequency.

1.3. APICAST BUILT-IN WILDCARD ROUTING

The built-in APIcast gateways that accompany your on-premise 3scale deployment support wildcard domain routing at the subdomain level. This feature allows you to name a portion of your subdomain for your production and staging gateway public base URLs. To use this feature, you must have enabled it during the 3scale installation.



NOTE

Ensure that you are using the OpenShift Container Platform version that supports Wildcard Routing. For information on the supported versions, see [Supported Configurations](#).

The AMP does not provide DNS capabilities, so your specified public base URL must match the DNS configuration specified in the **WILDCARD_DOMAIN** parameter of the OpenShift cluster on which it was deployed.

1.3.1. Modify wildcards

Perform the following steps to modify your wildcards:

1. Log in to your AMP.
2. Navigate to your API gateway settings page: **APIs** → your API → **Integration** → **edit APIcast configuration**
3. Modify the staging and production public base URLs with a string prefix of your choice, adhere to these requirements:
 - API endpoints must not begin with a numeric character

The following is an example of a valid wildcard for a staging gateway on the domain **example.com**:

```
apiname-staging.example.com
```

More Information

For information on routing, see the [OpenShift documentation](#).

1.4. SCALING UP 3SCALE ON-PREMISE

1.4.1. Scaling up storage

As your APIcast deployment grows, you may need to increase the amount of storage available. How you scale up storage depends on which type of file system you are using for your persistent storage.

If you are using a network file system (NFS), you can scale up your persistent volume using the **oc edit pv** command:

```
oc edit pv <pv_name>
```

If you are using any other storage method, you must scale up your persistent volume manually using one of the methods listed in the following sections.

1.4.1.1. Method 1: Backup and swap persistent volumes

1. Back up the data on your existing persistent volume.
2. Create and attach a target persistent volume, scaled for your new size requirements.
3. Create a pre-bound persistent volume claim. Specify the size of your new PVC The persistent volume name using the **volumeName** field.
4. Restore data from your backup onto your newly created PV.
5. Modify your deployment configuration with the name of your new PV:

```
oc edit dc/system-app
```

6. Verify your new PV is configured and working correctly.
7. Delete your previous PVC to release its claimed resources.

1.4.1.2. Method 2: Back up and redeploy 3scale

1. Back up the data on your existing persistent volume.
2. Shut down your 3scale pods.
3. Create and attach a target persistent volume, scaled for your new size requirements.
4. Restore data from your backup onto your newly created PV.
5. Create a pre-bound persistent volume claim. Specify:
 - a. The size of your new PVC
 - b. The persistent volume name using the **volumeName** field.
6. Deploy your AMP.yml.
7. Verify your new PV is configured and working correctly.
8. Delete your previous PVC to release its claimed resources.

1.4.2. Scaling up performance

1.4.2.1. Configuring 3scale on-premise deployments

By default, 3scale deployments run one process per pod. You can increase performance by running more processes per pod. Red Hat recommends running 1-2 processes per core on each node.

Perform the following steps to add more processes to a pod:

1. Log in to your OpenShift cluster.

```
oc login
```

2. Switch to your 3scale project.

```
oc project <project_name>
```

3. Set the appropriate environment variable to the desired number of processes per pod.
 - a. **APICAST_WORKERS** for APICast pods (Red Hat recommends to keep this environment variable unset to allow APICast to determine the number of workers by the number of CPUs available to the APICast pod)
 - b. **PUMA_WORKERS** for backend pods
 - c. **UNICORN_WORKERS** for system pods

```
oc set env dc/apicast-{production/staging} --overwrite APICAST_WORKERS=  
<number_of_processes>
```

```
oc set env dc/backend-listener --overwrite PUMA_WORKERS=<number_of_processes>
```

```
oc set env dc/system-app --overwrite UNICORN_WORKERS=<number_of_processes>
```

1.4.2.2. Vertical and horizontal hardware scaling

You can increase the performance of your AMP deployment on OpenShift by adding resources. You can add more compute nodes as pods to your OpenShift cluster (horizontal scaling) or you can allocate more resources to existing compute nodes (vertical scaling).

Horizontal Scaling

You can add more compute nodes as pods to your OpenShift. If the additional compute nodes match the existing nodes in your cluster, you do not have to reconfigure any environment variables.

Vertical Scaling

You can allocate more resources to existing compute nodes. If you allocate more resources, you must add additional processes to your pods to increase performance.



NOTE

Red Hat does not recommend mixing compute nodes of a different specification or configuration on your 3scale deployment.

1.4.2.3. Scaling up routers

As your traffic increases, you must ensure your OCP routers can adequately handle requests. If your routers are limiting the throughput of your requests, you must scale up your router nodes.

1.4.2.4. Further reading

- Scaling tasks, adding hardware compute nodes to OpenShift
- Adding Compute Nodes
- Routers

1.5. OPERATIONS TROUBLESHOOTING

1.5.1. Accessing your logs

Each component's deployment configuration contains logs for access and exceptions. If you encounter issues with your deployment, check these logs for details.

Follow these steps to access logs in 3scale:

1. Find the identifier (ID) of the pod you want logs for:

```
oc get pods
```

2. Enter **oc logs** and the ID of your chosen pod:

```
oc logs <pod>
```

The system pod has two containers, each with a separate log. To access a container's log, specify the **--container** parameter with the **system-provider** and **system-developer**:

```
oc logs <pod> --container=system-provider
oc logs <pod> --container=system-developer
```

1.5.2. Checking job queues

Job Queues contain logs of information sent from the **system-sidekiq** pods. Use these logs to check if your cluster is processing data. You can query the logs using the OpenShift CLI:

```
oc get jobs
```

```
oc logs <job>
```

1.5.3. Preventing monotonic growth

To prevent monotonic growth, 3scale schedules by default, automatic purging of the following tables:

- *user_sessions* - clean up is triggered once a week, deletes records older than two weeks.
- *audits* - clean up is triggered once a day, deletes records older than three months.
- *log_entries* - clean up triggered once a day, deletes records older than six months.
- *event_store_events* - clean up is triggered once a week, deletes records older than a week.

With the exception of the above listed tables, the *alerts* table requires manual purging by the database administrator.

Database type	SQL command
MySQL	<pre>DELETE FROM alerts WHERE timestamp < NOW() - INTERVAL 14 DAY;</pre>

Database type	SQL command
PostgreSQL	<pre>DELETE FROM alerts WHERE timestamp < NOW() - INTERVAL '14 day';</pre>
Oracle	<pre>DELETE FROM alerts WHERE timestamp <= TRUNC(SYSDATE) - 14;</pre>

For other tables not specified in this section, the database administrator must manually clean the tables that the system does not automatically purge.

CHAPTER 2. USING THE 3SCALE TOOLBOX

The [3scale toolbox](#) is a Ruby client that lets you manage 3scale services from the command line.

2.1. INSTALLING THE TOOLBOX

The only officially supported method of installing the toolbox is on Red Hat Enterprise Linux using **yum** or **rpm**.

Prerequisites

You must enable access to the following repositories:

- **rhel-7-server-3scale-amp-2.5-rpms**
- **rhel-server-rhscl-7-rpms**

For example:

```
$ sudo subscription-manager repos --enable=rhel-7-server-3scale-amp-2.5-rpms --enable rhel-server-rhscl-7-rpms
```

Procedure

1. Install the toolbox:

```
$ yum install 3scale_toolbox
```

2. Verify the installation:

```
$ 3scale help
```

Additional resources

You can however use unsupported versions on Fedora Linux, Ubuntu Linux, Windows and macOS by downloading and installing the [.rpm](#), [.deb](#), [.msi](#) or [.pkg](#) file directly from GitHub.

2.2. IMPORTING SERVICES

Import services from a CSV file by specifying the following fields in this order (you also need to include these headers in your CSV file):

```
service_name,endpoint_name,endpoint_http_method,endpoint_path,auth_mode,endpoint_system_name,type
```

You will need the following information:

- 3scale admin account: **{3SCALE_ADMIN}**
- domain your 3scale instance is running on: **{DOMAIN_NAME}** (if you are using hosted APICast this is 3scale.net)
- [access key](#) of your account: **{ACCESS_KEY}**

- CSV file of services (**examples/import_example.csv**)

Import the services by running:

```
$ 3scale import csv --destination=https://{ACCESS_KEY}@{3SCALE_ADMIN}-admin.
{DOMAIN_NAME} --file=examples/import_example.csv
```

2.3. COPYING SERVICES

Create a new service based on an existing one from the same account or from another account. When you copy a service, the relevant ActiveDocs are also copied.

You need the following information:

- service id you want to copy: **{SERVICE_ID}**
- 3scale admin account: **{3SCALE_ADMIN}**
- domain your 3scale instance is running on: **{DOMAIN_NAME}** (if you are using hosted APICast this is 3scale.net)
- [access key](#) of your account: **{ACCESS_KEY}**
- access key of the destination account if you are copying to a different account: **{DEST_KEY}**
- name for the new service: **{NEW_NAME}**

```
$ 3scale copy service {SERVICE_ID} --source=https://{ACCESS_KEY}@{3SCALE_ADMIN}-admin.
{DOMAIN_NAME} --destination=https://{DEST_KEY}@{3SCALE_ADMIN}-admin.{DOMAIN_NAME} -
-target_system_name={NEW_NAME}
```

2.4. COPYING SERVICE SETTINGS ONLY

Bulk copy (also known as updating) the service settings, proxy settings, metrics, methods, application plans, application plan limits and mapping rules from one service to another which already exists.

You will need the following information:

- service id you want to copy: **{SERVICE_ID}**
- service id of the destination: **{DEST_ID}**
- 3scale admin account: **{3SCALE_ADMIN}**
- domain your 3scale instance is running on: **{DOMAIN_NAME}** (if you are using hosted APICast this is 3scale.net)
- [access key](#) of your account: **{ACCESS_KEY}**
- access key of the destination account: **{DEST_KEY}**

And can use the following optional flags:

- **-f** remove existing target service mapping rules before copying

- **-r** copy only mapping rules to target service

```
$ 3scale update [opts] service --source=https://{ACCESS_KEY}@{3SCALE_ADMIN}-admin.
{DOMAIN_NAME} --destination=https://{DEST_KEY}@{3SCALE_ADMIN}-admin.{DOMAIN_NAME}
{SERVICE_ID} {DEST_ID}
```

2.5. IMPORTING OPENAPI DEFINITIONS

You can import definitions based on OpenAPI Specification (OAS) in the following formats: **json** and **yaml**. To create a new service or to update an existing service, use the definitions from a local file or a web address.

The default service name for the import is taken from **info.title** in the OpenAPI definition, and you can override it using **--target_system_name=<NEW NAME>**. If that service name already exists, it will be updated. Conversely, if the service name does not exist, it will be created.

The following rules apply to every import:

- Definitions are validated as OpenAPI 2.0.
- All mapping rules in the 3scale service are deleted.
- In order to be replaced, all method names must be identical to methods defined in the OpenAPI definition (**operation.operationId**) by using exact pattern matching.
- Only methods included in the OpenAPI definition are modified.
- All methods that were present only in the OpenAPI definition are attached to the **Hits** metric.
- All mapping rules from the OpenAPI definition are imported. View these in **API > Integration**.

```
$ 3scale import openapi [opts] --destination=https://{DEST_KEY}@{3SCALE_ADMIN}-admin.
{DOMAIN_NAME}
```

2.5.1. Optional flags

-d --destination=<value>

3scale target instance. Format: **http[s]://<authentication>@3scale_domain**

-t --target_system_name=<value>

Target system name

2.6. MANAGING REMOTE WEB ADDRESSES

To facilitate work with 3scale instances, you can define in a config file the remote web addresses (URLs) along with authentication, and refer to them by a short name in any 3scale toolbox command.

The default location for the config file is **\$HOME/.3scalerc.yaml** but you can specify another location using the **THREESCALE_CLI_CONFIG** environment variable or the **--config-file <config_file>** option.

You can specify remotes using either an **access_token** or a **provider_key**:

- **http[s]://<access_token>@<3scale-instance-domain>**

- `http[s]://<provider_key>@<3scale-instance-domain>`

2.6.1. Listing remote web addresses

```
3scale remote list [--config-file <config_file>]
```

Shows the list of existing remotes (name, URL and authentication key).

Example

```
$ 3scale remote list
instance_a https://example_a.net 123456789
instance_b https://example_b.net 987654321
```

2.6.2. Adding remote web addresses

```
3scale remote add [--config-file <config_file>] <name> <url>
```

Adds a remote with short name **<name>** at **<url>**.

Example

```
3scale remote add instance_a https://123456789@example_a.net
```

2.6.3. Removing remote web addresses

```
3scale remote remove [--config-file <config_file>] <name>
```

Removes the remote with short name **<name>**.

Example

```
3scale remote remove instance_a
```

2.6.4. Renaming remote web addresses

```
3scale remote rename [--config-file <config_file>] <old_name> <new_name>
```

Renames remote with short name **<old_name>** to **<new_name>**.

Example

```
3scale remote rename instance_a instance_b
```

2.7. TROUBLESHOOTING SSL ISSUES

There is more information on [certificates](#) in the Red Hat documentation, but if you are getting issues related to self-signed SSL certificates, **SSL certificate problem: self signed certificate** for example, you can download and use the remote certificate:

1. Download the remote certificate using **openssl**

```
$ echo | openssl s_client -showcerts -servername self-signed.badssl.com -connect self-signed.badssl.com:443 2>/dev/null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > self-signed-cert.pem
```

2. Make sure it works by using it with **curl**

```
$ SSL_CERT_FILE=self-signed-cert.pem curl -v https://self-signed.badssl.com
```

If the certificate is working properly, you will not get the SSL error.

3. Prefix your **3scale** commands with **SSL_CERT_FILE=self-signed-cert.pem**:

```
$ SSL_CERT_FILE=self-signed-cert.pem 3scale import csv
```

CHAPTER 3. PROVISIONING 3SCALE SERVICES AND CONFIGURATIONS VIA THE OPERATOR (CAPABILITIES)

3.1. INTRODUCTION

This document provides information about provisioning 3scale services and configurations via the 3scale operator.



IMPORTANT

The 3scale operator is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

3.1.1. Prerequisites

- A 3scale 2.5 [On-Premises](#) instance
- You must have the [3scale operator installed](#).
- OpenShift Container Platform 3.11
 - A user account with administrator privileges in the OpenShift cluster



WARNING

When using the operator to update API configurations in 3scale, the custom resource definitions (CRDs) are the source of truth. If changes are made in the Admin UI, they will not persist and eventually be overridden by the definition in the CRD.

3.2. DEPLOYING CAPABILITIES RELATED CUSTOM RESOURCES

You will start to configure APIs, *metrics* and *mappingrules* in you newly created tenant by only using *Openshift Objects*.

1. Create your first API

```
apiVersion: capabilities.3scale.net/v1alpha1
kind: API
metadata:
  creationTimestamp: 2019-01-25T13:28:41Z
  generation: 1
  labels:
    environment: testing
```

```

name: api01
spec:
  planSelector:
    matchLabels:
      api: api01
  description: api01
  integrationMethod:
    apicastHosted:
      apiTestGetRequest: /
    authenticationSettings:
      credentials:
        apiKey:
          authParameterName: user-key
          credentialsLocation: headers
    errors:
      authenticationFailed:
        contentType: text/plain; charset=us-ascii
        responseBody: Authentication failed
        responseCode: 403
      authenticationMissing:
        contentType: text/plain; charset=us-ascii
        responseBody: Authentication Missing
        responseCode: 403
      hostHeader: ""
      secretToken: Shared_secret_sent_from_proxy_to_API_backend_9603f637ca51ccfe
  mappingRulesSelector:
    matchLabels:
      api: api01
    privateBaseURL: https://echo-api.3scale.net:443
  metricSelector:
    matchLabels:
      api: api01

```

In all the Selectors (*metric*, *plan*, *mappingrules*) we use a specific label **api: api01**, you can change that and add as many labels and play with the selectors to cover really complex scenarios.

2. Add a *Plan*

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: Plan
metadata:
  labels:
    api: api01
  name: plan01
spec:
  approvalRequired: false
  default: true
  costs:
    costMonth: 0
    setupFee: 0
  limitSelector:
    matchLabels:
      api: api01
  trialPeriod: 0

```

3. Add a metric called **metric01**

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: Metric
metadata:
  labels:
    api: api01
    name: metric01
spec:
  description: metric01
  unit: hit
  incrementHits: false

```

4. Set a limit with a limit of 10 hits per day for the metric

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: Limit
metadata:
  labels:
    api: api01
    name: plan01-metric01-day-10
spec:
  description: Limit for metric01 in plan01
  maxValue: 10
  metricRef:
    name: metric01
  period: day

```

5. Add a *MappingRule* to increment **metric01**

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: MappingRule
metadata:
  labels:
    api: api01
    name: metric01-get-path01
spec:
  increment: 1
  method: GET
  metricRef:
    name: metric01
  path: /path01

```

6. Bind using the binding object

You will use the credentials created by the *Tenant Controller*

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: Binding
metadata:
  name: mytestingbinding
spec:
  credentialsRef:
    name: ecorp-tenant-secret

```

```

APISelector:
  matchLabels:
    environment: testing

```

The binding object references the **ecorp-tenant-secret** and creates the API objects that are labeled as **environment: staging**.

- Navigate to your new 3scale tenant and check that everything has been created.



NOTE

For more information, check the reference documentation: [Capabilities CRD Reference](#).

3.3. DEPLOYING OPTIONAL TENANTS CUSTOM RESOURCE

Optionally, you may create other tenants deploying *Tenant* custom resource objects.

- Deploy a new tenant in your 3scale instance by creating a secret to store the administrator password:

```

$ cat ecorp-admin-secret.yaml
apiVersion: v1
kind: Secret
metadata:
  name: ecorp-admin-secret
type: Opaque
stringData:
  admin_password: <admin password value>

$ oc create -f ecorp-admin-secret.yaml
secret/ecorp-admin-secret created

```

- Create a new tenant *CR YAML* file with the following content:

```

apiVersion: capabilities.3scale.net/v1alpha1
kind: Tenant
metadata:
  name: ecorp-tenant
spec:
  username: admin
  systemMasterUrl: https://master.<wildcardDomain>
  email: admin@ecorp.com
  organizationName: ECorp
  masterCredentialsRef:
    name: system-seed
  passwordCredentialsRef:
    name: ecorp-admin-secret
  tenantSecretRef:
    name: ecorp-tenant-secret
  namespace: operator-test

```

**NOTE**

For more information about the *Tenant Custom Resource* fields and possible values, refer to the [Tenant CRD Reference](#) documentation.

```
export NAMESPACE="operator-test"
oc project ${NAMESPACE}
oc create -f <yaml-name>
```

- a. This should trigger the creation of a new tenant in your 3scale solution in the *operator-test* project.

Tenant *provider_key* and *admin domain URL* will be stored in a secret. You can specify the secret location by using **tenantSecretRef** tenant spec key.

3.4. DELETING CREATED CUSTOM RESOURCES

**WARNING**

Deleting the *APIManager* will delete the 3scale installation.

- Delete the *APIManager* custom resource and the 3scale solution elements that have been deployed from it.
 - Deleting the *APIManager* will delete all 3scale related objects in where it has been deployed:

```
oc delete -f <yaml-name-of-the-apimanager-custom-resource>
```

- Delete the 3scale operator, its associated roles and service accounts.

```
oc delete -f deploy/operator.yaml
oc delete -f deploy/role_binding.yaml
oc delete -f deploy/service_account.yaml
oc delete -f deploy/role.yaml
```

- Delete the *APIManager* and *Capabilities* related CRDs.

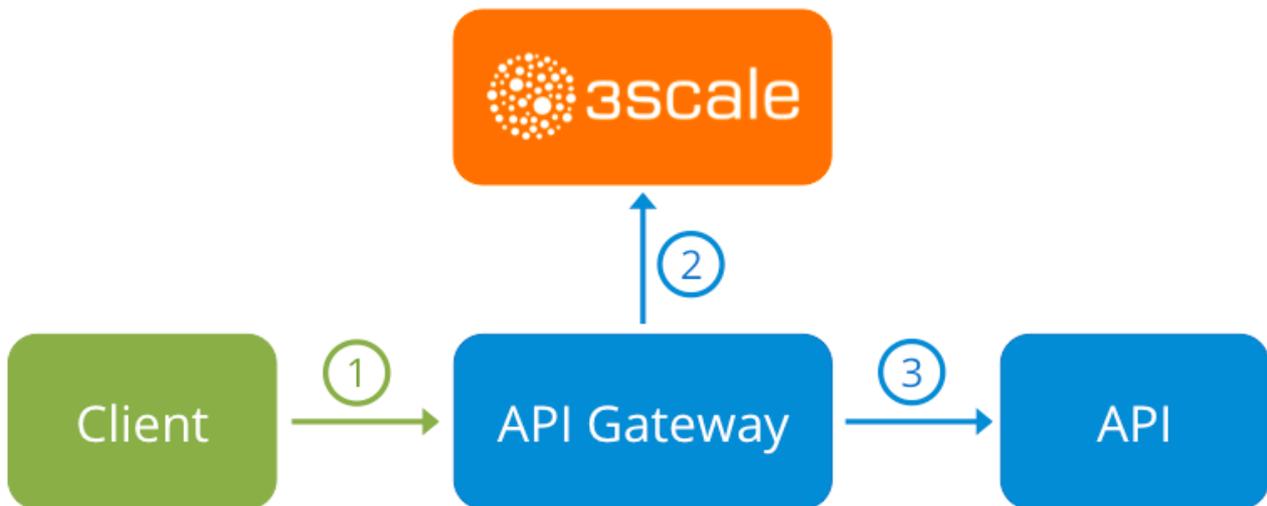
```
oc delete -f deploy/crds/
```

CHAPTER 4. TROUBLESHOOTING

This guide aims to help you identify and fix the cause of issues with your API infrastructure.

API Infrastructure is a lengthy and complex topic. However, at a minimum, you will have three moving parts in your Infrastructure:

1. The API gateway
2. 3scale
3. The API



Errors in any of these three elements results in your clients being unable to access your API. However, it is difficult to find the component that caused the failure. This guide gives you some tips to troubleshoot your infrastructure to identify the problem.

4.1. COMMON ISSUES

There are some evidences that can point to some very common issues with your integration with 3scale. These will vary depending on whether you are at the beginning of your API project, setting up your infrastructure, or are already live in production.

4.1.1. Integration issues

The following sections attempt to outline some common issues you may see in the APIcast error log during the initial phases of your integration with 3scale: at the beginning using APIcast Hosted and prior to go-live, running the self-managed APIcast.

4.1.1.1. APIcast Hosted

When you are first integrating your API with APIcast Hosted on the Service Integration screen, you might get some of the following errors shown on the page or returned by the test call you make to check for a successful integration.

- **Test request failed: execution expired**

Check that your API is reachable from the public internet. APIcast Hosted cannot be used with private APIs. If you don't want to make your API publicly available to integrate with APIcast

Hosted, you can set up a private secret between APIcast Hosted and your API to reject any calls not coming from the API gateway.

- The accepted format is 'protocol://address(:port)'. Remove any paths at the end of your API's private base URL. You can add these in the "mapping rules" pattern or at the beginning of the "API test GET request".
- **Test request failed with HTTP code XXX**
 - **405:** Check that the endpoint accepts GET requests. APIcast only supports GET requests to test the integration.
 - **403: Authentication parameters missing:** If your API already has some authentication in place, APIcast will be unable to make a test request.
 - **403: Authentication failed:** If this is not the first service you have created with 3scale, check that you have created an application under the service with credentials to make the test request. If it is the first service you are integrating, ensure that you have not deleted the test account or application that you created on signup.

4.1.1.2. APIcast self-managed

After you have successfully tested the integration with APIcast self-managed, you might want to host the API gateway yourself. Following are some errors you may encounter when you first install your self-managed gateway and call your API through it.

- **upstream timed out (110: Connection timed out) while connecting to upstream**
Check that there are no firewalls or proxies between the API Gateway and the public Internet that would prevent your self-managed gateway from reaching 3scale.
- **failed to get list of services: invalid status: 403 (Forbidden)**

```
2018/06/04 08:04:49 [emerg] 14#14: [lua] configuration_loader.lua:134: init(): failed to load
configuration, exiting (code 1)
2018/06/04 08:04:49 [warn] 22#22: *2 [lua] remote_v2.lua:163: call(): failed to get list of
services: invalid status: 403 (Forbidden) url: https://example-
admin.3scale.net/admin/api/services.json , context: ngx.timer
ERROR: /opt/app-root/src/src/apicast/configuration_loader.lua:57: missing configuration
```

Check that the Access Token that you used in the **THREESCALE_PORTAL_ENDPOINT** value is correct and that it has the Account Management API scope. Verify it with a **curl** command: **curl -v "https://example-admin.3scale.net/admin/api/services.json?access_token=<YOUR_ACCESS_TOKEN>"**

It should return a 200 response with a JSON body. If it returns an error status code, check the response body for details.

- **service not found for host apicast.example.com**

```
2018/06/04 11:06:15 [warn] 23#23: *495 [lua] find_service.lua:24: find_service(): service not
found for host apicast.example.com, client: 172.17.0.1, server: _, request: "GET / HTTP/1.1",
host: "apicast.example.com"
```

This error indicates that the Public Base URL has not been configured properly. You should ensure that the configured Public Base URL is the same that you use for the request to self-managed APIcast. After configuring the correct Public Base URL:

- Ensure that APICast is configured for "production" (default configuration for standalone APICast if not overridden with **THREESCALE_DEPLOYMENT_ENV** variable). Ensure that you promote the configuration to production.
- Restart APICast, if you have not configured auto-reloading of configuration using **APICAST_CONFIGURATION_CACHE** and **APICAST_CONFIGURATION_LOADER** environment variables.

Following are some other symptoms that may point to an incorrect APICast self-managed integration:

- **Mapping rules not matched / Double counting of API calls** Depending on the way you have defined the mapping between methods and actual URL endpoints on your API, you might find that sometimes methods either don't get matched or get incremented more than once per request. To troubleshoot this, make a test call to your API with the [3scale debug header](#). This will return a list of all the methods that have been matched by the API call.
- **Authentication parameters not found:** Ensure you are sending the parameters to the correct location as specified in the Service Integration screen. If you do not send credentials as headers, the credentials must be sent as query parameters for GET requests and body parameters for all other HTTP methods. Use the 3scale debug header to double-check the credentials that are being read from the request by the API gateway.

4.1.2. Production issues

It is rare to run into issues with your API gateway after you have fully tested your setup and have been live with your API for a while. However, here are some of the issues you might encounter in a live production environment.

4.1.2.1. Availability issues

Availability issues are normally characterised by **upstream timed out** errors in your nginx error.log; example:

```
upstream timed out (110: Connection timed out) while connecting to upstream, client: X.X.X.X,
server: api.example.com, request: "GET /RESOURCE?CREDENTIALS HTTP/1.1", upstream:
"http://Y.Y.Y.Y:80/RESOURCE?CREDENTIALS", host: "api.example.com"
```

If you are experiencing intermittent 3scale availability issues, following may be the reasons for this:

- You are resolving to an old 3scale IP that is no longer in use. The latest version of the API gateway configuration files defines 3scale as a variable to force IP resolution each time. For a quick fix, reload your NGINX instance. For a long-term fix, ensure that instead of defining the 3scale backend in an upstream block, you define it as a variable within each server block; example:

```
server {
    # Enabling the Lua code cache is strongly encouraged for production use. Here it is enabled
    .
    .
    .
    set $threescale_backend "https://su1.3scale.net:443";
```

When you refer to it:

```
location = /threescale_authrep {
```

```

internal;
set $provider_key "YOUR_PROVIDER_KEY";

proxy_pass $threescale_backend/transactions/authrep.xml?
provider_key=$provider_key&service_id=$service_id&$usage&$credentials&log%5Bcode%5
D=$arg_code&log%5Brequest%5D=$arg_req&log%5Bresponse%5D=$arg_resp;
}

```

- You are missing some 3scale IPs from your whitelist. Following is the current list of IPs that 3scale resolves to:
 - 75.101.142.93
 - 174.129.235.69
 - 184.73.197.122
 - 50.16.225.117
 - 54.83.62.94
 - 54.83.62.186
 - 54.83.63.187
 - 54.235.143.255

The above issues refer to problems with perceived 3scale availability. However, you might encounter similar issues with your API availability from the API gateway if your API is behind an AWS ELB. This is because NGINX, by default, does DNS resolution at start-up time and then caches the IP addresses. However, ELBs do not ensure static IP addresses and these might change frequently. Whenever the ELB changes to a different IP, NGINX is unable to reach it.

The solution for this is similar to the above fix for forcing runtime DNS resolution.

1. Set a specific DNS resolver such as Google DNS, by adding this line at the top of the **http** section: **resolver 8.8.8.8 8.8.4.4;**
2. Set your API base URL as a variable anywhere near the top of the **server** section. **set \$api_base "http://api.example.com:80";**
3. Inside the **location /** section, find the **proxy_pass** line and replace it with **proxy_pass \$api_base;**

4.1.3. Post-deploy issues

If you make changes to your API such as adding a new endpoint, you must ensure that you add a new method and URL mapping before downloading a new set of configuration files for your API gateway.

The most common problem when you have modified the configuration downloaded from 3scale will be code errors in the Lua, which will result in a **500 - Internal server error** such as:

```

curl -v -X GET "http://localhost/"
* About to connect() to localhost port 80 (#0)
* Trying 127.0.0.1... connected
> GET / HTTP/1.1

```

```

> User-Agent: curl/7.22.0 (x86_64-pc-linux-gnu) libcurl/7.22.0 OpenSSL/1.0.1 zlib/1.2.3.4 libidn/1.23
librtmp/2.3
> Host: localhost
> Accept: */*
>
< HTTP/1.1 500 Internal Server Error
< Server: openresty/1.5.12.1
< Date: Thu, 04 Feb 2016 10:22:25 GMT
< Content-Type: text/html
< Content-Length: 199
< Connection: close
<

<head><title>500 Internal Server Error</title></head>

<center><h1>500 Internal Server Error</h1></center>
<hr><center>openresty/1.5.12.1</center>

* Closing connection #0

```

You can see the nginx error.log to know the cause, such as:

```

2016/02/04 11:22:25 [error] 8980#0: *1 lua entry thread aborted: runtime error:
/home/pili/NGINX/troubleshooting/nginx.lua:66: bad argument #3 to '_newindex' (number expected,
got nil)
stack traceback:
coroutine 0:
  [C]: in function '_newindex'
  /home/pili/NGINX/troubleshooting/nginx.lua:66: in function 'error_authorization_failed'
  /home/pili/NGINX/troubleshooting/nginx.lua:330: in function 'authrep'
  /home/pili/NGINX/troubleshooting/nginx.lua:283: in function 'authorize'
  /home/pili/NGINX/troubleshooting/nginx.lua:392: in function 'while' sending to client, client:
127.0.0.1, server: api-2445581381726.staging.apicast.io, request: "GET / HTTP/1.1", host: "localhost"

```

In the access.log this will look like the following:

```

127.0.0.1 - - [04/Feb/2016:11:22:25 +0100] "GET / HTTP/1.1" 500 199 "-" "curl/7.22.0 (x86_64-pc-
linux-gnu) libcurl/7.22.0 OpenSSL/1.0.1 zlib/1.2.3.4 libidn/1.23 librtmp/2.3"

```

The above section gives you an overview of the most common, well-known issues that you might encounter at any stage of your 3scale journey.

If all of these have been checked and you are still unable to find the cause and solution for your issue, you should proceed to the more detailed [operational troubleshooting](#troubleshooting-checklists) sections below. Start at your API and work your way back to the client in order to try to identify the point of failure.

4.2. TROUBLESHOOTING 101

If you are experiencing failures when connecting to a server, whether that is the API gateway, 3scale, or your API, the following troubleshooting steps should be your first port of call:

4.2.1.1. Can we connect?

Use telnet to check the basic TCP/IP connectivity **telnet api.example.com 443**

- Success

```
telnet echo-api.3scale.net 80
Trying 52.21.167.109...
Connected to tf-lb-i2t5pgt2cfdnbdhf2c6qqoartm-829217110.us-east-1.elb.amazonaws.com.
Escape character is '^]'.
Connection closed by foreign host.
```

- Failure

```
telnet su1.3scale.net 443
Trying 174.129.235.69...
telnet: Unable to connect to remote host: Connection timed out
```

4.2.2. 2. Is it me or is it them?

Try to connect to the same server from different network locations, devices, and directions. For example, if your client is unable to reach your API, try to connect to your API from a machine that should have access such as the API gateway.

If any of the attempted connections succeed, you can rule out any problems with the actual server and concentrate your troubleshooting on the network between them, as this is where the problem will most likely be.

4.2.3. 3. Is it a DNS issue?

Try to connect to the server by using its IP address instead of its hostname e.g. **telnet 94.125.104.17 80** instead of **telnet apis.io 80**

This will rule out any problems with the DNS.

You can get the IP address for a server using **dig** for example for 3scale **dig su1.3scale.net** or **dig any su1.3scale.net** if you suspect there may be multiple IPs that a host may resolve to.

NB: Some hosts block `dig any`

4.2.4. 4. Is it an SSL issue?

You can use OpenSSL to test:

- Secure connections to a host or IP, such as from the shell prompt **openssl s_client -connect su1.3scale.net:443**

Output:

```
CONNECTED(00000003)
depth=1 C = US, O = GeoTrust Inc., CN = GeoTrust SSL CA - G3
verify error:num=20:unable to get local issuer certificate
---
Certificate chain
 0 s:/C=ES/ST=Barcelona/L=Barcelona/O=3scale Networks, S.L./OU=IT/CN=*.3scale.net
  i:/C=US/O=GeoTrust Inc./CN=GeoTrust SSL CA - G3
 1 s:/C=US/O=GeoTrust Inc./CN=GeoTrust SSL CA - G3
```

```

i:/C=US/O=GeoTrust Inc./CN=GeoTrust Global CA
---
Server certificate
-----BEGIN CERTIFICATE-----
MIIE8zCCA9ugAwIBAgIQcz2Y9JNxxH7f2zpOT0DajUjANBgkqhkiG9w0BAQsFADBE
...
TRUNCATED
...
3FZigX+OpWLVrjYsr0kZzX+HCerYMwc=
-----END CERTIFICATE-----
subject=/C=ES/ST=Barcelona/L=Barcelona/O=3scale Networks,
S.L./OU=IT/CN=*.3scale.net
issuer=/C=US/O=GeoTrust Inc./CN=GeoTrust SSL CA - G3
---
Acceptable client certificate CA names
/C=ES/ST=Barcelona/L=Barcelona/O=3scale Networks, S.L./OU=IT/CN=*.3scale.net
/C=US/O=GeoTrust Inc./CN=GeoTrust SSL CA - G3
Client Certificate Types: RSA sign, DSA sign, ECDSA sign
Requested Signature Algorithms:
RSA+SHA512:DSA+SHA512:ECDSA+SHA512:RSA+SHA384:DSA+SHA384:ECDSA+SHA384
:RSA+SHA256:DSA+SHA256:ECDSA+SHA256:RSA+SHA224:DSA+SHA224:ECDSA+SHA22
4:RSA+SHA1:DSA+SHA1:ECDSA+SHA1:RSA+MD5
Shared Requested Signature Algorithms:
RSA+SHA512:DSA+SHA512:ECDSA+SHA512:RSA+SHA384:DSA+SHA384:ECDSA+SHA384
:RSA+SHA256:DSA+SHA256:ECDSA+SHA256:RSA+SHA224:DSA+SHA224:ECDSA+SHA22
4:RSA+SHA1:DSA+SHA1:ECDSA+SHA1
Peer signing digest: SHA512
Server Temp Key: ECDH, P-256, 256 bits
---
SSL handshake has read 3281 bytes and written 499 bytes
---
New, TLSv1/SSLv3, Cipher is ECDHE-RSA-AES256-GCM-SHA384
Server public key is 2048 bit
Secure Renegotiation IS supported
Compression: NONE
Expansion: NONE
No ALPN negotiated
SSL-Session:
  Protocol : TLSv1.2
  Cipher   : ECDHE-RSA-AES256-GCM-SHA384
  Session-ID:
A85EFD61D3BFD6C27A979E95E66DA3EC8F2E7B3007C0166A9BCBDA5DCA5477B8
  Session-ID-ctx:
  Master-Key:
F7E898F1D996B91D13090AE9D5624FF19DFE645D5DEEE2D595D1B6F79B1875CF935B3
A4F6ECCA7A6D5EF852AE3D4108B
  Key-Arg  : None
  PSK identity: None
  PSK identity hint: None
  SRP username: None
  TLS session ticket lifetime hint: 300 (seconds)
  TLS session ticket:
0000 - a8 8b 6c ac 9c 3c 60 78-2c 5c 8a de 22 88 06 15  ..!..<`x,\."...
0010 - eb be 26 6c e6 7b 43 cc-ae 9b c0 27 6c b7 d9 13  ..&!.{C....!...
0020 - 84 e4 0d d5 f1 ff 4c 08-7a 09 10 17 f3 00 45 2c  .....L.z.....E,
0030 - 1b e7 47 0c de dc 32 eb-ca d7 e9 26 33 26 8b 8e  ..G...2....&3&..

```

```

0040 - 0a 86 ee f0 a9 f7 ad 8a-f7 b8 7b bc 8c c2 77 7b .....{...w{
0050 - ae b7 57 a8 40 1b 75 c8-25 4f eb df b0 2b f6 b7 ..W.@.u.%O...+..
0060 - 8b 8e fc 93 e4 be d6 60-0f 0f 20 f1 0a f2 cf 46 .....`.. ....F
0070 - b0 e6 a1 e5 31 73 c2 f5-d4 2f 57 d1 b0 8e 51 cc ....1s.../W...Q.
0080 - ff dd 6e 4f 35 e4 2c 12-6c a2 34 26 84 b3 0c 19 ..nO5.,l.4&....
0090 - 8a eb 80 e0 4d 45 f8 4a-75 8e a2 06 70 84 de 10 ....ME.Ju...p...

```

```

Start Time: 1454932598
Timeout : 300 (sec)
Verify return code: 20 (unable to get local issuer certificate)

```

```
---
```

- SSLv3 support (NOT supported by 3scale)
openssl s_client -ssl3 -connect su.3scale.net:443

Output

```

CONNECTED(00000003)
140735196860496:error:14094410:SSL routines:ssl3_read_bytes:ssl3 alert handshake
failure:s3_pkt.c:1456:SSL alert number 40
140735196860496:error:1409E0E5:SSL routines:ssl3_write_bytes:ssl handshake
failure:s3_pkt.c:644:
---
no peer certificate available
---
No client certificate CA names sent
---
SSL handshake has read 7 bytes and written 0 bytes
---
New, (NONE), Cipher is (NONE)
Secure Renegotiation IS NOT supported
Compression: NONE
Expansion: NONE
No ALPN negotiated
SSL-Session:
    Protocol : SSLv3
    Cipher   : 0000
    Session-ID:
    Session-ID-ctx:
    Master-Key:
    Key-Arg  : None
    PSK identity: None
    PSK identity hint: None
    SRP username: None
    Start Time: 1454932872
    Timeout  : 7200 (sec)
    Verify return code: 0 (ok)
---

```

For more details, see the [OpenSSL man pages](#).

4.3. TROUBLESHOOTING CHECKLISTS

To identify where an issue with requests to your API might lie, go through the following checks.

4.3.1. API

To confirm that the API is up and responding to requests, make the same request directly to your API (not going through the API gateway). You should ensure that you are sending the same parameters and headers as the request that goes through the API gateway. If you are unsure of the exact request that is failing, capture the traffic between the API gateway and your API.

If the call succeeds, you can rule out any problems with the API, otherwise you should troubleshoot your API further.

4.3.2. API Gateway > API

To rule out any network issues between the API gateway and the API, make the same call as before – directly to your API – from your API gateway server.

If the call succeeds, you can move on to troubleshooting the API gateway itself.

4.3.3. API gateway

There are a number of steps to go through to check that the API gateway is working correctly.

4.3.3.1. 1. Is the API gateway up and running?

Log in to the machine where the gateway is running. If this fails, your gateway server might be down.

After you have logged in, check that the NGINX process is running. For this, run **ps ax | grep nginx** or **htop**.

NGINX is running if you see **nginx master process** and **nginx worker process** in the list.

4.3.3.2. 2. Are there any errors in the gateway logs?

Following are some common errors you might see in the gateway logs, for example in error.log:

- API gateway can't connect to API

```
upstream timed out (110: Connection timed out) while connecting to upstream, client:
X.X.X.X, server: api.example.com, request: "GET /RESOURCE?CREDENTIALS HTTP/1.1",
upstream: "http://Y.Y.Y.Y:80/RESOURCE?CREDENTIALS", host: "api.example.com"
```

- API gateway cannot connect to 3scale

```
2015/11/20 11:33:51 [error] 3578#0: *1 upstream timed out (110: Connection timed out) while
connecting to upstream, client: 127.0.0.1, server: , request: "GET /api/activities.json?
user_key=USER_KEY HTTP/1.1", subrequest: "/threescale_authrep", upstream:
"https://54.83.62.186:443/transactions/authrep.xml?
provider_key=YOUR_PROVIDER_KEY&service_id=SERVICE_ID&usage[hits]=1&user_key=U
SER_KEY&log%5Bcode%5D=", host: "localhost"
```

4.3.4. API gateway > 3scale

Once you are sure the API gateway is running correctly, the next step is troubleshooting the connection between the API gateway and 3scale.

4.3.4.1. 1. Can the API gateway reach 3scale?

If you are using NGINX as your API gateway, the following message displays in the NGINX error logs when the gateway is unable to contact 3scale.

```
2015/11/20 11:33:51 [error] 3578#0: *1 upstream timed out (110: Connection timed out) while
connecting to upstream, client: 127.0.0.1, server: , request: "GET /api/activities.json?
user_key=USER_KEY HTTP/1.1", subrequest: "/threescale_authrep", upstream:
"https://54.83.62.186:443/transactions/authrep.xml?
provider_key=YOUR_PROVIDER_KEY&service_id=SERVICE_ID&usage[hits]=1&user_key=USER_KE
Y&log%5Bcode%5D=", host: "localhost"
```

Here, note the upstream value. This IP corresponds to one of the IPs that the 3scale service resolves to. This implies that there is a problem reaching 3scale. You can do a reverse DNS lookup to check the domain for an IP by calling **nslookup**.

For example, because the API gateway is unable to reach 3scale, it does not mean that 3scale is down. One of the most common reasons for this would be firewall rules preventing the API gateway from connecting to 3scale.

There may be network issues between the gateway and 3scale that could cause connections to timeout. In this case, you should go through the steps in [troubleshooting generic connectivity issues](#) to identify where the problem lies.

To rule out networking issues, use traceroute or MTR to check the routing and packet transmission. You can also run the same command from a machine that is able to connect to 3scale and your API gateway and compare the output.

Additionally, to see the traffic that is being sent between your API gateway and 3scale, you can use tcpdump as long as you temporarily switch to using the HTTP endpoint for the 3scale service (**su1.3scale.net**).

4.3.4.2. 2. Is the API gateway resolving 3scale addresses correctly?

Ensure you have the resolver directive added to your nginx.conf.

For example, in nginx.conf:

```
http {
    lua_shared_dict api_keys 10m;
    server_names_hash_bucket_size 128;
    lua_package_path "::$prefix/?.lua;";
    init_by_lua 'math.randomseed(ngx.time()); cjson = require("cjson");

    resolver 8.8.8.8 8.8.4.4;
```

You can substitute the Google DNS (8.8.8.8 and 8.8.4.4) with your preferred DNS.

To check DNS resolution from your API gateway, call nslookup as follows with the specified resolver IP:

```
nslookup su1.3scale.net 8.8.8.8
;; connection timed out; no servers could be reached
```

The above example shows the response returned if Google DNS cannot be reached. If this is the case, you must update the resolver IPs. You might also see the following alert in your nginx error.log:

```
2016/05/09 14:15:15 [alert] 9391#0: send() failed (1: Operation not permitted) while resolving,
resolver: 8.8.8.8:53
```

Finally, run **dig any su1.3scale.net** to see the IP addresses currently in operation for the 3scale Service Management API. Note that this is not the entire range of IP addresses that might be used by 3scale. Some may be swapped in and out for capacity reasons. Additionally, you may add more domain names for the 3scale service in the future. For this you should always test against the specific address that are supplied to you during integration, if applicable.

4.3.4.3. 3. Is the API gateway calling 3scale correctly?

If you want to check the request your API gateway is making to 3scale for troubleshooting purposes only you can add the following snippet to the 3scale authrep location in **nginx.conf** (**/threescale_authrep** for API Key and App_id authentication modes):

```
body_filter_by_lua_block{
    if ngx.req.get_headers()["X-3scale-debug"] == ngx.var.provider_key then
        local resp = ""
        ngx.ctx.buffered = (ngx.ctx.buffered or "") .. string.sub(ngx.arg[1], 1, 1000)
        if ngx.arg[2] then
            resp = ngx.ctx.buffered
        end

        ngx.log(0, ngx.req.raw_header())
        ngx.log(0, resp)
    end
}
```

This snippet will add the following extra logging to the nginx error.log when the **X-3scale-debug** header is sent, e.g. **curl -v -H 'X-3scale-debug: YOUR_PROVIDER_KEY' -X GET "https://726e3b99.ngrok.com/api/contacts.json?access_token=7c6f24f5"**

This will produce the following log entries:

```
2016/05/05 14:24:33 [] 7238#0: *57 [lua] body_filter_by_lua:7: GET /api/contacts.json?
access_token=7c6f24f5 HTTP/1.1
Host: 726e3b99.ngrok.io
User-Agent: curl/7.43.0
Accept: */*
X-Forwarded-Proto: https
X-Forwarded-For: 2.139.235.79

    while sending to client, client: 127.0.0.1, server: pili-virtualbox, request: "GET /api/contacts.json?
access_token=7c6f24f5 HTTP/1.1", subrequest: "/threescale_authrep", upstream:
"https://54.83.62.94:443/transactions/oauth_authrep.xml?
provider_key=REDACTED&service_id=REDACTED&usage[hits]=1&access_token=7c6f24f5", host:
"726e3b99.ngrok.io"
2016/05/05 14:24:33 [] 7238#0: *57 [lua] body_filter_by_lua:8: <?xml version="1.0" encoding="UTF-
8"?><error code="access_token_invalid">access_token "7c6f24f5" is invalid: expired or never
defined</error> while sending to client, client: 127.0.0.1, server: pili-virtualbox, request: "GET
/api/contacts.json?access_token=7c6f24f5 HTTP/1.1", subrequest: "/threescale_authrep", upstream:
"https://54.83.62.94:443/transactions/oauth_authrep.xml?
provider_key=REDACTED&service_id=REDACTED&usage[hits]=1&access_token=7c6f24f5", host:
"726e3b99.ngrok.io"
```

The first entry (2016/05/05 14:24:33 [] 7238#0: *57 [lua] body_filter_by_lua:7:) prints out the request headers sent to 3scale, in this case: Host, User-Agent, Accept, X-Forwarded-Proto and X-Forwarded-For.

The second entry (2016/05/05 14:24:33 [] 7238#0: *57 [lua] body_filter_by_lua:8:) prints out the response from 3scale, in this case: **<error code="access_token_invalid">access_token "7c6f24f5" is invalid: expired or never defined</error>**.

Both will print out the original request (**GET /api/contacts.json?access_token=7c6f24f5**) and subrequest location (**/threescale_authrep**) as well as the upstream request (**upstream: "https://54.83.62.94:443/transactions/threescale_authrep.xml?provider_key=REDACTED&service_id=REDACTED&usage[hits]=1&access_token=7c6f24f5"**.)

This last value allows you to see which of the 3scale IPs have been resolved and also the exact request made to 3scale.

4.3.5. 3scale

4.3.5.1. 1. Is 3scale available?

Check [3scale's status page](#) or [@3scalestatus](#) on Twitter.

4.3.5.2. 2. Is 3scale returning an error?

It is also possible that 3scale is available but is returning an error to your API gateway which would prevent calls going through to your API. Try to make the authorization call directly in 3scale and check the response. If you get an error, check the [#troubleshooting-api-error-codes](#) [Error Codes] section to see what the issue is.

4.3.5.3. 3. Use the 3scale debug headers

You can also turn on the 3scale debug headers by making a call to your API with the **X-3scale-debug** header, example:

```
curl -v -X GET "https://api.example.com/endpoint?user_key" X-3scale-debug:
YOUR_PROVIDER_KEY
```

This will return the following headers with the API response:

```
X-3scale-matched-rules: /, /api/contacts.json
< X-3scale-credentials: access_token=TOKEN_VALUE
< X-3scale-usage: usage[hits]=2
< X-3scale-hostname: HOSTNAME_VALUE
```

4.3.5.4. 4. Check the integration errors

You can also check the integration errors on your Admin Portal to check for any issues reporting traffic to 3scale. See https://YOUR_DOMAIN-admin.3scale.net/apiconfig/errors.

One of the reasons for integration errors can be sending credentials in the headers with [underscores_in_headers](#) directive not enabled in server block.

4.3.6. Client API gateway

4.3.6.1. 1. Is the API gateway reachable from the public internet?

Try directing a browser to the IP address (or domain name) of your gateway server. If this fails, ensure that you have opened the firewall on the relevant ports.

4.3.6.2. 2. Is the API gateway reachable by the client?

If possible, try to connect to the API gateway from the client using one of the methods outlined earlier (telnet, curl, etc.) If the connection fails, the problem lies in the network between the two.

Otherwise, you should move on to troubleshooting the client making the calls to the API.

4.3.7. Client

4.3.7.1. 1. Test the same call using a different client

If a request is not returning the expected result, test with a different HTTP client. For example, if you are calling an API with a Java HTTP client and you see something wrong, cross-check with cURL.

You can also call the API through a proxy between the client and the gateway to capture the exact parameters and headers being sent by the client.

4.3.7.2. 2. Inspect the traffic sent by client

Use a tool like Wireshark to see the requests being made by the client. This will allow you to identify if the client is making calls to the API and the details of the request.

4.4. OTHER ISSUES

4.4.1. ActiveDocs issues

Sometimes calls that work when you call the API from the command line fail when going through ActiveDocs.

To enable ActiveDocs calls to work, we send these out through a proxy on our side. This proxy will add certain headers that can sometimes cause issues on the API if they are not expected. To identify if this is the case, try the following steps:

4.4.1.1. 1. Use `petstore.swagger.io`

Swagger provides a hosted swagger-ui at `petstore.swagger.io` which you can use to test your Swagger spec and API going through the latest version of swagger-ui. If both swagger-ui and ActiveDocs fail in the same way, you can rule out any issues with ActiveDocs or the ActiveDocs proxy and focus the troubleshooting on your own spec. Alternatively, you can check the swagger-ui GitHub repo for any known issues with the current version of swagger-ui.

4.4.1.2. 2. Check that firewall allows connections from ActiveDocs proxy

We recommend to not whitelist IP address for clients using your API. The ActiveDocs proxy uses floating IP addresses for high availability and there is currently no mechanism to notify of any changes to these IPs.

4.4.1.3. 3. Call the API with incorrect credentials

One way to identify whether the ActiveDocs proxy is working correctly is to call your API with invalid credentials. This will help you to confirm or rule out any problems with both the ActiveDocs proxy and your API gateway.

If you get a 403 code back from the API call (or from the code you have configured on your gateway for invalid credentials), the problem lies with your API because the calls are reaching your gateway.

4.4.1.4. 4. Compare calls

To identify any differences in headers and parameters between calls made from ActiveDocs versus outside of ActiveDocs, it can sometimes be helpful to run your calls through some service (APItools on premise, Runscope, etc.) that allows you to inspect and compare your HTTP calls before sending them to your API. This will allow you to identify any potential headers and/or parameters in the request that could be causing issues on your side.

4.5. APPENDIX

4.5.1. Logging in NGINX

For a comprehensive guide on this, see the [NGINX Logging and Monitoring](#) docs.

4.5.1.1. Enabling debugging log

To find out more about enabling debugging log, see the [NGINX debugging log documentation](#).

4.5.2. 3scale error codes

To double-check the error codes that are returned by the 3scale Service Management API endpoints, see the [3scale API Documentation](#) page by following these steps:

1. Click the question mark (?) icon, which is in the upper-right corner of the Admin Portal.
2. Choose **3scale API Docs**.

The following is a list HTTP response codes returned by 3scale, and the conditions under which they are returned:

- **400:** Bad request. This can be because of:
 - Invalid encoding
 - Payload too large
 - Content type is invalid (for POST calls). Valid values for the **Content-Type** header are: **application/x-www-form-urlencoded**, **multipart/form-data**, or empty header.
- **403:**
 - Credentials are not valid
 - Sending body data to 3scale for a GET request
- **404:** Non-existent entity referenced, such as applications, metrics, etc.
- **409:**

- Usage limits exceeded
- Application is not active
- Application key is invalid or missing (for **app_id/app_key** authentication method)
- Referrer is not allowed or missing (when referrer filters are enabled and required)
- **422**: Missing required parameters

Most of these error responses will also contain an XML body with a machine readable error category and a human readable explanation.

When using the standard API gateway configuration, any return code different from 200 provided by 3scale can result in a response to the client with one of the following codes:

- 403
- 404