Red Hat 3scale API Management 2.2

Infrastructure

Learn more about deploying Red Hat 3scale API Management on different platforms.

Last Updated: 2021-08-03

Red Hat 3scale API Management 2.2 Infrastructure

Learn more about deploying Red Hat 3scale API Management on different platforms.

Legal Notice

Abstract

This guide documents deployment and infrastructure management with Red Hat 3scale API Management 2.2.

Table of Contents

CHAPTER 1. UPGRADE 3SCALE 2.1 TO 2.2

Perform the steps in this document to upgrade Red Hat 3scale API Management on-premises deployment from version 2.1 to 2.2.

1.1. PREREQUISITES:

- 3scale On-Premises 2.1
- OpenShift CLI
- 3scale AMP 2.2 templates
- Access and permissions to your OpenShift server and project
- A persistent volume with enough space to hold a backup of the MySQL database

WARNING

This process can cause disruption in the service. Make sure to have a maintenance window.

1.2. SELECT THE PROJECT

1. Take a backup of your OpenShift cluster.
2. From a terminal session, log in to your OpenShift cluster:

```
oc login https://<YOUR_OPENSHIFT_CLUSTER>:8443
```

3. Select the project that you want to upgrade:

```
oc project <YOUR_AMP_21_PROJECT>
```

1.3. GATHER THE NEEDED VALUES

You need the following parameters for the new 3scale API Management 2.2 Multitenancy feature. You may choose to either specify new values for these parameters or keep the default ones.

Parameter	Description
MASTER_USER	Username for the Master Admin Portal. Default: "master"
MASTER_PASSWORD	Password for the Master Admin Portal. Automatically generated if not specified

Parameter	Description
MASTER_ACCESS_TOKEN	Access Token for master automatically generated during the upgrade. However, MASTER_ACCESS_TOKEN is not added to the system seed by default.
APICAST_REGISTRY_URL	The URL to point to APIcast policies registry management. Default: http://apicast-staging:8090/policies

1. Gather the following values from the system components of your current 2.1 deployment:
   - DATABASE_URL
   - THREESCALE_SUPERDOMAIN
   - TENANT_NAME
   - APICAST_ACCESS_TOKEN
   - ADMIN_ACCESS_TOKEN
   - USER_LOGIN
   - USER_PASSWORD
   - EVENTS_SHARED_SECRET
   - APICAST_BACKEND_ROOT_ENDPOINT
   - CONFIG_INTERNAL_API_USER
   - CONFIG_INTERNAL_API_PASSWORD
   - SECRET_KEY_BASE
   - BACKEND_ROUTE
2. Export them from the current deployment into the active shell:

```
export `oc env dc/system-app --list | grep -E
'^(DATABASE_URL|THREESCALE_SUPERDOMAIN|TENANT_NAME|APICAST_ACCESS_
TOKEN|ADMIN_ACCESS_TOKEN|USER_LOGIN|USER_PASSWORD|EVENTS_SHARED_
SECRET|APICAST_BACKEND_ROOT_ENDPOINT|CONFIG_INTERNAL_API_USER|CON
FIG_INTERNAL_API_PASSWORD|SECRET_KEY_BASE|BACKEND_ROUTE)=' | tr "\n" ' ' `
```

3. Optionally, to query individual values from the OpenShift CLI, run the following **oc get** command, where **<variable_name>** is the name of the variable you want to query:

```
oc get "-o=custom-columns=NAMES:.spec.template.spec.containers[0].env[?(.name==\"
<variable_name>\")].value" dc/system-app
```

4. Gather the following values from the **system-mysql** component of your current 2.1 deployment:
   - MYSQL_USER
   - MYSQL_PASSWORD
   - MYSQL_DATABASE
   - MYSQL_ROOT_PASSWORD
5. Export these values from the current deployment into the active shell:

```
export `oc env dc/system-mysql --list | grep -E
'^(MYSQL_USER|MYSQL_PASSWORD|MYSQL_DATABASE|MYSQL_ROOT_PASSWORD)
=' | tr "\n" ' ' `
```

6. Optionally, to query individual values from the OpenShift CLI, run the following **oc get** command, where **<variable_name>** is the name of the variable you want to query:

```
oc get "-o=custom-columns=NAMES:.spec.template.spec.containers[0].env[?(.name==\"
<variable_name>\")].value" dc/system-mysql
```

7. Gather the following values from the APIcast component of your current 2.1 deployment:
   - APICAST_MANAGEMENT_API
   - OPENSSL_VERIFY
   - APICAST_RESPONSE_CODES
8. Export these values from the current deployment into the active shell:

```
export `oc env dc/apicast-production --list | grep -E
'^(APICAST_MANAGEMENT_API|OPENSSL_VERIFY|APICAST_RESPONSE_CODES)=' |
tr "\n" ' ' `
```

9. Optionally, to query individual values from the OpenShift CLI, run the following **oc get** command, where **<variable_name>** is the name of the variable you want to query:

```
oc get "-o=custom-columns=NAMES:.spec.template.spec.containers[0].env[?(.name==\"
<variable_name>\")].value" dc/apicast-production
```

1.4. CONFIGURE NEW VARIABLE VALUES

1. Set the value for the new version of the AMP release:

```
export AMP_RELEASE=2.2.0
```

2. Set the values for the new **optional** parameters introduced in AMP 2.2. These parameters are described in the beginning of the gather needed values section. Use the **export** command, replacing the values in the parenthesis:

```
export MASTER_ACCESS_TOKEN=<MASTER_ACCESS_TOKEN>
export APICAST_REGISTRY_URL=<APICAST_REGISTRY_URL>
```

Regarding MASTER_USER and MASTER_PASSWORD, consider the following:

- If you want to use the default values, no action is required.
- If you have specified values for these environment variables, export them with the following commands:

```
export MASTER_USER=<MASTER_USER>
export MASTER_PASSWORD=<MASTER_PASSWORD>
```

3. Confirm that the necessary values gathered in the gather needed values section are exported to the active shell and that the new values are set in this section:

```
echo AMP_RELEASE=$AMP_RELEASE

echo DATABASE_URL=$DATABASE_URL
echo THREESCALE_SUPERDOMAIN=$THREESCALE_SUPERDOMAIN
echo TENANT_NAME=$TENANT_NAME
echo APICAST_ACCESS_TOKEN=$APICAST_ACCESS_TOKEN
echo ADMIN_ACCESS_TOKEN=$ADMIN_ACCESS_TOKEN
echo USER_LOGIN=$USER_LOGIN
echo USER_PASSWORD=$USER_PASSWORD
echo EVENTS_SHARED_SECRET=$EVENTS_SHARED_SECRET
echo
APICAST_BACKEND_ROOT_ENDPOINT=$APICAST_BACKEND_ROOT_ENDPOINT
echo CONFIG_INTERNAL_API_USER=$CONFIG_INTERNAL_API_USER
echo CONFIG_INTERNAL_API_PASSWORD=$CONFIG_INTERNAL_API_PASSWORD
echo SECRET_KEY_BASE=$SECRET_KEY_BASE
echo BACKEND_ROUTE=$BACKEND_ROUTE

echo MYSQL_USER=$MYSQL_USER
echo MYSQL_PASSWORD=$MYSQL_PASSWORD
echo MYSQL_DATABASE=$MYSQL_DATABASE
echo MYSQL_ROOT_PASSWORD=$MYSQL_ROOT_PASSWORD

echo APICAST_MANAGEMENT_API=$APICAST_MANAGEMENT_API
echo OPENSSL_VERIFY=$OPENSSL_VERIFY
echo APICAST_RESPONSE_CODES=$APICAST_RESPONSE_CODES

echo MASTER_USER=$MASTER_USER
echo MASTER_PASSWORD=$MASTER_PASSWORD
echo MASTER_ACCESS_TOKEN=$MASTER_ACCESS_TOKEN
echo APICAST_REGISTRY_URL=$APICAST_REGISTRY_URL
```

1.5. UPGRADE THE DATABASE POD

To upgrade the database, create a backup of the pod and deploy a new pod.

1.5.1. Create the backup

1. Create a persistent volume with enough storage to hold the MySQL database.
2. To create a persistent volume claim with enough storage to hold the MySQL database, run the following command, replacing the **<size>** value with an appropriate size for your database:

```
echo "apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysql-backup
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: <size>" | oc create -f -
```

This command crates a persistent volume claim named **mysql-backup**.

3. To create a pod to house the backup database, run the following command:

```
echo "apiVersion: v1
kind: Pod
metadata:
  name: mysql-backup
  labels:
    name: mysql-backup
spec:
  containers:
  - name: mysql-backup
    image: registry.access.redhat.com/rhscl/mysql-57-rhel7:5.7-5
    args:
    - sleep
    - infinity
    volumeMounts:
      - mountPath: /backup
        name: mysql-backup
  volumes:
  - name: mysql-backup
    persistentVolumeClaim:
      claimName: mysql-backup" | oc create -f -
```

4. Wait until the pod is created and log in to it using the following **oc rsh** command and take the backup:

```
oc rsh mysql-backup /opt/rh/rh-mysql57/root/usr/bin/mysqldump -h system-mysql -u
${MYSQL_USER} -p${MYSQL_PASSWORD} system -r /backup/backup.sql
```

5. To verify the contents of your backup, check the backup file size against the original file size by running the following **oc rsh** command:

```
oc rsh mysql-backup ls -lha /backup/backup.sql
```

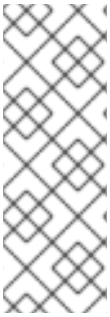

NOTE

If in the next steps you experience a failure, you can redeploy your database from your backup by using the following **oc rsh** command:

```
oc rsh mysql-backup /bin/bash -c "/usr/bin/cat /backup/backup.sql | /opt/rh/rh-mysql57/root/usr/bin/mysql -h system-mysql -uroot -p${MYSQL_ROOT_PASSWORD} system"
```

1.5.2. Perform the upgrade

1. Delete the system service with the following **oc delete** command. Note that your application will be down from this point because the application loses connection with the database:

```
oc delete service system-mysql
```

2. To patch the MySQL DeploymentConfig, run the following **oc patch** command:

```
oc patch dc/system-mysql -p "spec:
  template:
    spec:
      containers:
      - name: system-mysql
        image: registry.access.redhat.com/rhscl/mysql-57-rhel7:5.7-5
        args:
          - /opt/rh/rh-mysql57/root/usr/libexec/mysqld
          - '--datadir=/var/lib/mysql/data/'
"
```

3. Ensure that the new pod is deployed successfully before continuing.

4. To fetch the new pod details, run the following **oc get** commands:

   a. To fetch the pod name:

```
oc get pods -l deploymentconfig=system-mysql
```

   b. To fetch the pod IP address:

```
oc get pods  -o=custom-columns=IP:.status.podIP -l deploymentconfig=system-mysql
```

5. Log in to the pod using the following **oc rsh** command, substitute the <pod_name> and <pod_ip> with the name and IP address from the preceding steps:

```
oc rsh <pod_name> /opt/rh/rh-mysql57/root/usr/bin/mysql_upgrade -h <pod_ip> -u root -p${MYSQL_ROOT_PASSWORD}
```

6. To patch the mysql Deployment Config with the changes done on the 2.2 version, run the following **oc patch** command:

```
oc patch dc/system-mysql -p "
metadata:
  labels:
```

```
    app: System
spec:
  template:
    spec:
      containers:
      - name: system-mysql
        args:
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            memory: 2Gi
          requests:
            cpu: 250m
            memory: 512Mi
"
```

1.5.3. Change the MySQL character set and collation

1. To create the **mysql-extra-conf** ConfigMap with the mysql-charset.cnf config file, run the following **oc create** command:

```
echo "kind: ConfigMap
apiVersion: v1
metadata:
  name: mysql-extra-conf
data:
  mysql-charset.cnf: |
    [client]
    default-character-set = utf8

    [mysql]
    default-character-set = utf8

    [mysqld]
    character-set-server = utf8
    collation-server = utf8_unicode_ci" | oc create -f -
```

2. To create the **mysql-main-conf** ConfigMap, run the following **oc create** command:

```
echo 'kind: ConfigMap
apiVersion: v1
metadata:
  name: mysql-main-conf
data:
  my.cnf: |
    !include /etc/my.cnf
    !includedir /etc/my-extra.d' | oc create -f -
```

3. To configure system-mysql starting with the created configmaps on last steps, run the following **oc patch** command:

```
oc patch dc/system-mysql -p "spec:
  template:
    spec:
```

```
      containers:
        - name: system-mysql
          env:
            - name: MYSQL_USER
              value: "${MYSQL_USER}"
            - name: MYSQL_PASSWORD
              value: "${MYSQL_PASSWORD}"
            - name: MYSQL_DATABASE
              value: "${MYSQL_DATABASE}"
            - name: MYSQL_ROOT_PASSWORD
              value: "${MYSQL_ROOT_PASSWORD}"
            - name: MYSQL_LOWER_CASE_TABLE_NAMES
              value: '1'
            - name: MYSQL_DEFAULTS_FILE
              value: "/etc/my-extra/my.cnf"
          volumeMounts:
            - name: 'mysql-storage'
              mountPath: /var/lib/mysql/data
            - name: 'mysql-extra-conf'
              mountPath: /etc/my-extra.d
            - name: 'mysql-main-conf'
              mountPath: /etc/my-extra
      volumes:
        - name: 'mysql-storage'
          persistentVolumeClaim:
            claimName: 'mysql-storage'
        - name: 'mysql-extra-conf'
          configMap:
            name: 'mysql-extra-conf'
        - name: 'mysql-main-conf'
          configMap:
            name: 'mysql-main-conf'
"
```

4. Wait until the deployment is complete on system-mysql and run the following **oc get** command to fetch the new MySQL pod name:

```
oc get pods -l deploymentconfig=system-mysql
```

5. To fetch the new MySQL pod IP address, run the following **oc get** command:

```
oc get pods  -o=custom-columns=IP:.status.podIP -l deploymentconfig=system-mysql
```

6. To change the character set on the database and all the tables, run the **oc rsh** command, specifying the previously fetched **<pod_name>** and **<pod_ip>**:

```
oc rsh <pod_name> /bin/bash -c "echo ALTER DATABASE system CHARACTER SET utf8
COLLATE utf8_general_ci | mysql -h <pod_ip> -u root -p${MYSQL_ROOT_PASSWORD} --
default-character-set=utf8"
oc rsh <pod_name> /bin/bash -c "/opt/rh/rh-mysql57/root/usr/bin/mysql -h <pod_ip> -u root -
p${MYSQL_ROOT_PASSWORD} --default-character-set=utf8 -B -N -e 'SHOW TABLES'
system | awk '{print \"SET foreign_key_checks = 0; ALTER TABLE\", \$1, \"CONVERT TO
CHARACTER SET utf8 COLLATE utf8_general_ci; SET foreign_key_checks = 1; \"}' |
/opt/rh/rh-mysql57/root/usr/bin/mysql -h <pod_ip> -u root -p${MYSQL_ROOT_PASSWORD}
--default-character-set=utf8 system"
```

7. To create the **system-mysql** service, run the following **oc create** command:

```
echo "kind: Service
apiVersion: v1
metadata:
  name: 'system-mysql'
spec:
  ports:
    - name: system-mysql
      protocol: TCP
      port: 3306
      targetPort: 3306
      nodePort: 0
  selector:
    name: 'system-mysql'" | oc create -f -
```

1.5.4. Delete the backups

1. Verify the updated database and ensure that the pods are running.
2. To delete the backup pod and persistent volume claim, run the following **oc delete** command:

```
oc delete pod/mysql-backup
oc delete pvc/mysql-backup
```

1.6. CREATE NEW ROUTES AND SERVICES FOR SYSTEM

After you have configured the new variable values, run the following **oc create** command to create new routes and services:

```
echo "
apiVersion: v1
kind: Service
metadata:
  name: system-master
  annotations:
    service.alpha.openshift.io/dependencies: '[{\"name\": \"system-developer\", \"kind\": \"Service\"}]'
spec:
  ports:
  - port: 3000
    protocol: TCP
    targetPort: master
    name: http
  selector:
    name: system-app
" | oc create -f -
```

```
echo "
apiVersion: v1
kind: Route
metadata:
  name: system-master-admin-route
spec:
```

```
  host: master-account-admin.${THREESCALE_SUPERDOMAIN}
  to:
    kind: Service
    name: system-master
  port:
    targetPort: http
  tls:
    termination: edge
    insecureEdgeTerminationPolicy: Allow
" | oc create -f -
```

1.7. PATCH SYSTEM COMPONENTS

Continue your in-place upgrade using the **oc patch** command. The **oc patch** command allows you to patch your deployment configurations, image streams and ConfigMaps.

In this section of the upgrade, you must patch the system config map. You must also patch deployment configurations for the following pods:

- system-app
- system-resque
- system-sidekiq
- system-sphinx

Follow these steps to patch config maps and deployment configurations:

1. To patch the **system** ConfigMap, run the following **oc patch** command:

```
oc patch cm/system -p "
data:
  zync.yml: |
    production:
      endpoint: 'http://zync:8080'
      authentication:
        token: \"<%= ENV.fetch('ZYNC_AUTHENTICATION_TOKEN') %>\"
      connect_timeout: 5
      send_timeout: 5
      receive_timeout: 10
      root_url:
  rolling_updates.yml: |
    production:
      old_charts: false
      new_provider_documentation: false
      proxy_pro: false
      instant_bill_plan_change: false
      service_permissions: true
      async_apicast_deploy: false
      duplicate_application_id: true
      duplicate_user_key: true
      plan_changes_wizard: false
      require_cc_on_signup: false
      apicast_per_service: true
      new_notification_system: true
```

```
cms_api: false
apicast_v2: true
forum: false
published_service_plan_signup: true
apicast_oidc: true
policies: true"
```

2. To patch the **system-resque** deployment configuration, consider the following:

   - If you want to use the default values for MASTER_USER and MASTER_PASSWORD environment variables, do not describe them in the **oc patch** command below.
   - Alternatively, if you want to specify values for MASTER_USER and MASTER_PASSWORD, include them in the **oc patch** command below.

```
oc patch dc/system-resque -p "
metadata:
  labels:
    app: System
spec:
  template:
    spec:
      containers:
      - env:
        - name: RAILS_ENV
          value: \"production\"
        - name: DATABASE_URL
          value: \"${DATABASE_URL}\"
        - name: FORCE_SSL
          value: \"true\"
        - name: THREESCALE_SUPERDOMAIN
          value: \"${THREESCALE_SUPERDOMAIN}\"
        - name: MASTER_USER
          value: \"${MASTER_USER}\"
        - name: MASTER_PASSWORD
          value: \"${MASTER_PASSWORD}\"
        - name: TENANT_NAME
          value: \"${TENANT_NAME}\"
        - name: APICAST_ACCESS_TOKEN
          value: \"${APICAST_ACCESS_TOKEN}\"
        - name: ADMIN_ACCESS_TOKEN
          value: \"${ADMIN_ACCESS_TOKEN}\"
        - name: PROVIDER_PLAN
          value: 'enterprise'
        - name: USER_LOGIN
          value: \"${USER_LOGIN}\"
        - name: USER_PASSWORD
          value: \"${USER_PASSWORD}\"
        - name: RAILS_LOG_TO_STDOUT
          value: \"true\"
        - name: RAILS_LOG_LEVEL
          value: \"info\"
        - name: THINKING_SPHINX_ADDRESS
          value: \"system-sphinx\"
        - name: THINKING_SPHINX_PORT
          value: \"9306\"
```

```
- name: THINKING_SPHINX_CONFIGURATION_FILE
  value: \"/tmp/sphinx.conf\"
- name: EVENTS_SHARED_SECRET
  value: \"${EVENTS_SHARED_SECRET}\"
- name: THREESCALE_SANDBOX_PROXY_OPENSSL_VERIFY_MODE
  value: \"VERIFY_NONE\"
- name: APICAST_BACKEND_ROOT_ENDPOINT
  value: \"${APICAST_BACKEND_ROOT_ENDPOINT}\"
- name: CONFIG_INTERNAL_API_USER
  value: \"${CONFIG_INTERNAL_API_USER}\"
- name: CONFIG_INTERNAL_API_PASSWORD
  value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
- name: SECRET_KEY_BASE
  value: \"${SECRET_KEY_BASE}\"
- name: AMP_RELEASE
  value: \"${AMP_RELEASE}\"
- name: ZYNC_AUTHENTICATION_TOKEN
  valueFrom:
    secretKeyRef:
      name: zync
      key: ZYNC_AUTHENTICATION_TOKEN
- name: SMTP_ADDRESS
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: address
- name: SMTP_USER_NAME
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: username
- name: SMTP_PASSWORD
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: password
- name: SMTP_DOMAIN
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: domain
- name: SMTP_PORT
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: port
- name: SMTP_AUTHENTICATION
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: authentication
- name: SMTP_OPENSSL_VERIFY_MODE
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: openssl.verify.mode
```

```
    - name: BACKEND_ROUTE
      value: \"${BACKEND_ROUTE}\"
    - name: SSL_CERT_DIR
      value: \"/etc/pki/tls/certs\"
    - name: APICAST_REGISTRY_URL
      value: \"${APICAST_REGISTRY_URL}\"
    image: registry.access.redhat.com/3scale-amp22/system:1.7
    imagePullPolicy: IfNotPresent
    name: system-resque
    resources:
      limits:
        cpu: 150m
        memory: 450Mi
      requests:
        cpu: 100m
        memory: 300Mi
  - env:
    - name: RAILS_ENV
      value: \"production\"
    - name: DATABASE_URL
      value: \"${DATABASE_URL}\"
    - name: FORCE_SSL
      value: \"true\"
    - name: THREESCALE_SUPERDOMAIN
      value: \"${THREESCALE_SUPERDOMAIN}\"
    - name: MASTER_USER
      value: \"${MASTER_USER}\"
    - name: MASTER_PASSWORD
      value: \"${MASTER_PASSWORD}\"
    - name: TENANT_NAME
      value: \"${TENANT_NAME}\"
    - name: APICAST_ACCESS_TOKEN
      value: \"${APICAST_ACCESS_TOKEN}\"
    - name: ADMIN_ACCESS_TOKEN
      value: \"${ADMIN_ACCESS_TOKEN}\"
    - name: PROVIDER_PLAN
      value: 'enterprise'
    - name: USER_LOGIN
      value: \"${USER_LOGIN}\"
    - name: USER_PASSWORD
      value: \"${USER_PASSWORD}\"
    - name: RAILS_LOG_TO_STDOUT
      value: \"true\"
    - name: RAILS_LOG_LEVEL
      value: \"info\"
    - name: THINKING_SPHINX_ADDRESS
      value: \"system-sphinx\"
    - name: THINKING_SPHINX_PORT
      value: \"9306\"
    - name: THINKING_SPHINX_CONFIGURATION_FILE
      value: \"/tmp/sphinx.conf\"
    - name: EVENTS_SHARED_SECRET
      value: \"${EVENTS_SHARED_SECRET}\"
    - name: THREESCALE_SANDBOX_PROXY_OPENSSL_VERIFY_MODE
      value: \"VERIFY_NONE\"
    - name: APICAST_BACKEND_ROOT_ENDPOINT
```

```
    value: \"${APICAST_BACKEND_ROOT_ENDPOINT}\"
  - name: CONFIG_INTERNAL_API_USER
    value: \"${CONFIG_INTERNAL_API_USER}\"
  - name: CONFIG_INTERNAL_API_PASSWORD
    value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
  - name: SECRET_KEY_BASE
    value: \"${SECRET_KEY_BASE}\"
  - name: AMP_RELEASE
    value: \"${AMP_RELEASE}\"
  - name: ZYNC_AUTHENTICATION_TOKEN
    valueFrom:
      secretKeyRef:
        name: zync
        key: ZYNC_AUTHENTICATION_TOKEN
  - name: SMTP_ADDRESS
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: address
  - name: SMTP_USER_NAME
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: username
  - name: SMTP_PASSWORD
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: password
  - name: SMTP_DOMAIN
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: domain
  - name: SMTP_PORT
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: port
  - name: SMTP_AUTHENTICATION
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: authentication
  - name: SMTP_OPENSSL_VERIFY_MODE
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: openssl.verify.mode
  - name: BACKEND_ROUTE
    value: \"${BACKEND_ROUTE}\"
  - name: SSL_CERT_DIR
    value: \"/etc/pki/tls/certs\"
  - name: APICAST_REGISTRY_URL
    value: \"${APICAST_REGISTRY_URL}\"
  image: registry.access.redhat.com/3scale-amp22/system:1.7
```

```
      imagePullPolicy: IfNotPresent
      name: system-scheduler
      resources:
        limits:
          cpu: 150m
          memory: 250Mi
        requests:
          cpu: 50m
          memory: 200Mi
"
```

3. To patch the **system-sidekiq** deployment configuration, consider the following:

   - If you want to use the default values for MASTER_USER and MASTER_PASSWORD environment variables, do not describe them in the **oc patch** command below.

   - Alternatively, if you want to specify values for MASTER_USER and MASTER_PASSWORD, include them in the **oc patch** command below.

```
oc patch dc/system-sidekiq -p "
spec:
  template:
    spec:
      containers:
      - name: system-sidekiq
        volumeMounts:
"
```

```
oc patch dc/system-sidekiq -p "
metadata:
  labels:
    app: System
spec:
  template:
    spec:
      containers:
      - env:
        - name: RAILS_ENV
          value: \"production\"
        - name: DATABASE_URL
          value: \"${DATABASE_URL}\"
        - name: FORCE_SSL
          value: \"true\"
        - name: THREESCALE_SUPERDOMAIN
          value: \"${THREESCALE_SUPERDOMAIN}\"
        - name: MASTER_USER
          value: \"${MASTER_USER}\"
        - name: MASTER_PASSWORD
          value: \"${MASTER_PASSWORD}\"
        - name: TENANT_NAME
          value: \"${TENANT_NAME}\"
        - name: APICAST_ACCESS_TOKEN
          value: \"${APICAST_ACCESS_TOKEN}\"
        - name: ADMIN_ACCESS_TOKEN
          value: \"${ADMIN_ACCESS_TOKEN}\"
        - name: PROVIDER_PLAN
```

```
    value: 'enterprise'
  - name: USER_LOGIN
    value: \"${USER_LOGIN}\"
  - name: USER_PASSWORD
    value: \"${USER_PASSWORD}\"
  - name: RAILS_LOG_TO_STDOUT
    value: \"true\"
  - name: RAILS_LOG_LEVEL
    value: \"info\"
  - name: THINKING_SPHINX_ADDRESS
    value: \"system-sphinx\"
  - name: THINKING_SPHINX_PORT
    value: \"9306\"
  - name: THINKING_SPHINX_CONFIGURATION_FILE
    value: \"/tmp/sphinx.conf\"
  - name: EVENTS_SHARED_SECRET
    value: \"${EVENTS_SHARED_SECRET}\"
  - name: THREESCALE_SANDBOX_PROXY_OPENSSL_VERIFY_MODE
    value: \"VERIFY_NONE\"
  - name: APICAST_BACKEND_ROOT_ENDPOINT
    value: \"${APICAST_BACKEND_ROOT_ENDPOINT}\"
  - name: CONFIG_INTERNAL_API_USER
    value: \"${CONFIG_INTERNAL_API_USER}\"
  - name: CONFIG_INTERNAL_API_PASSWORD
    value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
  - name: SECRET_KEY_BASE
    value: \"${SECRET_KEY_BASE}\"
  - name: AMP_RELEASE
    value: \"${AMP_RELEASE}\"
  - name: ZYNC_AUTHENTICATION_TOKEN
    valueFrom:
      secretKeyRef:
        name: zync
        key: ZYNC_AUTHENTICATION_TOKEN
  - name: SMTP_ADDRESS
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: address
  - name: SMTP_USER_NAME
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: username
  - name: SMTP_PASSWORD
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: password
  - name: SMTP_DOMAIN
    valueFrom:
      configMapKeyRef:
        name: smtp
        key: domain
  - name: SMTP_PORT
    valueFrom:
```

```
          configMapKeyRef:
            name: smtp
            key: port
        - name: SMTP_AUTHENTICATION
          valueFrom:
            configMapKeyRef:
              name: smtp
              key: authentication
        - name: SMTP_OPENSSL_VERIFY_MODE
          valueFrom:
            configMapKeyRef:
              name: smtp
              key: openssl.verify.mode
        - name: BACKEND_ROUTE
          value: \"${BACKEND_ROUTE}\"
        - name: SSL_CERT_DIR
          value: \"/etc/pki/tls/certs\"
        - name: APICAST_REGISTRY_URL
          value: \"${APICAST_REGISTRY_URL}\"
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        volumeMounts:
        - name: system-storage
          mountPath: /opt/system/public/system
        - name: system-config
          mountPath: /opt/system-extra-configs
        - name: system-tmp
          mountPath: /tmp
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        imagePullPolicy: IfNotPresent
        name: system-sidekiq
        resources:
          limits:
            cpu: 1000m
            memory: 2Gi
          requests:
            cpu: 100m
            memory: 500Mi
      volumes:
        - name: system-tmp
          emptyDir:
            medium: Memory
        - name: system-storage
          persistentVolumeClaim:
            claimName: system-storage
        - name: system-config
          configMap:
            name: system
            items:
            - key: zync.yml
              path: zync.yml
            - key: rolling_updates.yml
              path: rolling_updates.yml
"
```

4. To patch the **system-app** deployment configuration, consider the following:

- If you want to use the default values for MASTER_USER and MASTER_PASSWORD environment variables, do not describe them in the **oc patch** command below.

- Alternatively, if you want to specify values for MASTER_USER and MASTER_PASSWORD, include them in the **oc patch** command below.

```
oc patch dc/system-app -p "
spec:
  template:
    spec:
      containers:
      - name: system-provider
        volumeMounts:
      - name: system-developer
        volumeMounts:
"
```

```
oc patch dc/system-app -p "
metadata:
  labels:
    app: System
spec:
  strategy:
    rollingParams:
      pre:
        execNewPod:
          containerName: system-master
          env:
          - name: RAILS_ENV
            value: \"production\"
          - name: DATABASE_URL
            value: \"${DATABASE_URL}\"
          - name: FORCE_SSL
            value: \"true\"
          - name: THREESCALE_SUPERDOMAIN
            value: \"${THREESCALE_SUPERDOMAIN}\"
          - name: MASTER_USER
            value: \"${MASTER_USER}\"
          - name: MASTER_PASSWORD
            value: \"${MASTER_PASSWORD}\"
          - name: TENANT_NAME
            value: \"${TENANT_NAME}\"
          - name: APICAST_ACCESS_TOKEN
            value: \"${APICAST_ACCESS_TOKEN}\"
          - name: ADMIN_ACCESS_TOKEN
            value: \"${ADMIN_ACCESS_TOKEN}\"
          - name: PROVIDER_PLAN
            value: 'enterprise'
          - name: USER_LOGIN
            value: \"${USER_LOGIN}\"
          - name: USER_PASSWORD
            value: \"${USER_PASSWORD}\"
          - name: RAILS_LOG_TO_STDOUT
            value: \"true\"
          - name: RAILS_LOG_LEVEL
            value: \"info\"
```

```
- name: THINKING_SPHINX_ADDRESS
  value: \"system-sphinx\"
- name: THINKING_SPHINX_PORT
  value: \"9306\"
- name: THINKING_SPHINX_CONFIGURATION_FILE
  value: \"/tmp/sphinx.conf\"
- name: EVENTS_SHARED_SECRET
  value: \"${EVENTS_SHARED_SECRET}\"
- name: THREESCALE_SANDBOX_PROXY_OPENSSL_VERIFY_MODE
  value: \"VERIFY_NONE\"
- name: APICAST_BACKEND_ROOT_ENDPOINT
  value: \"${APICAST_BACKEND_ROOT_ENDPOINT}\"
- name: CONFIG_INTERNAL_API_USER
  value: \"${CONFIG_INTERNAL_API_USER}\"
- name: CONFIG_INTERNAL_API_PASSWORD
  value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
- name: SECRET_KEY_BASE
  value: \"${SECRET_KEY_BASE}\"
- name: AMP_RELEASE
  value: \"${AMP_RELEASE}\"
- name: ZYNC_AUTHENTICATION_TOKEN
  valueFrom:
    secretKeyRef:
      name: zync
      key: ZYNC_AUTHENTICATION_TOKEN
- name: SMTP_ADDRESS
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: address
- name: SMTP_USER_NAME
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: username
- name: SMTP_PASSWORD
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: password
- name: SMTP_DOMAIN
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: domain
- name: SMTP_PORT
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: port
- name: SMTP_AUTHENTICATION
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: authentication
- name: SMTP_OPENSSL_VERIFY_MODE
```

```
            valueFrom:
              configMapKeyRef:
                name: smtp
                key: openssl.verify.mode
            - name: BACKEND_ROUTE
              value: \"${BACKEND_ROUTE}\"
            - name: SSL_CERT_DIR
              value: \"/etc/pki/tls/certs\"
            - name: APICAST_REGISTRY_URL
              value: \"${APICAST_REGISTRY_URL}\"
            command:
            - bash
            - -c
            - bundle exec rake boot openshift:deploy
MASTER_ACCESS_TOKEN="${MASTER_ACCESS_TOKEN}"
      post:
        execNewPod:
          containerName: system-master
  template:
    spec:
      containers:
      - args:
        env:
        - name: RAILS_ENV
          value: \"production\"
        - name: DATABASE_URL
          value: \"${DATABASE_URL}\"
        - name: FORCE_SSL
          value: \"true\"
        - name: THREESCALE_SUPERDOMAIN
          value: \"${THREESCALE_SUPERDOMAIN}\"
        - name: MASTER_USER
          value: \"${MASTER_USER}\"
        - name: MASTER_PASSWORD
          value: \"${MASTER_PASSWORD}\"
        - name: TENANT_NAME
          value: \"${TENANT_NAME}\"
        - name: APICAST_ACCESS_TOKEN
          value: \"${APICAST_ACCESS_TOKEN}\"
        - name: ADMIN_ACCESS_TOKEN
          value: \"${ADMIN_ACCESS_TOKEN}\"
        - name: PROVIDER_PLAN
          value: 'enterprise'
        - name: USER_LOGIN
          value: \"${USER_LOGIN}\"
        - name: USER_PASSWORD
          value: \"${USER_PASSWORD}\"
        - name: RAILS_LOG_TO_STDOUT
          value: \"true\"
        - name: RAILS_LOG_LEVEL
          value: \"info\"
        - name: THINKING_SPHINX_ADDRESS
          value: \"system-sphinx\"
        - name: THINKING_SPHINX_PORT
          value: \"9306\"
        - name: THINKING_SPHINX_CONFIGURATION_FILE
```

```
  value: \"/tmp/sphinx.conf\"
- name: EVENTS_SHARED_SECRET
  value: \"${EVENTS_SHARED_SECRET}\"
- name: THREESCALE_SANDBOX_PROXY_OPENSSL_VERIFY_MODE
  value: \"VERIFY_NONE\"
- name: APICAST_BACKEND_ROOT_ENDPOINT
  value: \"${APICAST_BACKEND_ROOT_ENDPOINT}\"
- name: CONFIG_INTERNAL_API_USER
  value: \"${CONFIG_INTERNAL_API_USER}\"
- name: CONFIG_INTERNAL_API_PASSWORD
  value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
- name: SECRET_KEY_BASE
  value: \"${SECRET_KEY_BASE}\"
- name: AMP_RELEASE
  value: \"${AMP_RELEASE}\"
- name: ZYNC_AUTHENTICATION_TOKEN
  valueFrom:
    secretKeyRef:
      name: zync
      key: ZYNC_AUTHENTICATION_TOKEN
- name: SMTP_ADDRESS
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: address
- name: SMTP_USER_NAME
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: username
- name: SMTP_PASSWORD
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: password
- name: SMTP_DOMAIN
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: domain
- name: SMTP_PORT
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: port
- name: SMTP_AUTHENTICATION
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: authentication
- name: SMTP_OPENSSL_VERIFY_MODE
  valueFrom:
    configMapKeyRef:
      name: smtp
      key: openssl.verify.mode
- name: BACKEND_ROUTE
```

```
          value: \"${BACKEND_ROUTE}\"
        - name: SSL_CERT_DIR
          value: \"/etc/pki/tls/certs\"
        - name: APICAST_REGISTRY_URL
          value: \"${APICAST_REGISTRY_URL}\"
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        imagePullPolicy: IfNotPresent
        args: [ 'env', 'TENANT_MODE=master', 'PORT=3002', 'container-entrypoint',
'bundle', 'exec', 'unicorn', '-c', 'config/unicorn.rb' ]
        command:
        name: system-master
        resources:
          limits:
            cpu: 1000m
            memory: 800Mi
          requests:
            cpu: 50m
            memory: 600Mi
        livenessProbe:
          timeoutSeconds: 10
          initialDelaySeconds: 20
          tcpSocket:
            port: master
          periodSeconds: 10
        readinessProbe:
          httpGet:
            path: /check.txt
            port: master
            scheme: HTTP
            httpHeaders:
            - name: X-Forwarded-Proto
              value: https
          initialDelaySeconds: 30
          timeoutSeconds: 10
          periodSeconds: 30
        ports:
        - containerPort: 3002
          protocol: TCP
          name: master
        volumeMounts:
        - name: system-storage
          mountPath: /opt/system/public/system
        - name: system-config
          mountPath: /opt/system-extra-configs
      - name: system-provider
        env:
        - name: MASTER_USER
          value: ${MASTER_USER}
        - name: MASTER_PASSWORD
          value: ${MASTER_PASSWORD}
        - name: AMP_RELEASE
          value: ${AMP_RELEASE}
        - name: APICAST_REGISTRY_URL
          value: ${APICAST_REGISTRY_URL}
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        imagePullPolicy: IfNotPresent
```

```
        resources:
          limits:
            cpu: 1000m
            memory: 800Mi
          requests:
            cpu: 50m
            memory: 600Mi
        command:
        args: [ 'env', 'TENANT_MODE=provider', 'PORT=3000', 'container-entrypoint',
'bundle', 'exec', 'unicorn', '-c', 'config/unicorn.rb' ]
        volumeMounts:
        - name: system-storage
          mountPath: /opt/system/public/system
        - name: system-config
          mountPath: /opt/system-extra-configs
      - name: system-developer
        env:
        - name: MASTER_USER
          value: ${MASTER_USER}
        - name: MASTER_PASSWORD
          value: ${MASTER_PASSWORD}
        - name: AMP_RELEASE
          value: ${AMP_RELEASE}
        - name: APICAST_REGISTRY_URL
          value: ${APICAST_REGISTRY_URL}
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        imagePullPolicy: IfNotPresent
        command:
        args: [ 'env', 'PORT=3001', 'container-entrypoint', 'bundle', 'exec', 'unicorn', '-c',
'config/unicorn.rb' ]
        volumeMounts:
        - name: system-storage
          readOnly: true
          mountPath: /opt/system/public/system
        - name: system-config
          mountPath: /opt/system-extra-configs
  triggers:
    - type: ConfigChange
    - type: ImageChange
      imageChangeParams:
        automatic: true
        containerNames:
        - system-provider
        - system-developer
        - system-master
        from:
          kind: ImageStreamTag
          name: amp-system:latest
"
```

5. To patch the **amp-system** image, run the following **oc patch** command:

```
oc patch is/amp-system -p "
spec:
  tags:
    - name: 2.2.0
```

```
      annotations:
        openshift.io/display-name: AMP system 2.2.0
      from:
        kind: DockerImage
        name: 'registry.access.redhat.com/3scale-amp22/system:1.7'
    - name: latest
      from:
        kind: ImageStreamTag
        name: 2.2.0
"
```

6. To patch the **system-sphinx** deployment configuration, run the following **oc patch** command:

```
oc patch dc/system-sphinx -p "
metadata:
  labels:
    app: System
spec:
  template:
    spec:
      containers:
      - imagePullPolicy: IfNotPresent
        image: registry.access.redhat.com/3scale-amp22/system:1.7
        name: system-sphinx
        resources:
          limits:
            cpu: 1000m
            memory: 512Mi
          requests:
            cpu: 80m
            memory: 250Mi
"
```

7. To patch the **system-redis** deployment configuration, run the following **oc patch** command:

```
oc patch dc/system-redis -p '
metadata:
  labels:
    app: System
spec:
  template:
    spec:
      containers:
      - imagePullPolicy: IfNotPresent
        name: system-redis
        command:
          - "/opt/rh/rh-redis32/root/usr/bin/redis-server"
        args:
          -  "/etc/redis.d/redis.conf"
          -  "--daemonize"
          -  "no"
        resources:
          limits:
            memory: 32Gi
            cpu: 500m
```

```
          requests:
            cpu: 150m
            memory: 256Mi
        volumeMounts:
        - name: system-redis-storage
          mountPath: "/var/lib/redis/data"
        - name: redis-config
          mountPath: /etc/redis.d/
'
```

8. To patch the **system-memcache** deployment configuration, run the following **oc patch** command:

```
oc patch dc/system-memcache -p "
metadata:
  labels:
    app: System
spec:
  template:
    spec:
      containers:
      - imagePullPolicy: IfNotPresent
        name: memcache
        resources:
          limits:
            cpu: 250m
            memory: 96Mi
          requests:
            cpu: 50m
            memory: 64Mi
"
```

1.8. PATCH BACKEND COMPONENTS

1. To patch the **backend-cron** deployment configuration, run the following **oc patch** command:

```
oc patch dc/backend-cron -p "
metadata:
  labels:
    app: Backend
spec:
  template:
    spec:
      containers:
      - name: backend-cron
        env:
        - name: CONFIG_REDIS_PROXY
          value: redis://backend-redis:6379/0
        - name: CONFIG_REDIS_SENTINEL_HOSTS
          value: ""
        - name: CONFIG_REDIS_SENTINEL_ROLE
          value: ""
        - name: CONFIG_QUEUES_MASTER_NAME
          value: redis://backend-redis:6379/1
        - name: CONFIG_QUEUES_SENTINEL_HOSTS
```

```
          value: ""
        - name: CONFIG_QUEUES_SENTINEL_ROLE
          value: ""
        - name: RACK_ENV
          value: "production"
        image: registry.access.redhat.com/3scale-amp22/backend:1.6
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 150m
            memory: 80Mi
          requests:
            cpu: 50m
            memory: 40Mi
"
```

2. To patch the **backend-worker** deployment configuration, run the following **oc patch** command:

```
oc patch dc/backend-worker -p "
metadata:
  labels:
    app: Backend
spec:
  template:
    spec:
      containers:
      - name: backend-worker
        env:
          - name: CONFIG_REDIS_PROXY
            value: redis://backend-redis:6379/0
          - name: CONFIG_REDIS_SENTINEL_HOSTS
          - name: CONFIG_REDIS_SENTINEL_ROLE
          - name: CONFIG_QUEUES_MASTER_NAME
            value: redis://backend-redis:6379/1
          - name: CONFIG_QUEUES_SENTINEL_HOSTS
          - name: CONFIG_QUEUES_SENTINEL_ROLE
          - name: RACK_ENV
            value: \"production\"
          - name: PUMA_WORKERS
            value: \"16\"
          - name: CONFIG_EVENTS_HOOK
            value: http://system-master:3000/master/events/import
          - name: CONFIG_EVENTS_HOOK_SHARED_SECRET
            value: ${EVENTS_SHARED_SECRET}
        image: registry.access.redhat.com/3scale-amp22/backend:1.6
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 1000m
            memory: 300Mi
          requests:
            cpu: 150m
            memory: 50Mi
"
```

3. To patch the **backend-listener** deployment configuration, run the following **oc patch** command:

```
oc patch dc/backend-listener -p "
metadata:
  labels:
    app: Backend
spec:
  template:
    spec:
      containers:
      - name: backend-listener
        env:
          - name: CONFIG_REDIS_PROXY
            value: redis://backend-redis:6379/0
          - name: CONFIG_REDIS_SENTINEL_HOSTS
          - name: CONFIG_REDIS_SENTINEL_ROLE
            value: ""
          - name: CONFIG_QUEUES_MASTER_NAME
            value: redis://backend-redis:6379/1
          - name: CONFIG_QUEUES_SENTINEL_HOSTS
          - name: CONFIG_QUEUES_SENTINEL_ROLE
            value: ""
          - name: RACK_ENV
            value: \"production\"
          - name: CONFIG_INTERNAL_API_USER
            value: \"${CONFIG_INTERNAL_API_USER}\"
          - name: CONFIG_INTERNAL_API_PASSWORD
            value: \"${CONFIG_INTERNAL_API_PASSWORD}\"
          - name: PUMA_WORKERS
            value: \"16\"
        image: registry.access.redhat.com/3scale-amp22/backend:1.6
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 1000m
            memory: 700Mi
          requests:
            cpu: 500m
            memory: 550Mi
"
```

4. To patch the **amp-backend** image stream, run the following **oc patch** command:

```
oc patch is/amp-backend -p "
spec:
  tags:
    - name: 2.2.0
      annotations:
        openshift.io/display-name: AMP backend
      from:
        kind: DockerImage
        name: 'registry.access.redhat.com/3scale-amp22/backend:1.6'
    - name: latest
      from:
```

```
        kind: ImageStreamTag
        name: 2.2.0
"
```

5. To patch the **backend-redis** deployment configuration, run the following **oc patch** command:

```
oc patch dc/backend-redis -p '
metadata:
  labels:
    app: Backend
spec:
  template:
    spec:
      containers:
      - name: backend-redis
        command:
          - "/opt/rh/rh-redis32/root/usr/bin/redis-server"
        args:
          -  "/etc/redis.d/redis.conf"
          -  "--daemonize"
          -  "no"
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 2000m
            memory: 32Gi
          requests:
            cpu: 1000m
            memory: 1024Mi
        volumeMounts:
        - name: backend-redis-storage
          mountPath: "/var/lib/redis/data"
        - name: redis-config
          mountPath: /etc/redis.d/
'
```

1.9. PATCH APICAST

1. To patch the **apicast-staging** deployment configuration, run the following **oc patch** command:

```
oc patch dc/apicast-staging -p "
metadata:
  labels:
    app: APIcast
spec:
  template:
    spec:
      containers:
      - name: apicast-staging
        env:
        - name: THREESCALE_PORTAL_ENDPOINT
          value: \"http://${APICAST_ACCESS_TOKEN}@system-
master:3000/master/api/proxy/configs\"
        - name: APICAST_CONFIGURATION_LOADER
          value: \"lazy\"
```

```
        - name: APICAST_CONFIGURATION_CACHE
          value: \"0\"
        - name: THREESCALE_DEPLOYMENT_ENV
          value: \"sandbox\"
        - name: APICAST_MANAGEMENT_API
          value: \"${APICAST_MANAGEMENT_API}\"
        - name: BACKEND_ENDPOINT_OVERRIDE
          value: http://backend-listener:3000
        - name: OPENSSL_VERIFY
          value: \"${OPENSSL_VERIFY}\"
        - name: APICAST_RESPONSE_CODES
          value: \"${APICAST_RESPONSE_CODES}\"
        - name: REDIS_URL
          value: \"redis://system-redis:6379/2\"
        image: registry.access.redhat.com/3scale-amp22/apicast-gateway:1.8
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 100m
            memory: 128Mi
          requests:
            cpu: 50m
            memory: 64Mi
"
```

2. To patch the **apicast-production** deployment configuration, run the following **oc patch** command:

```
oc patch dc/apicast-production -p "
metadata:
  labels:
    app: APIcast
spec:
  template:
    spec:
      containers:
      - name: apicast-production
        env:
        - name: THREESCALE_PORTAL_ENDPOINT
          value: \"http://${APICAST_ACCESS_TOKEN}@system-
master:3000/master/api/proxy/configs\"
        - name: APICAST_CONFIGURATION_LOADER
          value: \"boot\"
        - name: APICAST_CONFIGURATION_CACHE
          value: \"300\"
        - name: THREESCALE_DEPLOYMENT_ENV
          value: \"production\"
        - name: APICAST_MANAGEMENT_API
          value: \"${APICAST_MANAGEMENT_API}\"
        - name: BACKEND_ENDPOINT_OVERRIDE
          value: http://backend-listener:3000
        - name: OPENSSL_VERIFY
          value: \"${APICAST_OPENSSL_VERIFY}\"
        - name: APICAST_RESPONSE_CODES
          value: \"${APICAST_RESPONSE_CODES}\"
        - name: REDIS_URL
```

```
          value: \"redis://system-redis:6379/1\"
        image: registry.access.redhat.com/3scale-amp22/apicast-gateway:1.8
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 1000m
            memory: 128Mi
          requests:
            cpu: 500m
            memory: 64Mi
"
```

3. To patch the **amp-apicast** image stream, run the following **oc patch** command:

```
oc patch is/amp-apicast -p "
spec:
  tags:
    - name: 2.2.0
      annotations:
        openshift.io/display-name: AMP apicast
      from:
        kind: DockerImage
        name: 'registry.access.redhat.com/3scale-amp22/apicast-gateway:1.8'
    - name: latest
      from:
        kind: ImageStreamTag
        name: 2.2.0
"
```

4. To patch the **apicast-wildcard-router** deployment configuration, run the following **oc patch** command:

```
oc patch dc/apicast-wildcard-router -p "
metadata:
  labels:
    app: APIcast
spec:
  template:
    spec:
      containers:
      - name: apicast-wildcard-router
        env:
        - name: API_HOST
          value: \"http://${APICAST_ACCESS_TOKEN}@system-master:3000\"
        image: registry.access.redhat.com/3scale-amp22/wildcard-router:1.6
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            cpu: 500m
            memory: 64Mi
          requests:
            cpu: 120m
            memory: 32Mi
"
```

5. To patch the **amp-wildcard-router** image stream, run the following **oc patch** command:

```
oc patch is/amp-wildcard-router -p "
spec:
  tags:
    - name: 2.2.0
      annotations:
        openshift.io/display-name: AMP wildcard router
      from:
        kind: DockerImage
        name: 'registry.access.redhat.com/3scale-amp22/wildcard-router:1.6'
    - name: latest
      from:
        kind: ImageStreamTag
        name: 2.2.0
"
```

1.10. PATCH ZYNC COMPONENTS

1. To patch the **zync-database** deployment configuration, run the following **oc patch** command:

```
oc patch dc/zync-database -p "
metadata:
  labels:
    app: Zync
spec:
  template:
    spec:
      containers:
      - name: postgresql
        imagePullPolicy: IfNotPresent
        resources:
          limits:
            memory: 2Gi
            cpu: 250m
          requests:
            cpu: 50m
            memory: 250Mi
"
```

2. To patch the **zync** deployment configuration, run the following **oc patch** command:

```
oc patch dc/zync -p "
metadata:
  labels:
    app: Zync
spec:
  template:
    spec:
      containers:
      - name: zync
        image: 'registry.access.redhat.com/3scale-amp22/zync:1.6'
        resources:
          limits:
```

```
          cpu: 1
          memory: 512Mi
        requests:
          cpu: 150m
          memory: 250Mi
"
```

3. To patch the **zync** image stream, run the following **oc patch** command:

```
oc patch is/amp-zync -p "
spec:
  tags:
    - name: 2.2.0
      annotations:
        openshift.io/display-name: AMP zync
      from:
        kind: DockerImage
        name: 'registry.access.redhat.com/3scale-amp22/zync:1.6'
    - name: latest
      from:
        kind: ImageStreamTag
        name: 2.2.0
"
```

1.11. VERIFY UPGRADE

After you have performed the upgrade procedure, verify the success of your upgrade operation by checking the version number in the lower-right corner of your 3scale Admin Portal.

NOTE

It may take some time for your redeployment operations to complete in OpenShift.

CHAPTER 2. API DEPLOYMENT ON MICROSOFT AZURE

Since APIs are platform agnostic, they can be deployed on any platform. This tutorial is fast web API deployment on Microsoft Azure. You will use the Ruby Grape gem to create the API interface, an NGINX proxy, Thin server, and Capistrano to deploy using the command line.

For the purpose of this tutorial, you can use any Ruby-based API running on Thin server, or you can clone the Echo-API.

2.1. CREATE AND CONFIGURE MICROSOFT AZURE VM

Start to generate a *X509 certificate with a 2048-bit RSA keypair* to ssh into your Azure VM. It will be useful when you will set up your VM.

To generate this type of key, you can run the following command:

```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout myPrivateKey.key -out myCert.pem
```

Now, get started by creating your Microsoft Azure account. For this tutorial, you can use the free trial option. Once the Azure account is created, go to the Dashboard on the Virtual Machines tab. There, you will be guided to create your first VM. Choose the *from gallery* option and select an *Ubuntu Server 12.04 LTS*.

On step 2 you will be able to upload the pem you created earlier, you should not be prompted for your password again.

In steps 3 and 4, choose the options that best suit your needs.

It will take a couple of minutes for your VM to be ready. When it is, you will be able to access its dashboard where you can monitor activity (CPU, disk, network) of your VM and upgrade its size.

The VM comes with a few packages installed, so you'll need to access it to install other components. Once the key is created, you can ssh to your VM.

```
ssh -i myPrivateKey.key -p 22 username@servicename.cloudapp.net
```

Once in the VM, run the following commands to install everything you need:

```
sudo apt-get -y update
sudo apt-get -y upgrade
sudo apt-get -y install ruby1.9.3 build-essential libsqlite3-dev libpcre3 libpcre3-dev libssl-dev openssl
libreadline6 libreadline6-dev libxml2-dev libxslt1-dev
```

You can check that Ruby installation is complete by running:

```
ruby -v
```

It should output something like *ruby 1.9.3p194 (2012-04-20 revision 35410) [x86_64-linux]* .

You also need to install **bundler** and **thin**:

```
sudo gem install bundler
sudo gem install thin
```

Now, you should have everything you need on the VM. Go back to its dashboard and click on the *endpoints* tab. There, add the **HTTP** endpoint on port **80**, and the fields should autofill.

2.2. INSTALL OPENRESTY

In order to streamline this step, we recommend that you install the fantastic **OpenResty** web application. It's the standard NGINX core bundled with almost all the necessary third-party NGINX modules built in.

On your Azure VM Compile and install NGINX:

```
cd ~
sudo wget http://agentzh.org/misc/nginx/ngx_openresty-VERSION.tar.gz
sudo tar -zxvf ngx_openresty-VERSION.tar.gz
cd ngx_openresty-VERSION/
sudo ./configure --prefix=/opt/openresty --with-luajit --with-http_iconv_module -j2
sudo make
sudo make install
```

2.3. CONFIGURE YOUR GITHUB REPO

This tutorial uses GitHub to host the code. If you don't already have a repo for your API, make sure to create one and host it on github.com. If you're not familiar with Git and GitHub, check out this great tutorial.

To use Git on your VM and have access to your GitHub repo, you need to generate an SSH key on your VM and add it to Github as explained here.

2.3.1. Warning

Hosting your code on a public GitHub repo makes it vulnerable. Make sure it does not contain any sensitive information such as provider keys before pushing it publicly.

2.4. CONFIGURE YOUR API

This is how the system will work:

1. Thin server will be launched on port 8000.
2. The upstream **YOURAPINAME** is listening on localhost:8000.
3. Upcoming connections on port 80 (as defined in the **server** section) are "redirected" to **YOURAPINAME**.

2.4.1. On 3scale

Rather than reinvent the wheel and implement rate limits, access controls, and analytics from scratch, you'll use 3scale. If you don't have an account yet, sign up here, activate it, and log in to the new instance through the links provided. The first time you log in, choose the option for some sample data to be created, so you'll have some API keys to use later. Go through the tour to get a glimpse of the systems functionality (optional) and then go ahead with implementation.

To get some instant results, start with the API gateway in the staging environment, which can be used while in development. Then configure an NGINX proxy, which can scale up for full production deployments.

There is some documentation on configuring the API proxy here and more advanced configuration options here.

Once you sign in to your 3scale account, launch your API on the main Dashboard screen or Go to API→Select the service (API)→Integration in the sidebar→Proxy https://www.3scale.net/2015/06/how-to-deploy-an-api-amazon-ec2/

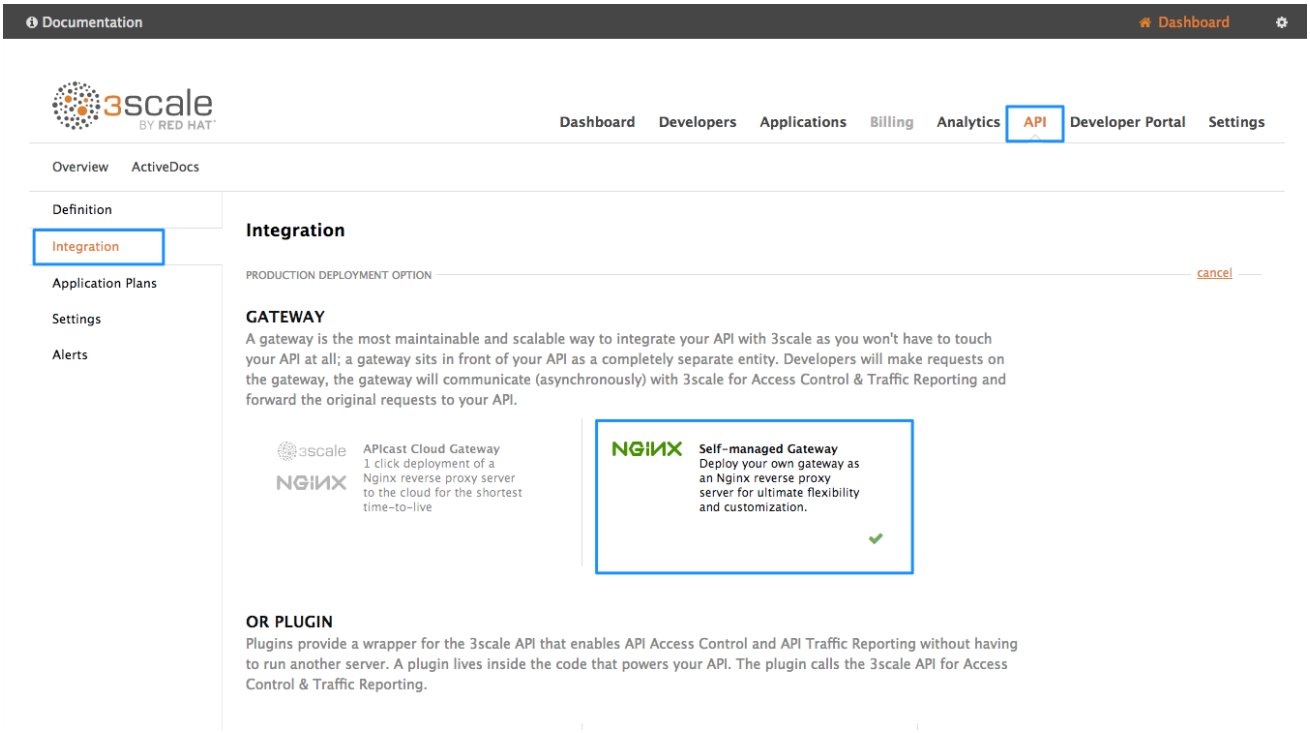

Set the address of your API backend -

```
`http://YOURAPP.cloudapp.net:80`
```

1. After creating some app credentials in 3scale, you can test your API by hitting the staging API gateway endpoint:

```
`https://XXX.staging.apicast.io/v1/words/awesome.json?
app_id=APP_ID&app_key=APP_KEY`
```

where, **XXX** is specific to your staging API gateway and **APP_ID** and **APP_KEY** are the ID and key of one of the sample applications you created when you first logged in to your 3scale account. (If you missed that step, just create a developer account and an application within that account.)

Try it without app credentials, next with incorrect credentials. Then once authenticated, within and over any rate limits that you've defined. Once it's working to your satisfaction, download the config files for NGINX.

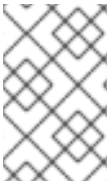

NOTE

Any time you have errors, check whether you can access the API directly: your-public-dns:3000/v1/words/awesome.json. If it's not available, check whether the AWS instance is running and whether the Thin server is running on the instance.*

There, you will be able to change your API backend address to **http://YOURAPP.cloudapp.net:80**.

Once you're done, click on **Download your nginx config**. That will download an archive containing the **.conf** and **.lua** file you're going to use to configure your app.

Modify the **.conf** accordingly:

If the API gateway and the API are on the same VM, delete the block:

```
upstream backend_YOURAPP.cloudapp.net{
server ....
}
```

...and replace it with...

```
upstream YOURAPINAME {
server 127.0.0.1:8000;
}
```

WARNING

YOURAPINAME can only contain URL valid characters as defined in RFC 3986.

In the **.lua** file, modify the line **ngx.var.proxy_pass = "http://backend_YOURAPP.cloudapp.net"**.

With **ngx.var.proxy_pass = "http://YOURAPINAME"** in all cases.

Replace **server_name api.2445580546262.proxy.3scale.net;** with

server_name YOURSERVICENAME.cloudapp.net;

In the **server** block, add this on top:

```
root /home/USERNAME/apps/YOURAPINAME/current;
access_log /home/USERNAME/apps/YOURAPINAME/current/log/thin.log;
error_log /home/USERNAME/apps/YOURAPINAME/current/log/error.log;
```

Replace **access_by_lua_file lua_tmp.lua;**

...with... **access_by_lua_file /opt/openresty/nginx/conf/lua_tmp.lua;**

Before **post_action /out_of_band_authrep_action;** add:

```
proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
proxy_set_header Host $http_host;
```

Finally, rename those files **nginx.conf** and **tmp_lua.lua**.

2.4.2. Capistrano setup

Use **Capistrano** to deploy the API. Capistrano is an automation tool, which will let you set up tasks for your deployments and execute them using a command line interface. Capistrano is used on your local machine to deploy on your remote VM.

To install Capistrano, add this line to your gem file: **gem 'capistrano'**

Run the following command locally to install the new gems and set up Capistrano: **bundle capify.**

Copy **nginx.conf** and **tmp_lua.lua** into /**config**.

2.5. CAPISTRANO SETUP

When you ran the **capify** command, you created two files, **Capfile** and **deploy.rb**. In **deploy.rb**, you describe all the commands necessary to deploy your app.

In /**config** edit **deploy.rb** and replace the content with the following:

```
require "bundler/capistrano"
set :application, "YOURAPINAME"
set :user,"USERNAME"
set :scm, :git
set :repository, "git@github.com:GITHUBUSERNAME/REPO.git"
set :branch, "master"

set :use_sudo, false

server "VNDNSname", :web, :app, :db, primary: true

set :deploy_to, "/home/#{user}/apps/#{application}"
default_run_options[:pty] = true
ssh_options[:forward_agent] = false
ssh_options[:port] = 22
ssh_options[:keys] = ["/PATH/TO/myPrivateKey.key"]

namespace :deploy do
   task :start, :roles => [:web, :app] do
     run "cd #{deploy_to}/current && nohup bundle exec thin start -C config/production_config.yml -R
config.ru"
     sudo "/opt/openresty/nginx/sbin/nginx -p /opt/openresty/nginx/ -c
/opt/openresty/nginx/conf/nginx.conf"
   end

   task :stop, :roles => [:web, :app] do
     run "kill -QUIT cat /opt/openresty/nginx/logs/nginx.pid"
     run "cd #{deploy_to}/current && nohup bundle exec thin stop -C config/production_config.yml -R
config.ru"
   end

   task :restart, :roles => [:web, :app] do
     deploy.stop
     deploy.start
   end

   task :setup_config, roles: :app do
     sudo "ln -nfs #{current_path}/config/nginx.conf /opt/openresty/nginx/conf/nginx.conf"
     sudo "ln -nfs #{current_path}/config/lua_tmp.lua /opt/openresty/nginx/conf/lua_tmp.lua"
```

```
      sudo "mkdir -p #{shared_path}/config"
    end
    after "deploy:setup", "deploy:setup_config"
end
```

This will ensure that Capistrano doesn't try to run **rake:migrate**. (This is not a Rails project!)

```
task :cold do
  deploy.update
  deploy.start
end
```

In above text, replace the following:

- **VNDNSname** with your .cloudapp.net DNS.
- **YOURAPINAME** with your applicationame.
- **USERNAME** with the username used to login into the VM.
- **GITHUBUSERNAME** with your Github username.
- **REPO** with your Github repo name.
- **/PATH/TO** with the path to access the SSH key created before.

The above works well if you don't have a database in your API. If you do have a database, comment the lines:

```
task :cold do
  deploy.update
  deploy.start
end
```

You also need to add a file **production_config.yml** in **/config** to configure the Thin server.

```
environment: production
chdir: /home/USERNAME/apps/YOURAPINAME/current/
address: 127.0.0.1
user: USERNAME
port: 8000
pid: /home/USERNAME/apps/YOURAPINAME/current/tmp/thin.pid
rackup: /home/USERNAME/apps/YOURAPINAME/current/config.ru
log: /home/USERNAME/apps/YOURAPINAME/current/log/thin.log
max_conns: 1024
timeout: 30
max_persistent_conns: 512
daemonize: true
```

Again, change usernames and paths accordingly.

Commit the changes on the project and upload them to GitHub.

```
git add .
git commit -m "adding config files"
git push
```

You are almost done.

2.6. DEPLOY

From your local development machine, run the following command to set up the remote Azure VM:

```
cap deploy:setup
```

You should not be prompted for a password if the path to your ssh key is correct.

Capistrano will connect to your VM and create an **apps** directory under the **home** directory of the user account.

Now, you can deploy your API to the VM and launch Thin server using the command: **cap deploy:cold**

This command should get the latest commit on your GitHub. Launch OpenResty and Thin server.

Your API should now be available on the URL:

MYAPI.cloudapp.net/path/to/resources

2.6.1. Troubleshooting

If you are not able to access to your API, ssh to your VM and check that you can call it on **localhost** using **curl**. Like this:

```
curl -X GET http://localhost:8000/v2/words/hello.json?app_id=APPID&app_key=APPKEY`
```

If it works, there is something wrong in nginx configuration.

You can check nginx logs on your VM with

```
cat /opt/openresty/nginx/logs/error.log
```

You should now have an API running on an Azure Linux instance.

Hope you enjoyed this tutorial. Please let us know if you have any questions or comments. We look forward to hearing from you.

CHAPTER 3. DEPLOY AN API ON AMAZON EC2 FOR AWS ROOKIES

At 3scale we find Amazon to be a fantastic platform for running APIs due to the complete control you have on the application stack. However, for people new to AWS, the learning curve is quite steep. So we put together our best practices into this short tutorial. Besides Amazon EC2, we'll use the Ruby Grape gem to create the API interface and an NGINX gateway to handle access control. Best of all everything in this tutorial is completely free.

3.1. PREREQUISITES

For the purpose of this tutorial you'll need a running API based on Ruby and Thin server. If you don't have one you can simply clone an example repo as described below in the "Deploying the Application" section.

We'll begin with the creation and configuration of the Amazon EC2 instance. If you already have an EC2 instance (micro or not), you can jump to the next step, "Preparing Instance for Deployment".

3.2. CREATE AND CONFIGURE EC2 INSTANCE

Start by signing up for the Amazon Elastic Compute Cloud (Amazon EC2). The free tier is enough to cover all your basic needs. Once the account is created, go to the EC2 dashboard under your AWS Management Console and click on the "launch instance" button. That will transfer you to a pop-up window where you'll continue the process:

- Choose the classic wizard
- Choose an AMI (Ubuntu Server 12.04.1 LTS 32bit, T1micro instance) leaving all the other settings for "instance details" as default
- Create a key pair and download it. This will be the key that you'll use to make an ssh connection to the server. It's VERY IMPORTANT!
- Add inbound rules for the firewall with source always 0.0.0.0/0 (HTTP, HTTPS, ALL ICMP, TCP port 3000 used by the Ruby Thin server)

3.3. PREPARE INSTANCE FOR DEPLOYMENT

Once the instance is created and running, you can connect there directly from the console (Windows users from PuTTY). Right click on your instance, connect, and choose **Connect with a standalone SSH Client**.

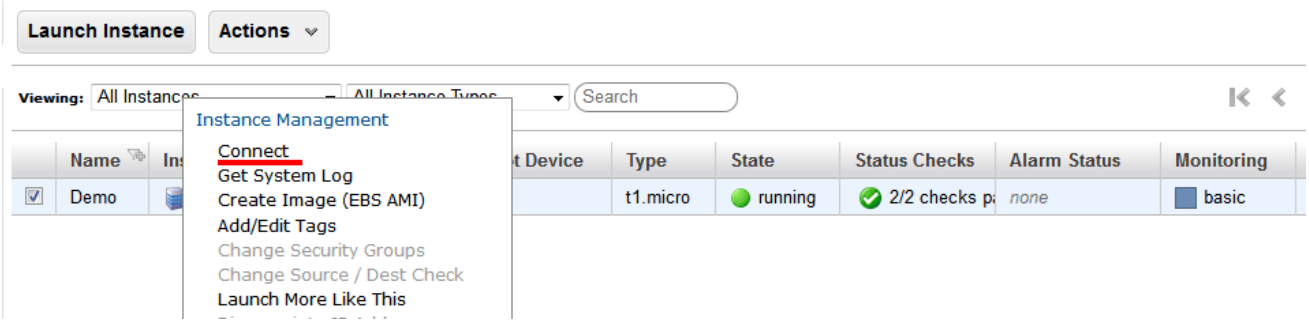

Follow the steps and change the username to "ubuntu" (instead of "root") in the given example.

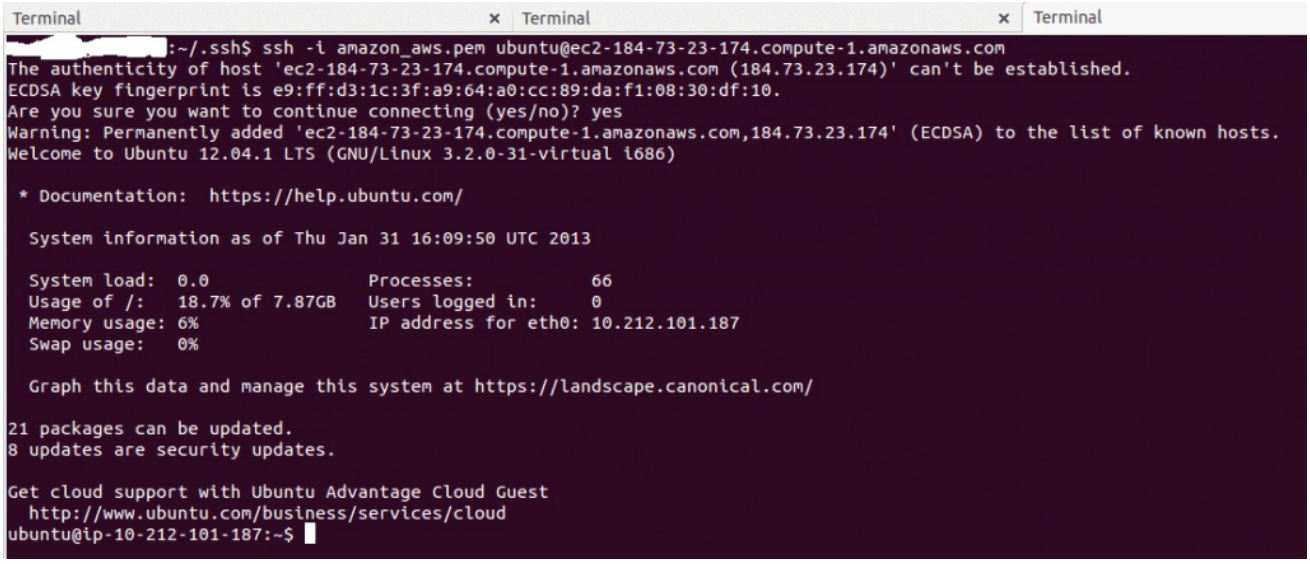

After executing this step you are connected to your instance. You'll have to install new packages. Some of them require root credentials, so you'll have to set a new root password: **sudo passwd root**. Then login as root: **su root**.

Now with root credentials, execute: **sudo apt-get update**

Switch back to your normal user with **exit** command and install all required packages:

- Install the libraries that will be required by rvm, Ruby, and Git:

```
sudo apt-get install build-essential git zlib1g-dev libssl-dev libreadline-gplv2-dev imagemagick
libxml2-dev libxslt1-dev openssl zlib1g libyaml-dev libxslt-dev autoconf libc6-dev ncurses-dev
automake libtool bison libpq-dev libpq5 libeditline-dev

  sudo apt-get install libreadline6 libreadline6-dev
```

- Install Git (on Linux rather than from Source)
- Install rvm
- Install Ruby

```
rvm install 1.9.3
rvm use 1.9.3 --default
```

3.4. DEPLOYING THE APPLICATION

Our example, the Sentiment API, is located on GitHub. Try cloning the repository:

```
git clone git@github.com:jerzyn/api-demo.git
```

You can review the code and tutorial on creating and deploying this app here and here. Note the changes – we're using only v1, as authentication will go through the gateway.

Now you can deploy the app by issuing **bundle install**.

Now you can start the thin server: **thin start**.

To access the API directly (without any security or access control) access: **your-public-ip:3000/v1/words/awesome.json** You can find your public IP in the AWS EC2 **Dashboard > Instances** in the details window of your instance.

3.4.1. Optional

If you want to assign a custom domain to your Amazon instance, you'll have to do one thing: Add an A record to the DNS record of your domain, mapping the domain to the public IP address.

Your domain provider should either give you some way to set the A record (the IPv4 address), or it will give you a way to edit the nameservers of your domain. If they don't allow you to set the A record directly find a DNS management service, register your domain as a zone there, and the service will give you the nameservers to enter in the admin panel of your domain provider. You can then add the A record for the domain. Some possible DNS management services include ZoneEdit (basic, free) or Amazon route 53.

At this point, your API is open to the world. This is good and bad—it's great that you're sharing, but bad that without rate limits a few apps could kill the resources of your server and you would have no insight into who is using your API and how it's being used. The solution is to add API management.

3.5. ENABLING API MANAGEMENT WITH 3SCALE

Rather than reinventing the wheel and implement rate limits, access controls, and analytics from scratch, you can leverage the 3scale API Management Platform. Sign up for a 3scale account if you haven't already, activate it, and log in through the links provided. The first time you log in, some sample data will be created for you so you'll have an API key to use later. You can go through the wizard to get an idea of the system's functionality (optional). Then start with the implementation.

To get some instant results, we'll start with the API gateway in the staging environment which can be used while in development. Then we'll configure an NGINX gateway that can scale up for full production deployments. Here's some documentation on the configuration of the API gateway, as well as more advanced configuration options.

Once you've signed in to your 3scale account, go to **Dashboard > API > Select the service (API) > Integration > edit integration settings** and then choose **APIcast Self-managed**.

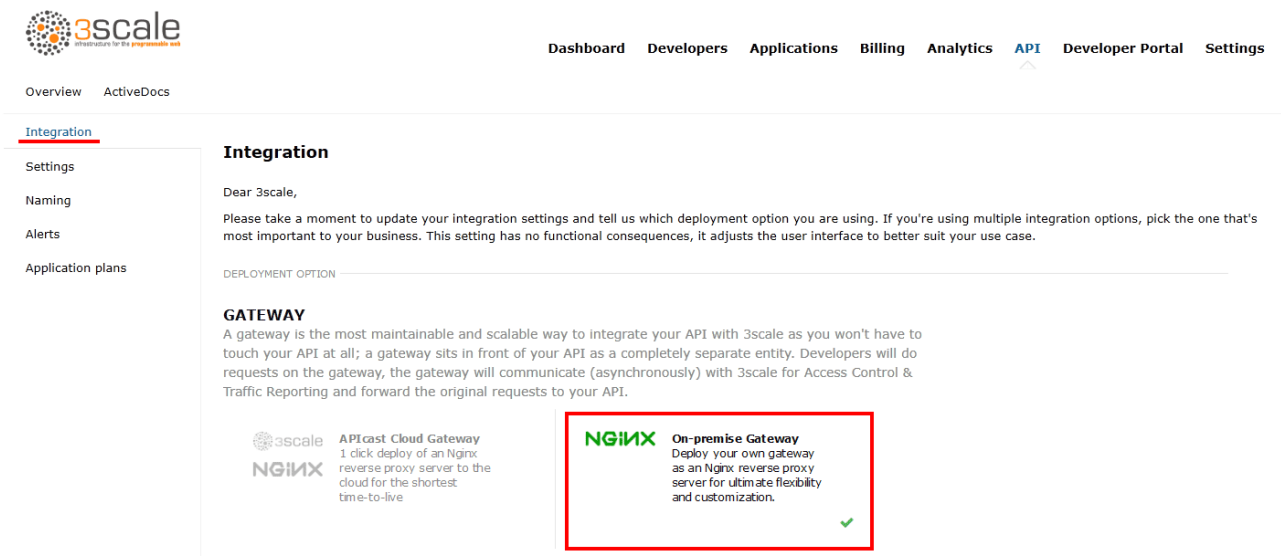
3scale
Dashboard Developers Applications Billing Analytics API Developer Portal Settings
Overview ActiveDocs
Integration
Settings
Naming
Alerts
Application plans
Integration
Dear 3scale,
Please take a moment to update your integration settings and tell us which deployment option you are using. If you're using multiple integration options, pick the one that's most important to your business. This setting has no functional consequences, it adjusts the user interface to better suit your use case.
DEPLOYMENT OPTION
GATEWAY
A gateway is the most maintainable and scalable way to integrate your API with 3scale as you won't have to touch your API at all; a gateway sits in front of your API as a completely separate entity. Developers will do requests on the gateway, the gateway will communicate (asynchronously) with 3scale for Access Control & Traffic Reporting and forward the original requests to your API.
3scale NGINX APIcast Cloud Gateway
1 click deploy of an Nginx reverse proxy server to the cloud for the shortest time-to-live
NGINX On-premise Gateway
Deploy your own gateway as an Nginx reverse proxy server for ultimate flexibility and customization.

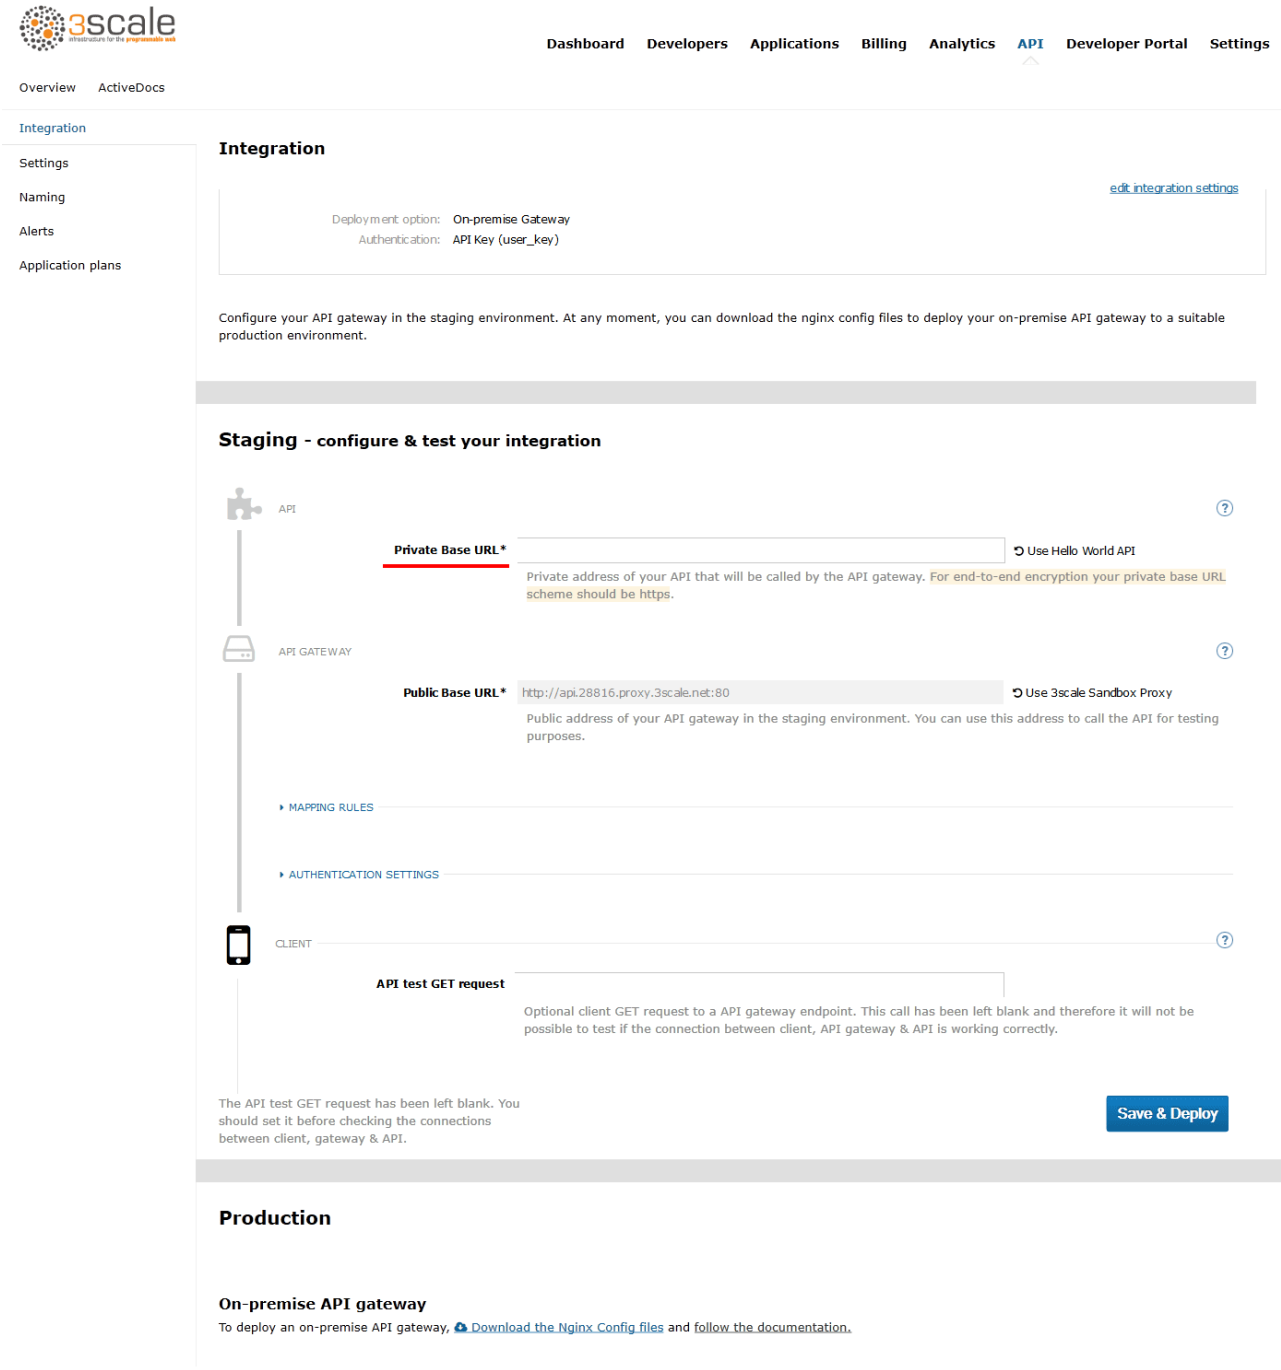
3scale
Dashboard Developers Applications Billing Analytics API Developer Portal Settings
Overview ActiveDocs
Integration
Settings
Naming
Alerts
Application plans
Integration
edit integration settings
Deployment option: On-premise Gateway
Authentication: API Key (user_key)
Configure your API gateway in the staging environment. At any moment, you can download the nginx config files to deploy your on-premise API gateway to a suitable production environment.
Staging - configure & test your integration
API
Private Base URL*
Use Hello World API
Private address of your API that will be called by the API gateway. For end-to-end encryption your private base URL scheme should be https.
API GATEWAY
Public Base URL* http://api.28816.proxy.3scale.net:80
Use 3scale Sandbox Proxy
Public address of your API gateway in the staging environment. You can use this address to call the API for testing purposes.
MAPPING RULES
AUTHENTICATION SETTINGS
CLIENT
API test GET request
Optional client GET request to a API gateway endpoint. This call has been left blank and therefore it will not be possible to test if the connection between client, API gateway & API is working correctly.
The API test GET request has been left blank. You should set it before checking the connections between client, gateway & API.
Save & Deploy
Production
On-premise API gateway
To deploy an on-premise API gateway, Download the Nginx Config files and follow the documentation.

Set the address of of your API backend. This has to be the public IP address unless the custom domain has been set, including http protocol and port 3000. Now you can save the changes to the API gateway in the staging environment to test your API by hitting the staging endpoint.

```
http://api.XXX.proxy.3scale.net/v1/words/awesome.json?user_key=USER_KEY
```

Where **XXX** is specific to your 3scale account and **USER_KEY** is the authentication key of one of the sample applications created when you first logged into your 3scale account. (If you missed that step just create a developer account and an application within that account.)

Try it without app credentials; next with incorrect credentials; and then once authenticated, within and over any rate limits you have defined. Once it's working to your satisfaction you can download the config files for NGINX.

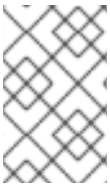

NOTE

Whenever you have errors, check whether you can access the API directly: your-public-dns:3000/v1/words/awesome.json. If that is not available, you need to check whether the AWS instance is running and whether the Thin server is running on the instance.

3.6. INSTALL AND DEPLOY APICAST (YOUR API GATEWAY)

Finally, to deploy install and deploy APIcast, follow the steps in the APIcast 2.0 self-managed tutorial for 'local' deploy.

You're almost finished! The last step is to start the NGINX gateway and put some traffic through it. If it's not running yet (remember the Thin server has to be started first), go to your EC2 instance terminal (the one you were connecting through ssh before) and start it now.

The last step will be verifying that the traffic goes through with a proper authorization. To do that, access:

```
http://your-public-ip/v1/words/awesome.json?app_id=APP_ID&app_key=APP_KEY
```

where APP_ID and APP_KEY are key and ID of the application you want to access through the API call.

Once everything is confirmed as working correctly, you'll want to block public access to the API backend on port 3000, which bypasses any access controls.

CHAPTER 4. BUILDING A 3SCALE API MANAGEMENT SYSTEM IMAGE WITH THE ORACLE DATABASE RELATIONAL DATABASE MANAGEMENT SYSTEM

By default, 3scale has a component called **system** which stores configuration data in a MySQL database. You have the option to override the default database and store your information in an external Oracle Database. Follow the steps in this document to build a custom system container image with your own Oracle Database client binaries and deploy 3scale to OpenShift.

4.1. BEFORE YOU BEGIN

4.1.1. Obtain Oracle software components

Before you can build the custom 3scale system container image, you must acquire a supported version of the following Oracle software components:

- Oracle Instant Client Package Basic or Basic Light
- Oracle Instant Client Package SDK
- Oracle Instant Client Package ODBC

4.1.2. Meet prerequisites

You must also meet the following prerequisites:

- A supported version of Oracle Database accessible from your OpenShift cluster
- Access to the Oracle Database **system** user for installation procedures
- Possess the Red Hat 3scale 2.2 amp.yml template

4.2. PREPARING ORACLE DATABASE

1. Create a new database
   The following settings are required for the Oracle Database to work with 3scale:

   ```
   ALTER SYSTEM SET max_string_size=extended SCOPE=SPFILE;
   ```

   ```
   ALTER SYSTEM SET compatible='12.2.0.1' SCOPE=SPFILE;
   ```

2. Collect the database details.
   Get the following information that will be needed for 3scale configuration:

   - Oracle Database URL
   - Oracle Database service name
   - Oracle Database **system** user name and password
   - Oracle Database service name

For information on creating a new database in Oracle Database, refer to the Oracle documentation.

4.3. BUILDING THE SYSTEM IMAGE

1. clone the 3scale-amp-openshift-templates github repository
2. place your Oracle Database Instant Client Package files into the **3scale-amp-openshift-templates/amp/system-oracle/oracle-client-files** directory
3. run the **oc new-app** command with the -f option and specify the **build.yml** OpenShift template

```
$ oc new-app -f build.yml
```

4. run the **oc new-app** command with the -f option, specifying the **amp.yml** OpenShift template, and the -p option, specifying the **WILDCARD_DOMAIN** parameter with the domain of your OpenShift cluster

```
$ oc new-app -f amp.yml -p WILDCARD_DOMAIN=example.com
```

5. enter the following shell **for** loop command, specifying the following information you collected in the Preparing Oracle Database section previously:
   - **{USER}**: the username that will represent 3scale in your Oracle Database
   - **{PASSWORD}**: the password for **USER**
   - **{ORACLE_DB_URL}**: the URL of your Oracle Database
   - **{DATABASE}**: the service name of the database you created in Oracle Database
   - **{PORT}**: the port number of your Oracle Database

```
for dc in system-app system-resque system-sidekiq system-sphinx; do oc env dc/$dc --overwrite DATABASE_URL="oracle-enhanced://{USER}:{PASSWORD}@{ORACLE_DB_URL}:{PORT}/{DATABASE}"; done
```

6. enter the following **oc patch** command, specifying the same **USER**, **PASSWORD**, **ORACLE_DB_URL**, **PORT**, and **DATABASE** values that you provided in the previous step above:

```
$ oc patch dc/system-app -p '[{"op": "replace", "path": "/spec/strategy/rollingParams/pre/execNewPod/env/1/value", "value": "oracle-enhanced://{USER}:{PASSWORD}@{ORACLE_DB_URL}:{PORT}/{DATABASE}"}]' --type=json
```

7. enter the following **oc patch** command, specifying your own Oracle Database **system** user password in the **SYSTEM_PASSWORD** field:

```
$ oc patch dc/system-app -p '[{"op": "add", "path": "/spec/strategy/rollingParams/pre/execNewPod/env/-", "value": {"name": "ORACLE_SYSTEM_PASSWORD", "value": "SYSTEM_PASSWORD"}}]' --type=json
```

8. enter the **oc start-build** command to build the new system image:

```
oc start-build 3scale-amp-system-oracle --from-dir=.
```

CHAPTER 5. 3SCALE AMP ON-PREMISES INSTALLATION GUIDE

In this guide you'll learn how to install 3scale 2.2 (on-premises) on OpenShift using OpenShift templates.

5.1. PREREQUISITES

- You must configure 3scale servers for UTC (Coordinated Universal Time).

5.2. 3SCALE AMP OPENSHIFT TEMPLATES

Red Hat 3scale API Management Platform (AMP) 2.2 provides an OpenShift template. You can use this template to deploy AMP onto OpenShift Container Platform.

The 3scale AMP template is composed of the following:

- Two built-in APIcast API gateways
- One AMP admin portal and developer portal with persistent storage

5.3. SYSTEM REQUIREMENTS

The 3scale Api Management OpenShift template requires the following:

5.3.1. Environment Requirements

3scale API Management requires an environment specified in supported configurations.

Persistent Volumes:

- 3 RWO (ReadWriteOnce) persistent volumes for Redis and MySQL persistence
- 1 RWX (ReadWriteMany) persistent volume for CMS and System-app Assets

The RWX persistent volume must be configured to be group writable. Refer to the OpenShift documentation for a list of persistent volume types which support the required access modes.

5.3.2. Hardware Requirements

Hardware requirements depend on your usage needs. Red Hat recommends you test and configure your environment to meet your specific requirements. Following are the recommendations when configuring your environment for 3scale on OpenShift:

- Compute optimized nodes for deployments on cloud environments (AWS c4.2xlarge or Azure Standard_F8).
- Very large installations may require a separate node (AWS M4 series or Azure Av2 series) for Redis if memory requirements exceed your current node's available RAM.
- Separate nodes between routing and compute tasks.
- Dedicate compute nodes to 3scale specific tasks.

- Set the **PUMA_WORKERS** variable of the backend listener to the number of cores in your compute node.

5.4. CONFIGURE NODES AND ENTITLEMENTS

Before you can deploy 3scale on OpenShift, you must configure your nodes and the entitlements required for your environment to fetch images from Red Hat.

Perform the following steps to configure entitlements:

1. Install Red Hat Enterprise Linux (RHEL) on each of your nodes.
2. Register your nodes with Red Hat using the Red Hat Subscription Manager (RHSM).
3. Attach your nodes to your 3scale subscription using RHSM.
4. Install OpenShift on your nodes, complying with the following requirements:
   - You must use a supported OpenShift version.
   - You must configure persistent storage on a file system that supports multiple writes.
5. Install the OpenShift command line interface.
6. Enable access to the **rhel-7-server-3scale-amp-2.2-rpms** repository using the subscription manager:

```
sudo subscription-manager repos --enable=rhel-7-server-3scale-amp-2.2-rpms
```

7. Install the **3scale-amp-template** AMP template. The template will be saved in **/opt/amp/templates**.

```
sudo yum install 3scale-amp-template
```

5.5. DEPLOY THE 3SCALE AMP ON OPENSHIFT USING A TEMPLATE

5.5.1. Prerequisites:

- An OpenShift cluster configured as specified in the Chapter 3, *Configure Nodes and Entitlements* section.
- A domain, preferably wildcard, that resolves to your OpenShift cluster.
- Access to the Red Hat container catalog.
- (Optional) A working SMTP server for email functionality.

Follow these procedures to install AMP onto OpenShift using a .yml template:

- Import the AMP Template
- Configure SMTP Variables (Optional)

5.5.2. Import the AMP Template

Perfrom the following steps to import the AMP template into your OpenShift cluster:

1. From a terminal session log in to OpenShift:

```
oc login
```

2. Select your project, or create a new project:

```
oc project <project_name>
```

```
oc new-project <project_name>
```

3. Enter the **oc new-app** command:

   - Specify the **--file** option with the path to the amp.yml file you downloaded as part of the configure nodes and entitlements section .
   - Specify the **--param** option with the **WILDCARD_DOMAIN** parameter set to the domain of your OpenShift cluster:
   - Optionally, specify the **--param** option with the **WILDCARD_POLICY** parameter set to **subdomain** to enable wildcard domain routing:
     Without Wildcard Routing:

```
oc new-app --file /opt/amp/templates/amp.yml --param WILDCARD_DOMAIN=
<WILDCARD_DOMAIN>
```

     With Wildcard Routing:

```
oc new-app --file /opt/amp/templates/amp.yml --param WILDCARD_DOMAIN=
<WILDCARD_DOMAIN> --param WILDCARD_POLICY=Subdomain
```

4. The terminal will show the master and tenant URLs and credentials for your newly created AMP admin portal. This output should include the following information:

   - master admin username
   - master password
   - master token information
   - tenant username
   - tenant password
   - tenant token information

```
Log in to https://user-admin.3scale-project.example.com as admin/xXxXyz123.

...

* With parameters:
 * ADMIN_PASSWORD=xXxXyz123 # generated
 * ADMIN_USERNAME=admin
 * TENANT_NAME=user
```

```
...

* MASTER_NAME=master
* MASTER_USER=master
* MASTER_PASSWORD=xXxXyz123 # generated

...

--> Success
Access your application via route 'user-admin.3scale-project.example.com'
Access your application via route 'master-admin.3scale-project.example.com'
Access your application via route 'backend-user.3scale-project.example.com'
Access your application via route 'user.3scale-project.example.com'
Access your application via route 'api-user-apicast-staging.3scale-project.example.com'
Access your application via route 'api-user-apicast-production.3scale-
project.example.com'
Access your application via route 'apicast-wildcard.3scale-project.example.com'

...
```

Make a note of these details for future reference.

NOTE

You may need to wait a few minutes for AMP to fully deploy on OpenShift for your login and credentials to work.

More Information

For information about wildcard domains on OpenShift, visit Using Wildcard Routes (for a Subdomain) .

5.5.3. Configure SMTP Variables (Optional)

OpenShift uses email to send notifications and invite new users. If you intend to use these features, you must provide your own SMTP server and configure SMTP variables in the SMTP config map.

Perform the following steps to configure the SMTP variables in the SMTP config map:

1. If you are not already logged in, log in to OpenShift:

```
oc login
```

2. Configure variables for the SMTP config map. Use the **oc patch** command, specify the **configmap** and **smtp** objects, followed by the **-p** option and write the following new values in JSON for the following variables:

Variable	Description
address	Allows you to specify a remote mail server as a relay
username	Specify your mail server username

password	Specify your mail server password
domain	Specify a HELO domain
port	Specify the port on which the mail server is listening for new connections
authentication	Specify the authentication type of your mail server. Allowed values: **plain** (sends the password in the clear), **login** (send password Base64 encoded), or **cram_md5** (exchange information and a cryptographic Message Digest 5 algorithm to hash important information)
openssl.verify.mode	Specify how OpenSSL checks certificates when using TLS. Allowed values: **none**, **peer**, **client_once**, or **fail_if_no_peer_cert**.

Example:

```
oc patch configmap smtp -p '{"data":{"address":"<your_address>"}}'
oc patch configmap smtp -p '{"data":{"username":"<your_username>"}}'
oc patch configmap smtp -p '{"data":{"password":"<your_password>"}}'
```

3. After you have set the configmap variables, redeploy the **system-app**, **system-resque**, and **system-sidekiq** pods:

```
oc rollout latest dc/system-app
oc rollout latest dc/system-resque
oc rollout latest dc/system-sidekiq
```

5.6. 3SCALE AMP TEMPLATE PARAMETERS

Template parameters configure environment variables of the AMP yml template during and after deployment.

Name	Description	Default Value	Required?
AMP_RELEASE	AMP release tag.	2.2.0	yes
ADMIN_PASSWORD	A randomly generated AMP administrator account password.	N/A	yes
ADMIN_USERNAME	AMP administrator account username.	admin	yes

APICAST_ACCESS_TOKEN	Read Only Access Token that APIcast will use to download its configuration.	N/A	yes
ADMIN_ACCESS_TOKEN	Admin Access Token with all scopes and write permissions for API access.	N/A	no
WILDCARD_DOMAIN	Root domain for the wildcard routes. For example, a root domain **example.com** will generate **3scale-admin.example.com**.	N/A	yes
WILDCARD_POLICY	Enable wildcard routes to built-in APIcast gateways by setting the value as "Subdomain"	none	yes
TENANT_NAME	Tenant name under the root that Admin UI will be available with -admin suffix.	3scale	yes
MYSQL_USER	Username for MySQL user that will be used for accessing the database.	mysql	yes
MYSQL_PASSWORD	Password for the MySQL user.	N/A	yes
MYSQL_DATABASE	Name of the MySQL database accessed.	system	yes
MYSQL_ROOT_PASSWORD	Password for Root user.	N/A	yes
SYSTEM_BACKEND_USERNAME	Internal 3scale API username for internal 3scale api auth.	3scale_api_user	yes
SYSTEM_BACKEND_PASSWORD	Internal 3scale API password for internal 3scale api auth.	N/A	yes

REDIS_IMAGE	Redis image to use	registry.access.redhat.com/rhscl/redis-32-rhel7:3.2	yes
MYSQL_IMAGE	Mysql image to use	registry.access.redhat.com/rhscl/mysql-57-rhel7:5.7-5	yes
SYSTEM_BACKEND_SHARED_SECRET	Shared secret to import events from backend to system.	N/A	yes
SYSTEM_APP_SECRET_KEY_BASE	System application secret key base	N/A	yes
APICAST_MANAGEMENT_API	Scope of the APIcast Management API. Can be disabled, status or debug. At least status required for health checks.	status	no
APICAST_OPENSSL_VERIFY	Turn on/off the OpenSSL peer verification when downloading the configuration. Can be set to true/false.	false	no
APICAST_RESPONSE_CODES	Enable logging response codes in APIcast.	true	no
APICAST_REGISTRY_URL	A URL which resolves to the location of APIcast policies	http://apicast-staging:8090/policies	yes
MASTER_USER	Master administrator account username	master	yes
MASTER_NAME	The subdomain value for the master admin portal, will be appended with the **-master** suffix	master	yes
MASTER_PASSWORD	A randomly generated master administrator password	N/A	yes
MASTER_ACCESS_TOKEN	A token with master level permissions for API calls	N/A	yes

5.7. USE APICAST WITH AMP ON OPENSHIFT

APIcast with AMP on OpenShift differs from APIcast with AMP hosted and requires unique configuration procedures.

This section explains how to deploy APIcast with AMP on OpenShift.

5.7.1. Deploy APIcast Templates on an Existing OpenShift Cluster Containing your AMP

AMP OpenShift templates contain two built-in APIcast API gateways by default. If you require more API gateways, or require separate APIcast deployments, you can deploy additional APIcast templates onto your OpenShift cluster.

Perform the following steps to deploy additional API gateways onto your OpenShift cluster:

1. Create an access token with the following configurations:
   - scoped to Account Management API
   - has read-only access
2. Log in to your APIcast Cluster:

```
oc login
```

3. Create a secret that allows APIcast to communicate with AMP. Specify **new-basicauth**, **apicast-configuration-url-secret**, and the **--password** parameter with the access token, tenant name, and wildcard domain of your AMP deployment:

```
oc secret new-basicauth apicast-configuration-url-secret --
password=https://<APICAST_ACCESS_TOKEN>@<TENANT_NAME>-admin.
<WILDCARD_DOMAIN>
```

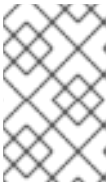

NOTE

TENANT_NAME is the name under the root that Admin UI will be available with. The default value for **TENANT_NAME** "3scale". If you used a custom value in your AMP deployment then you must use that value here.

4. Import the APIcast template by downloading the apicast.yml, located on the 3scale GitHub, and running the **oc new-app** command, specifying the **--file** option with the **apicast.yml** file:

```
oc new-app --file /path/to/file/apicast.yml
```

5.7.2. Connect APIcast from an OpenShift Cluster Outside of an OpenShift Cluster Containing your AMP

If you deploy APIcast onto a different OpenShift cluster, outside of your AMP cluster, you must connect over the public route.

1. Create an access token with the following configurations:
   - scoped to Account Management API

- has read-only access

2. Log in to your APIcast Cluster:

```
oc login
```

3. Create a secret that allows APIcast to communicate with AMP. Specify **new-basicauth**, **apicast-configuration-url-secret**, and the **--password** parameter with the access token, tenant name, and wildcard domain of your AMP deployment:

```
oc secret new-basicauth apicast-configuration-url-secret --
password=https://<APICAST_ACCESS_TOKEN>@<TENANT_NAME>-admin.
<WILDCARD_DOMAIN>
```

NOTE

TENANT_NAME is the name under the root that Admin UI will be available with. The default value for**TENANT_NAME** is "3scale". If you used a custom value in your AMP deployment then you must use that value here.

4. Deploy APIcast onto an OpenShift cluster outside of the OpenShift Cluster with the oc new-app command. Specify the **--file** option and the file path of your **apicast.yml** file:

```
oc new-app --file /path/to/file/apicast.yml
```

5. Update the apicast **BACKEND_ENDPOINT_OVERRIDE** environment variable set to the URL **backend.** followed by the wildcard domain of the OpenShift Cluster containing your AMP deployment:

```
oc env dc/apicast --overwrite BACKEND_ENDPOINT_OVERRIDE=https://backend-
<TENANT_NAME>.<WILDCARD_DOMAIN>
```

5.7.3. Connect APIcast from Other Deployments

Once you have deployed APIcast on other platforms, you can connect them to AMP on OpenShift by configuring the **BACKEND_ENDPOINT_OVERRIDE** environment variable in your AMP OpenShift Cluster:

1. Log in to your AMP OpenShift Cluster:

```
oc login
```

2. Configure the system-app object **BACKEND_ENDPOINT_OVERRIDE** environment variable: If you are using a native installation:

```
BACKEND_ENDPOINT_OVERRIDE=https://backend.<your_openshift_subdomain>
bin/apicast
```

If are using the Docker containerized environment:

```
docker run -e BACKEND_ENDPOINT_OVERRIDE=https://backend.
<your_openshift_subdomain>
```

5.7.4. Change Built-In APIcast Default Behavior

In external APIcast deployments, you can modify default behavior by changing the template parameters in the APIcast OpenShift template.

In built-in APIcast deployments, AMP and APIcast are deployed from a single template. You must modify environment variables after deployment if you wish to change the default behavior for the built-in APIcast deployments.

5.7.5. Connect Multiple APIcast Deployments on a Single OpenShift Cluster over Internal Service Routes

If you deploy multiple APIcast gateways into the same OpenShift cluster, you can configure them to connect using internal routes through the backend listener service instead of the default external route configuration.

You must have an OpenShift SDN plugin installed to connect over internal service routes. How you connect depends on which SDN you have installed.

ovs-subnet

If you are using the **ovs-subnet** OpenShift SDN plugin, follow these steps to connect over internal routes:

1. If not already logged in, log in to your OpenShift Cluster:

```
oc login
```

2. Enter the **oc new-app** command with the path to the **apicast.yml** file:
   - Specify the **--param** option with the **BACKEND_ENDPOINT_OVERRIDE** parameter set to the domain of your OpenShift cluster's AMP project:

```
oc new-app -f apicast.yml --param BACKEND_ENDPOINT_OVERRIDE=http://backend-listener.
<AMP_PROJECT>.svc.cluster.local:3000
```

ovs-multitenant

If you are using the 'ovs-multitenant' Openshift SDN plugin, follow these steps to connect over internal routes:

1. If not already logged in, log in to your OpenShift Cluster:

```
oc login
```

2. As admin, specify the **oadm** command with the **pod-network** and **join-projects** options to set up communication between both projects:

```
oadm pod-network join-projects --to=<AMP_PROJECT> <APICAST_PROJECT>
```

3. Enter the **oc new-app** cotion with the path to the **apicast.yml** file
   - Specify the **--param** option with the **BACKEND_ENDPOINT_OVERRIDE** parameter set to the domain of your OpenShift cluster's AMP project:

```
oc new-app -f apicast.yml --param BACKEND_ENDPOINT_OVERRIDE=http://backend-listener.
<AMP_PROJECT>.svc.cluster.local:3000
```

More information

For information on Openshift SDN and project network isolation, see: Openshift SDN

5.8. 7. TROUBLESHOOTING

This section contains a list of common installation issues and provides guidance for their resolution.

- Previous Deployment Leaves Dirty Persistent Volume Claims
- Incorrectly Pulling from the Docker Registry
- Permissions Issues for MySQL when Persistent Volumes are Mounted Locally
- Unable to Upload Logo or Images Because Persistent Volumes are not Writable by OpenShift
- Create Secure Routes on OpenShift
- APIcast on a Different Project from AMP Fails to Deploy Due to Problem with Secrets

5.8.1. Previous Deployment Leaves Dirty Persistent Volume Claims

Problem

A previous deployment attempt leaves a dirty Persistent Volume Claim (PVC) causing the MySQL container to fail to start.

Cause

Deleting a project in OpenShift does not clean the PVCs associated with it.

Solution

1. Find the PVC containing the erroneous MySQL data with **oc get pvc**:

```
# oc get pvc
NAME                    STATUS    VOLUME    CAPACITY   ACCESSMODES   AGE
backend-redis-storage   Bound     vol003    100Gi      RWO,RWX       4d
mysql-storage           Bound     vol006    100Gi      RWO,RWX       4d
system-redis-storage    Bound     vol008    100Gi      RWO,RWX       4d
system-storage          Bound     vol004    100Gi      RWO,RWX       4d
```

2. Stop the deployment of the system-mysql pod by clicking **cancel deployment** in the OpenShift UI.
3. Delete everything under the MySQL path to clean the volume.
4. Start a new **system-mysql** deployment.

5.8.2. Incorrectly Pulling from the Docker Registry

Problem

The following error occurs during installation:

```
svc/system-redis - 1EX.AMP.LE.IP:6379
  dc/system-redis deploys docker.io/rhscl/redis-32-rhel7:3.2-5.3
    deployment #1 failed 13 minutes ago: config change
```

Cause

OpenShift searches for and pulls container images by issuing the **docker** command. This command refers to the **docker.io** Docker registry instead of the **registry.access.redhat.com** Red Hat container registry.

This occurs when the system contains an unexpected version of the Docker containerized environment.

Solution

Use the appropriate version of the Docker containerized environment.

5.8.3. Permissions Issues for MySQL when Persistent Volumes are Mounted Locally

Problem

The system-msql pod crashes and does not deploy causing other systems dependant on it to fail deployment. The pod log displays the following error:

```
[ERROR] Can't start server : on unix socket: Permission denied
[ERROR] Do you already have another mysqld server running on socket: /var/lib/mysql/mysql.sock ?
[ERROR] Aborting
```

Cause

The MySQL process is started with inappropriate user permissions.

Solution

1. The directories used for the persistent volumes MUST have the write permissions for the root group. Having rw permissions for the root user is not enough as the MySQL service runs as a different user in the root group. Execute the following command as the root user:

   ```
   chmod -R g+w /path/for/pvs
   ```

2. Execute the following command to prevent SElinux from blocking access:

   ```
   chcon -Rt svirt_sandbox_file_t /path/for/pvs
   ```

5.8.4. Unable to Upload Logo or Images because Persistent Volumes are not Writable by OpenShift

Problem

Unable to upload a logo - **system-app** logs display the following error:

```
Errno::EACCES (Permission denied @ dir_s_mkdir - /opt/system/public//system/provider-name/2
```

Cause

Persistent volumes are not writable by OpenShift.

Solution

Ensure your persistent volume is writable by OpenShift. It should be owned by root group and be group writable.

5.8.5. Create Secure Routes on OpenShift

Problem

Test calls do not work after creation of a new service and routes on OpenShift. Direct calls via curl also fail, stating: **service not available**.

Cause

3scale requires HTTPS routes by default, and OpenShift routes are not secured.

Solution

Ensure the "secure route" checkbox is enabled in your OpenShift router settings.

5.8.6. APIcast on a Different Project from AMP Fails to Deploy due to Problem with Secrets

Problem

APIcast deploy fails (pod doesn't turn blue). The following error appears in the logs:

```
update acceptor rejected apicast-3: pods for deployment "apicast-3" took longer than 600 seconds to become ready
```

The following error appears in the pod:

```
Error synching pod, skipping: failed to "StartContainer" for "apicast" with RunContainerError: "GenerateRunContainerOptions: secrets \"apicast-configuration-url-secret\" not found"
```

Cause

The secret was not properly set up.

Solution

When creating a secret with APIcast v3, specify **apicast-configuration-url-secret**:

```
oc secret new-basicauth apicast-configuration-url-secret  --password=https://<ACCESS_TOKEN>@<TENANT_NAME>-admin.<WILDCARD_DOMAIN>
```

CHAPTER 6. RED HAT 3SCALE AMP 2.2 ON-PREMISES OPERATIONS AND SCALING GUIDE

6.1. INTRODUCTION

This document describes operations and scaling tasks of a Red Hat 3scale AMP 2.2 On-Premises installation.

6.1.1. Prerequisites

Before you can perform the steps in this guide, you must have installed and initially configured AMP On-Premises on a supported OpenShift version.

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

6.1.2. Further Reading

- Health and Liveness Monitoring
- OpenShift Documentation

6.2. RE-DEPLOYING APICAST

Once you have deployed AMP On-Premises and your chosen APIcast deployment method, you can test and promote system changes through your AMP dashboard. 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
4. 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, as well as configure how often your gateway retrieves published changes. Once 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
5. APIcast will retrieve and publish the promoted update at the configured frequency

6.3. APICAST BUILT-IN WILDCARD ROUTING (TECH PREVIEW)

The built-in APIcast gateways that accompany your on-preimses AMP 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 public base URLs. In order to use this feature, you must have enabled it during your on-premises installation.

NOTE

Wildcard routing is in tech preview. The current tech preview contains the following limitations:

- Any HTTP headers containing underscores will cause the service to fail with a 403 error code. As a workaround, remove underscores from all header names.
- You must set the template parameter **TENANT_NAME** to a value that does not start with a number

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.

6.3.1. 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, refer to the OpenShift documentation.

6.4. SCALING UP AMP ON PREMISES

6.4.1. 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 either of the following methods:

6.4.1.1. 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

6.4.1.2. Method 2. Back up and Redeploy AMP

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:
   - The size of your new PVC
   - 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.

6.4.2. Scaling up Performance

6.4.2.1. Configuring 3scale On-Premises Deployments

By default, 3scale deployments run 1 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 the desired number of processes per pod
   - **APICAST_WORKERS** for APIcast pods (Red Hat recommends no more than 2 per deployment)
   - **PUMA_WORKERS** for backend pods
   - **UNICORN_WORKERS** for system pods

```
oc env dc/apicast --overwrite APICAST_WORKERS=<number_of_processes>
```

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

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

6.4.2.2. 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. As long as your 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.

6.4.2.3. 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.

6.4.2.4. Further Reading

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

- Routers

6.5. OPERATIONS TROUBLESHOOTING

6.5.1. Access 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 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 2 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
```

6.5.2. Job Queues

Job Queues contain logs of information sent from the **system-resque** and **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>
```

CHAPTER 7. HOW TO DEPLOY A FULL-STACK API SOLUTION WITH FUSE, 3SCALE, AND OPENSHIFT

This tutorial describes how to get a full-stack API solution (API design, development, hosting, access control, monetization, etc.) using Red Hat JBoss xPaaS for OpenShift and 3scale API Management Platform - Cloud.

The tutorial is based on a collaboration between Red Hat and 3scale to provide a full-stack API solution. This solution includes design, development, and hosting of your API on the Red Hat JBoss xPaaS for OpenShift, combined with the 3scale API Management Platform for full control, visibility, and monetization features.

The API itself can be deployed on Red Hat JBoss xPaaS for OpenShift, which can be hosted in the cloud as well as on premise (that's the Red Hat part). The API management (the 3scale part) can be hosted on Amazon Web Services (AWS), using 3scale APIcast or OpenShift. This gives a wide range of different configuration options for maximum deployment flexibility.

The diagram below summarizes the main elements of this joint solution. It shows the whole integration chain including enterprise backend systems, middleware, API management, and API customers.

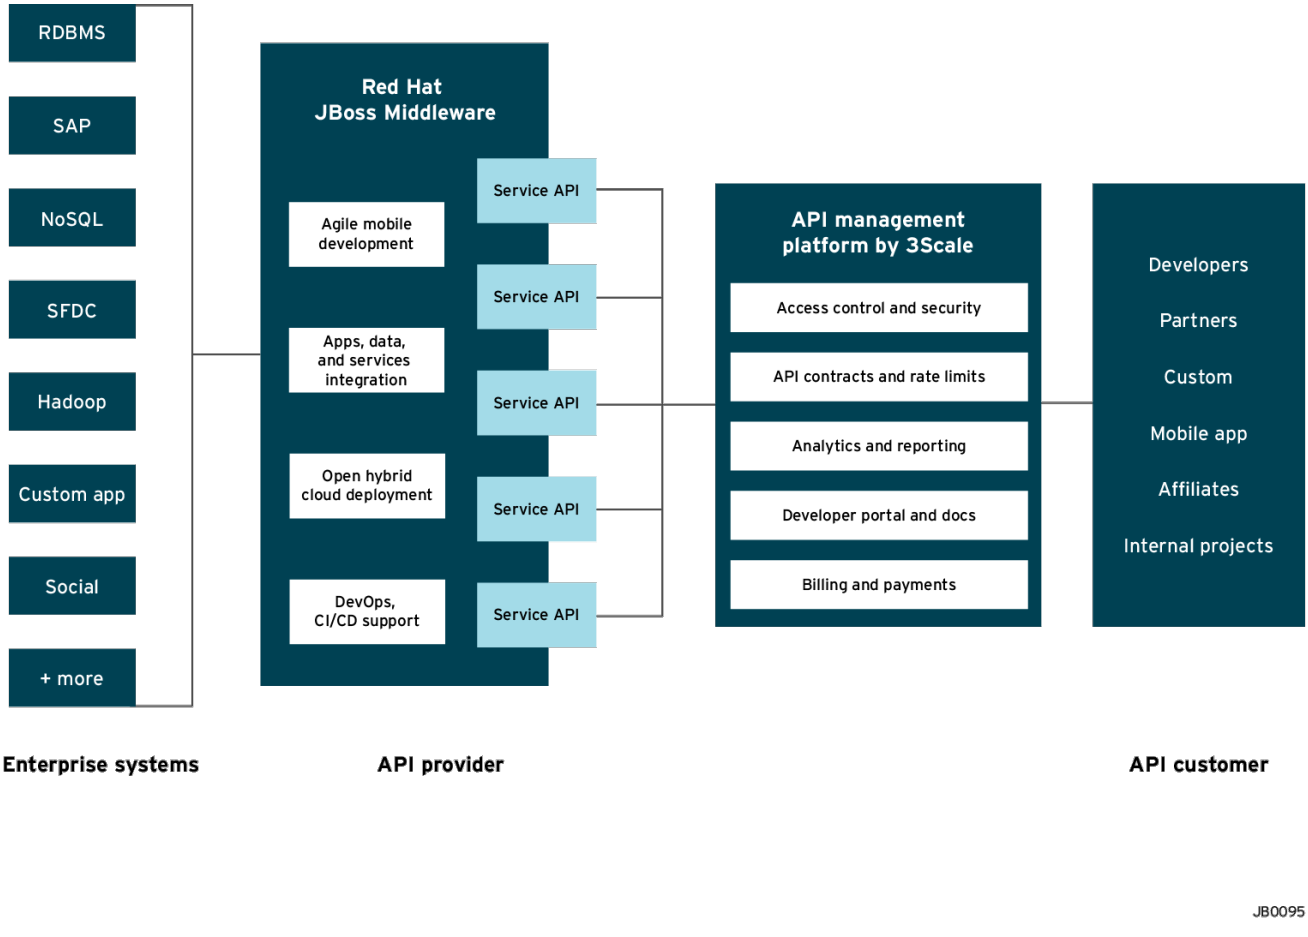

For specific support questions, please contact support.

This tutorial shows three different deployment scenarios step by step:

1. Scenario 1 – A Fuse on OpenShift application containing the API. The API is managed by 3scale with the API gateway hosted on Amazon Web Services (AWS) using the 3scale AMI.

2. Scenario 2 – A Fuse on OpenShift application containing the API. The API is managed by 3scale with the API gateway hosted on APIcast (3scale's cloud hosted API gateway).

3. Scenario 3 – A Fuse on OpenShift application containing the API. The API is managed by 3scale with the API gateway hosted on OpenShift

This tutorial is split into four parts:

- Part 1: Fuse on OpenShift setup to design and implement the API
- Part 2: Configuration of 3scale API Management
- Part 3: Integration of your API services
- Part 4: Testing the API and API management

The diagram below shows the roles the various parts play in this configuration.

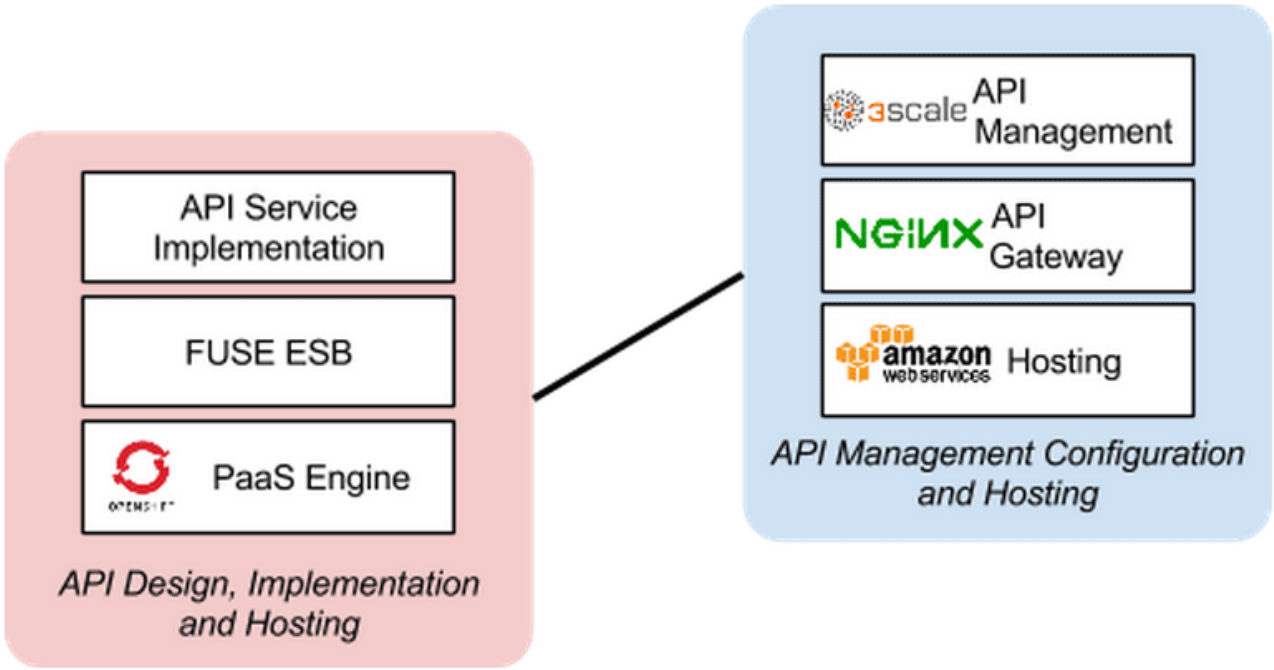

7.1. PART 1: FUSE ON OPENSHIFT SETUP

You will create a Fuse on OpenShift application that contains the API to be managed. You will use the REST quickstart that is included with Fuse 6.1. This requires a medium or large gear, as using the small gear will result in memory errors and/or horrible performance.

7.1.1. Step 1

Sign in to your OpenShift online account. Sign up for an OpenShift online account if you don't already have one.

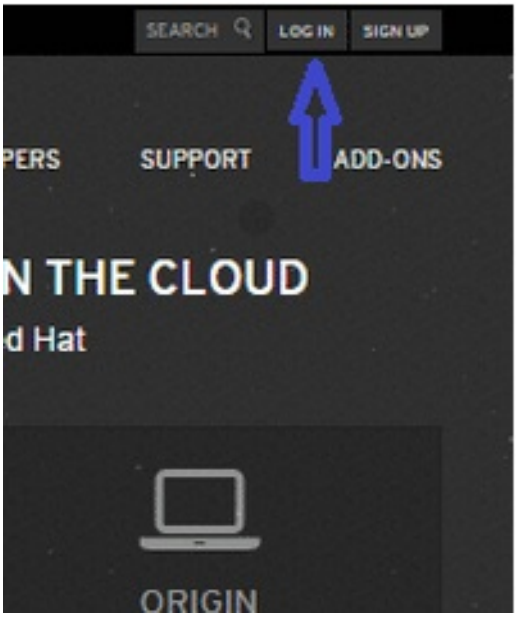

7.1.2. Step 2

Click the "add application" button after signing in.

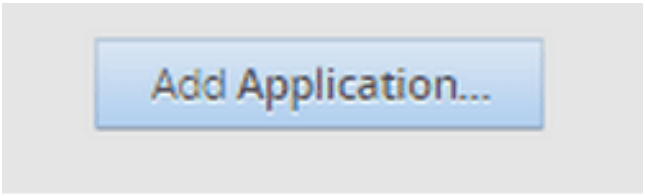

7.1.3. Step 3

Under xPaaS, select the Fuse type for the application.

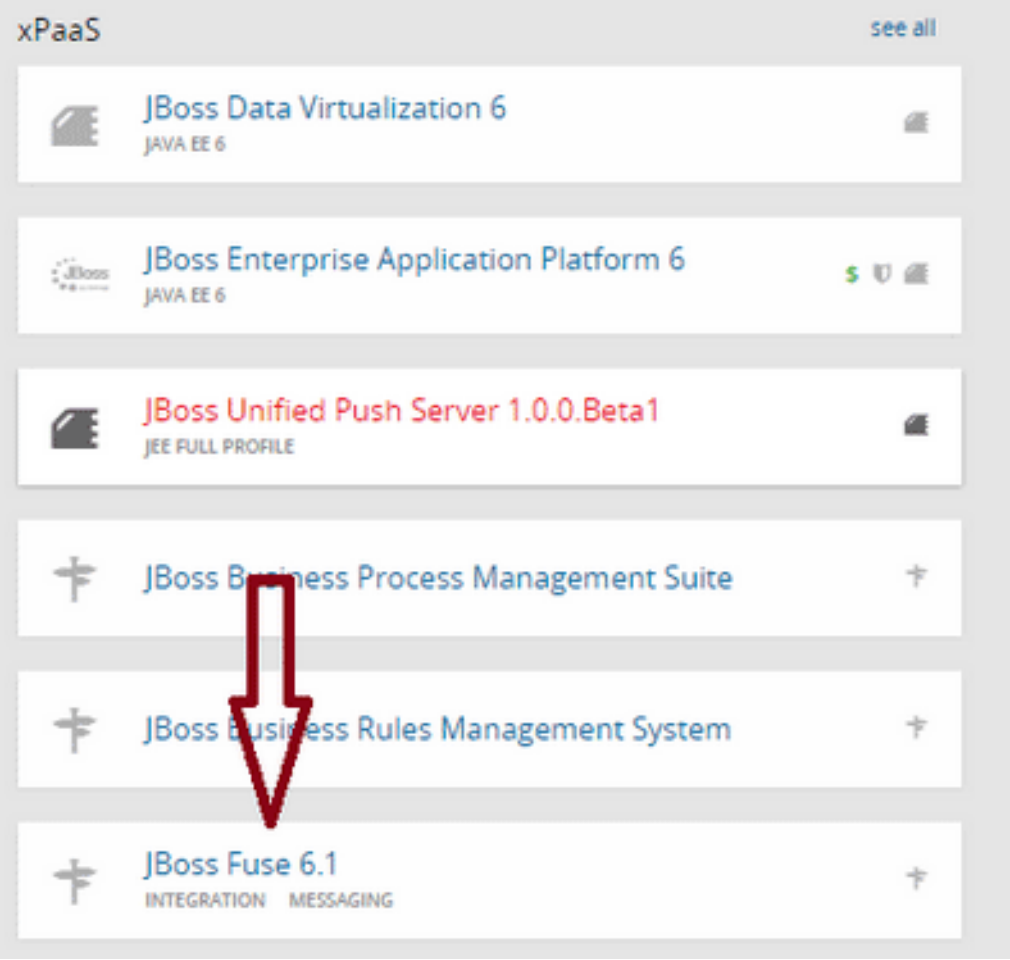

7.1.4. Step 4

Now configure the application. Enter the subdomain you'd like your application to show up under, such as "restapitest". This will give a full URL of the form "appname-domain.rhcloud.com" – in the example below "restapitest-ossmentor.rhcloud.com". Change the gear size to medium or large, which is required for the Fuse cartridge. Now click on "create application".

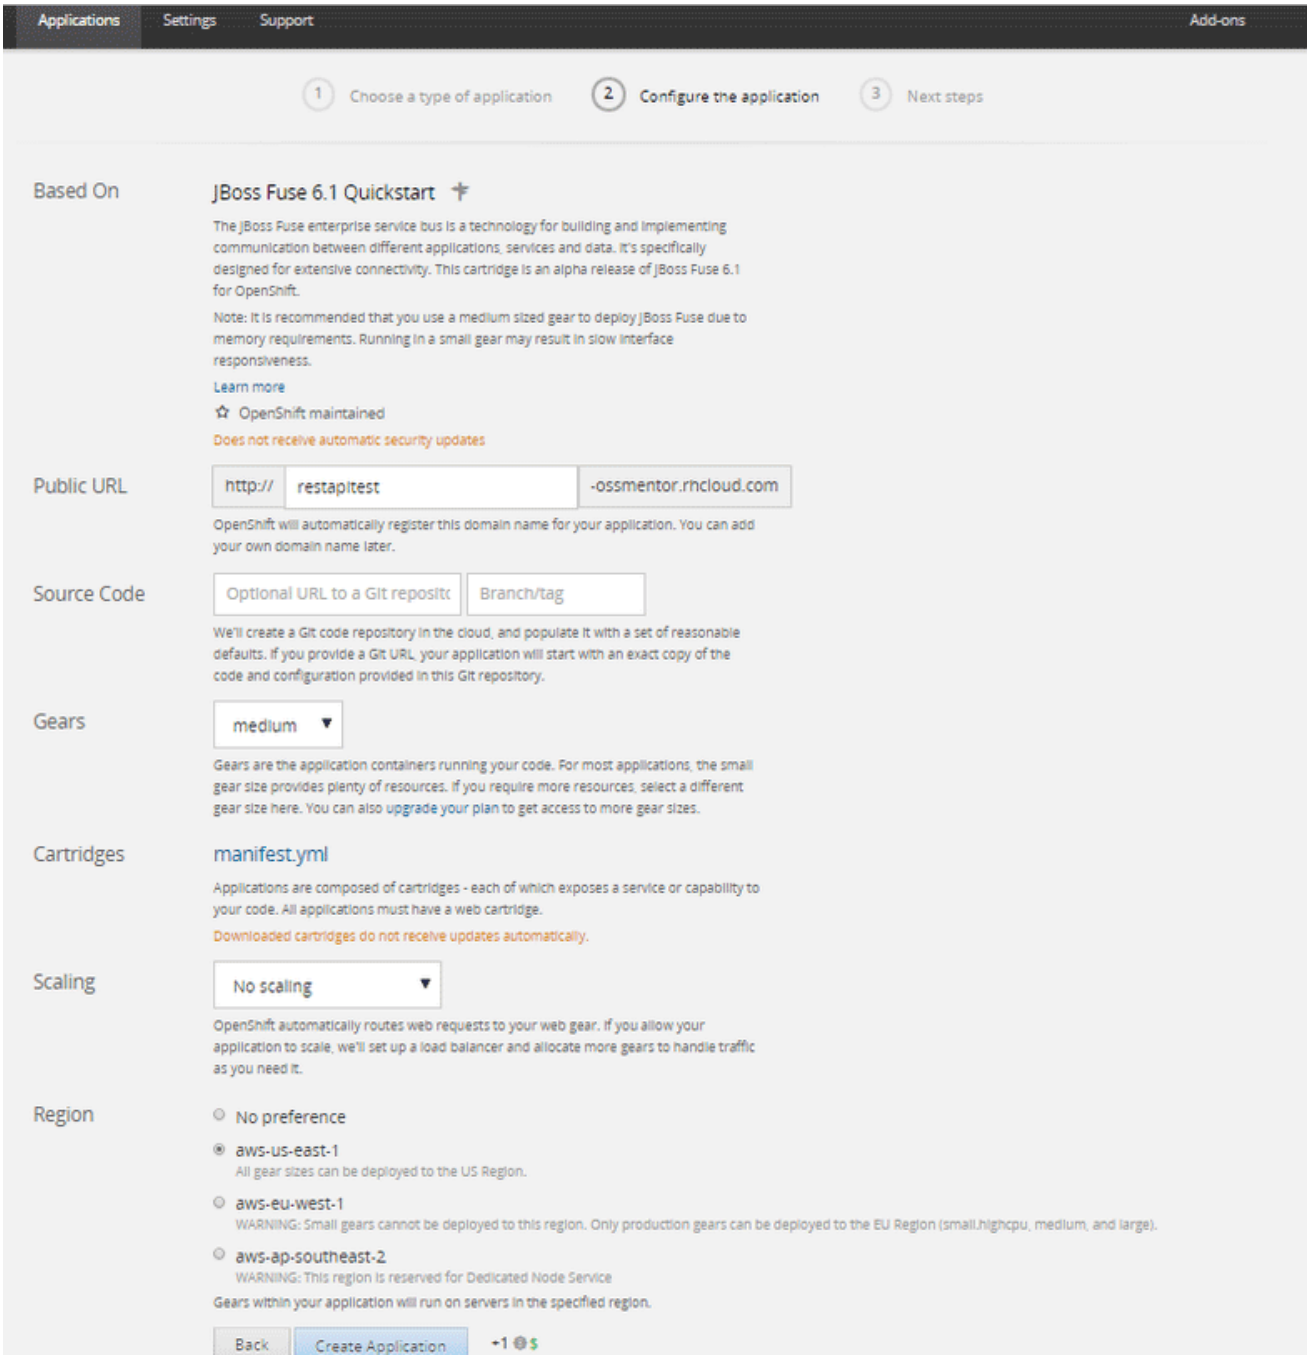

7.1.5. Step 5

Click "create application".

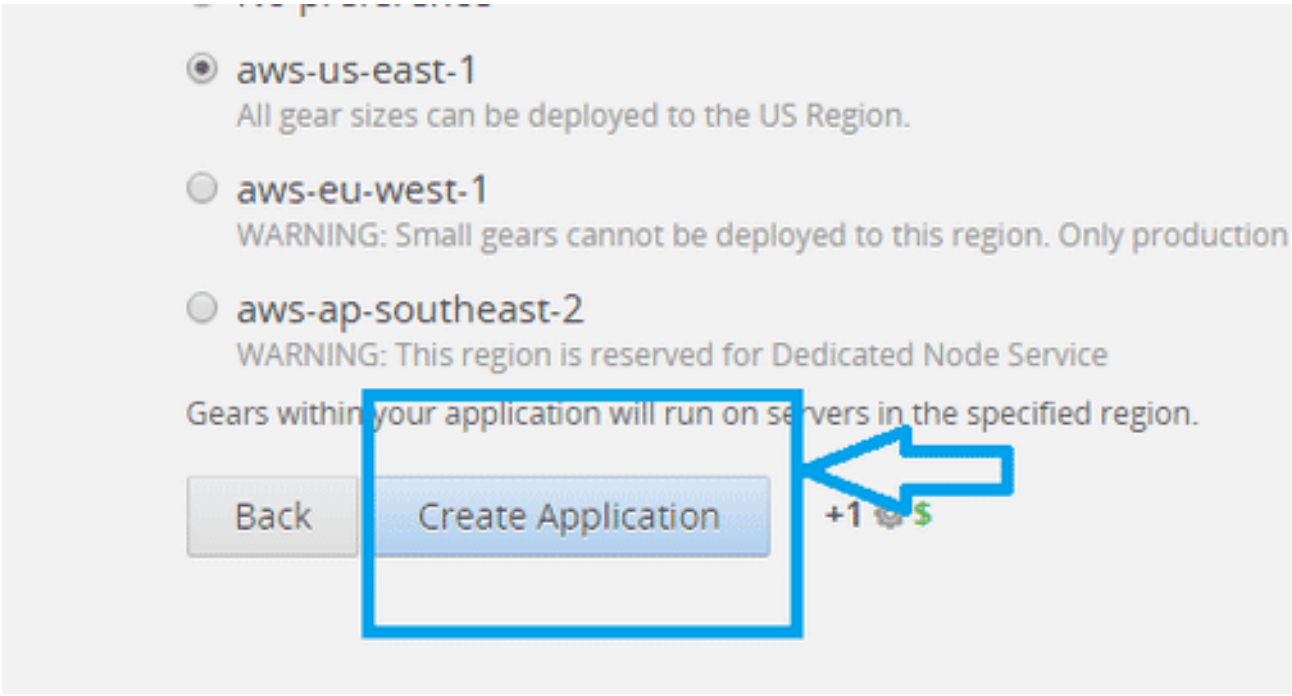

7.1.6. Step 6

Browse the application hawtio console and sign in.

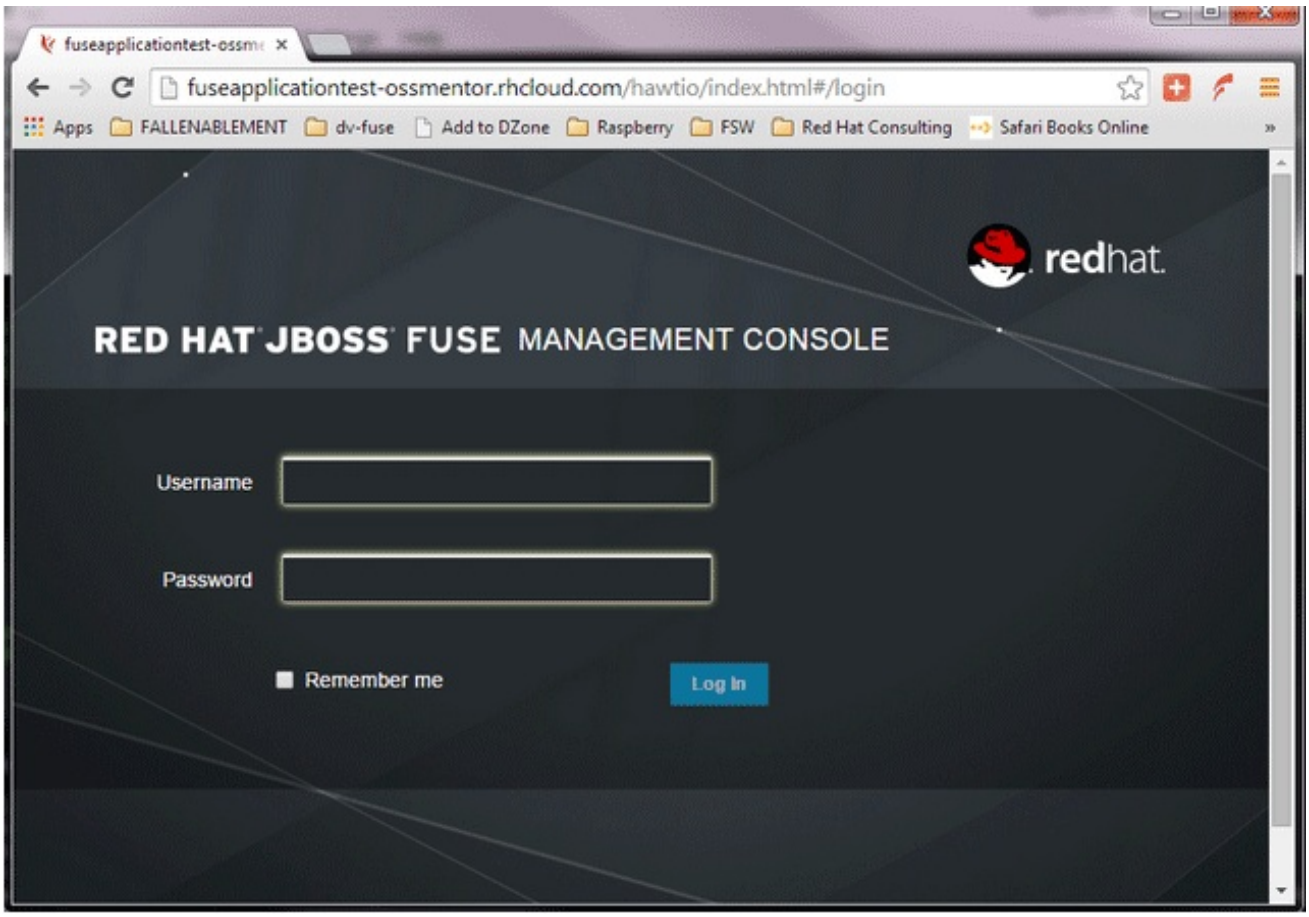

7.1.7. Step 7

After signing in, click on the "runtime" tab and the container, and add the REST API example.

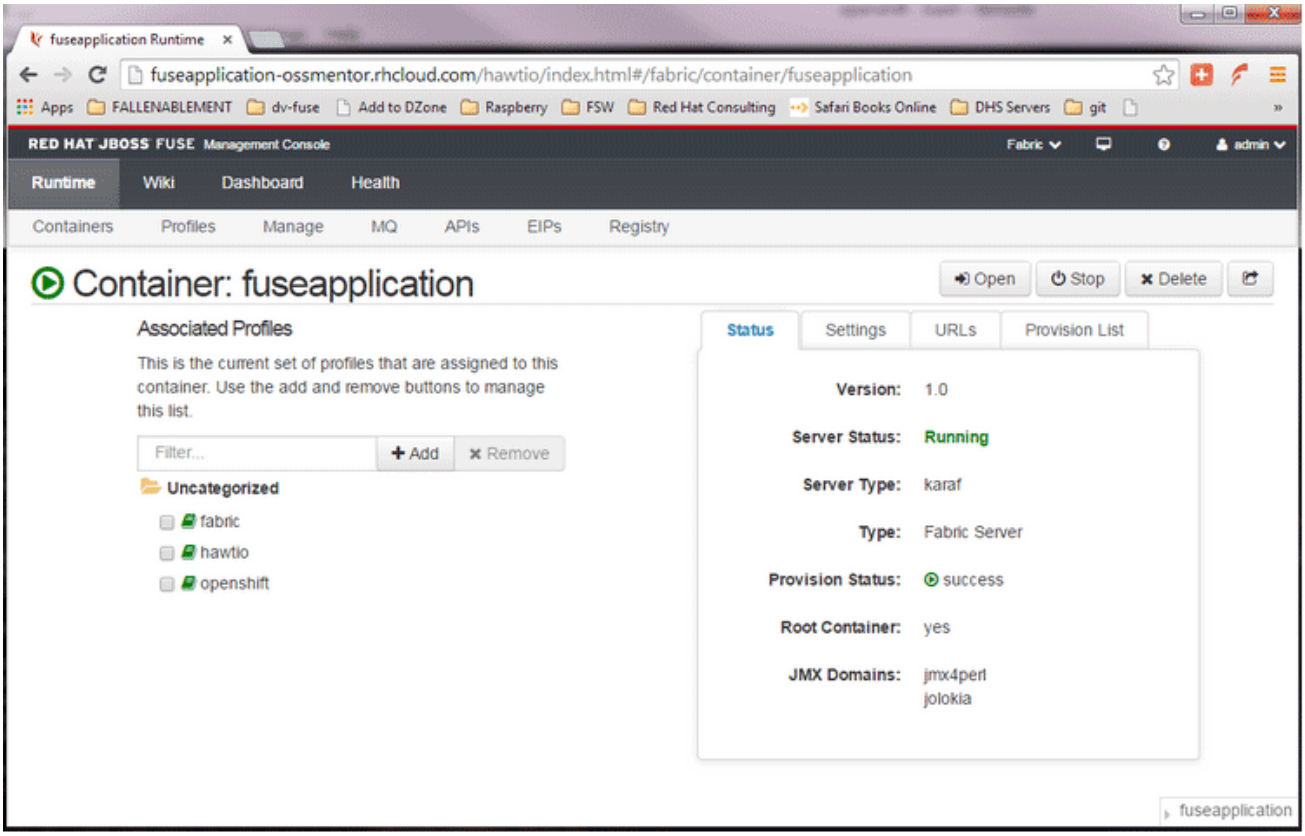

7.1.8. Step 8

Click on the "add a profile" button.

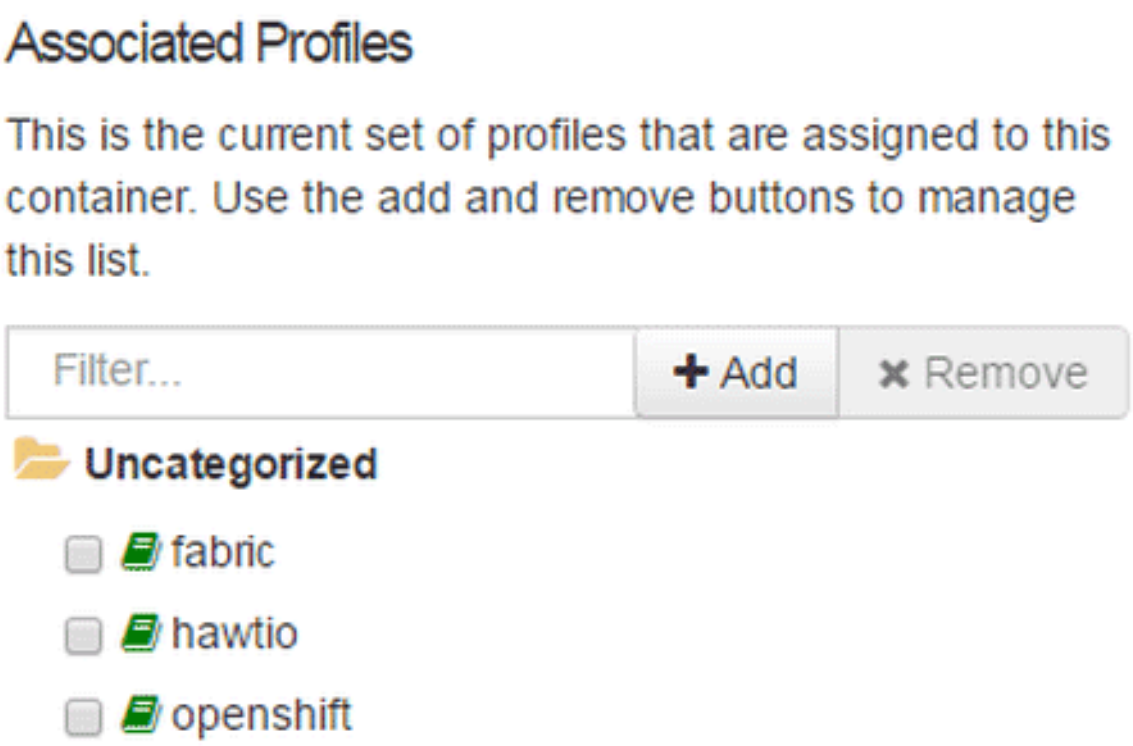

7.1.9. Step 9

Scroll down to examples/quickstarts and click the "REST" checkbox, then "add". The REST profile should show up on the container associated profile page.

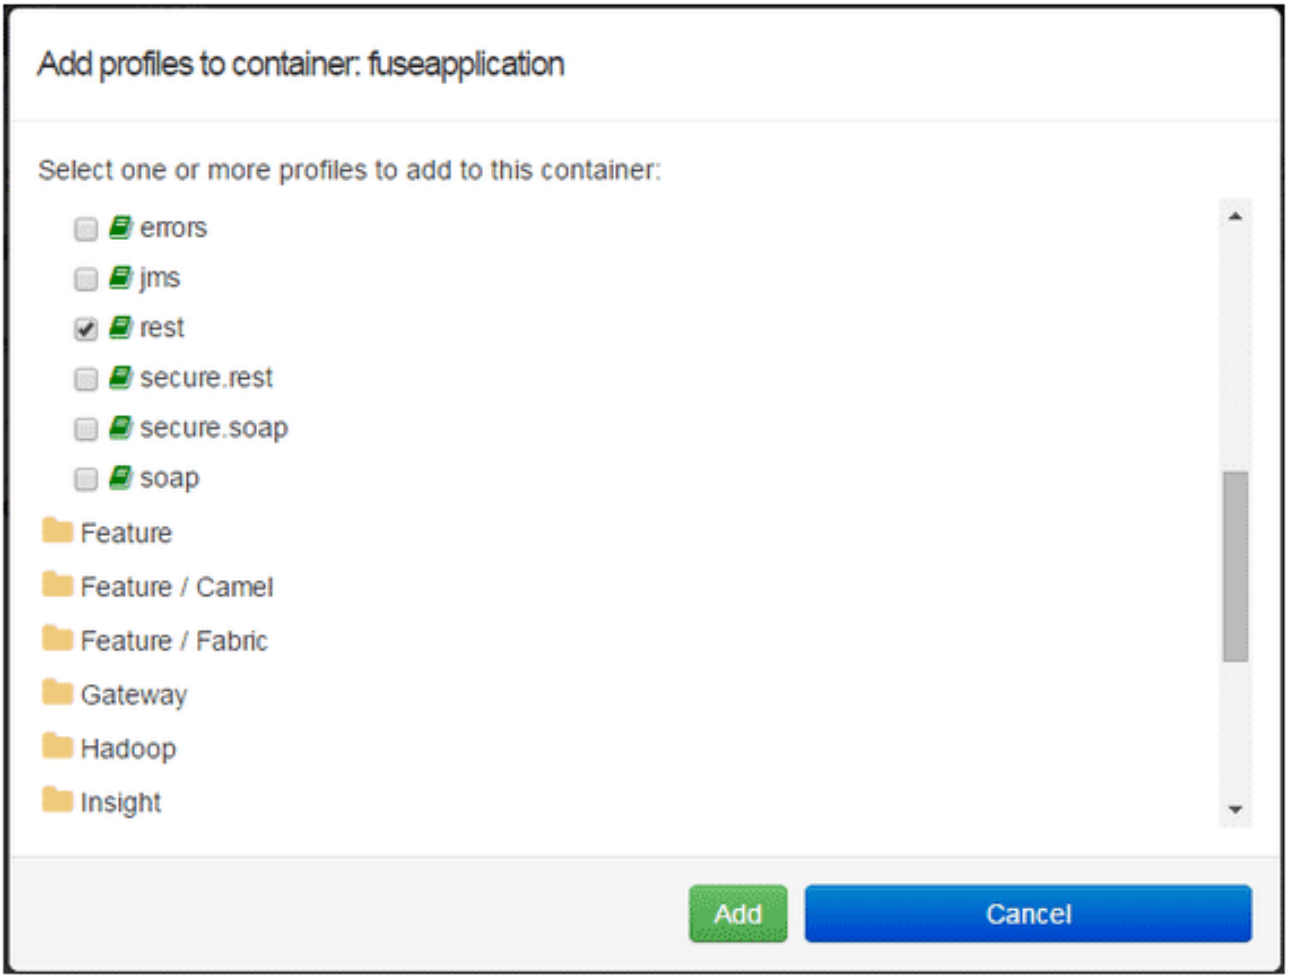

7.1.10. Step 10

Click on the runtime/APIs tab to verify the REST API profile.

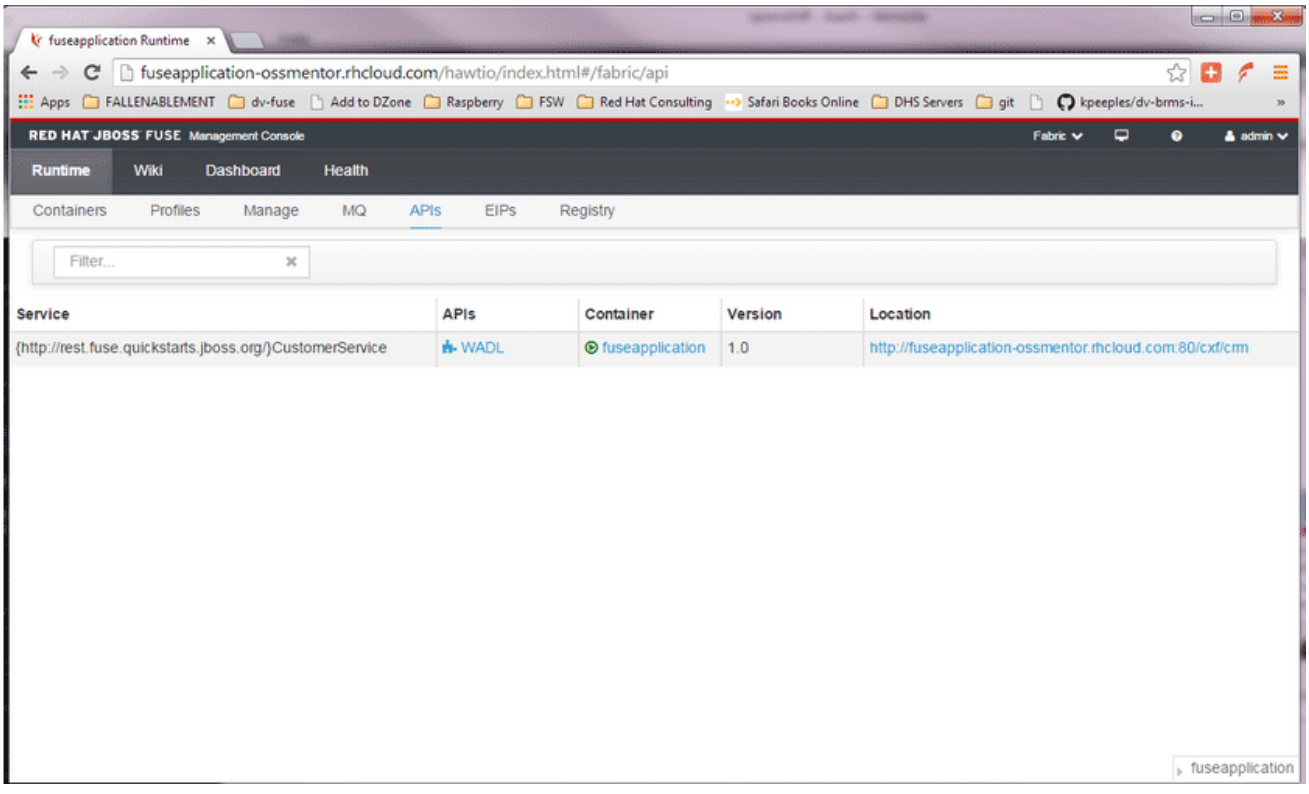

7.1.11. Step 11

Verify the REST API is working. Browse to customer 123, which will return the ID and name in XML format.

7.2. PART 2: CONFIGURE 3SCALE API MANAGEMENT

To protect the API that you just created in Part 1 using 3scale API Management, you first must conduct the according configuration, which is then later deployed according to one of the three scenarios presented.

Once you have your API set up on OpenShift, you can start setting it up on 3scale to provide the management layer for access control and usage monitoring.

7.2.1. Step 1

Log in to your 3scale account. You can sign up for a 3scale account at www.3scale.net if you don't already have one. When you log in to your account for the first time, follow the wizard to learn the basics about integrating your API with 3scale.

7.2.2. Step 2

In **API > Integration**, you can enter the public URL for the Fuse application on OpenShift that you just created, e.g. "restapitest-ossmentor.rhcloud.com" and click on **Test**. This will test your setup against the 3scale API Gateway in the staging environment. The staging API gateway allows you to test your 3scale setup before deploying your proxy configuration to AWS.

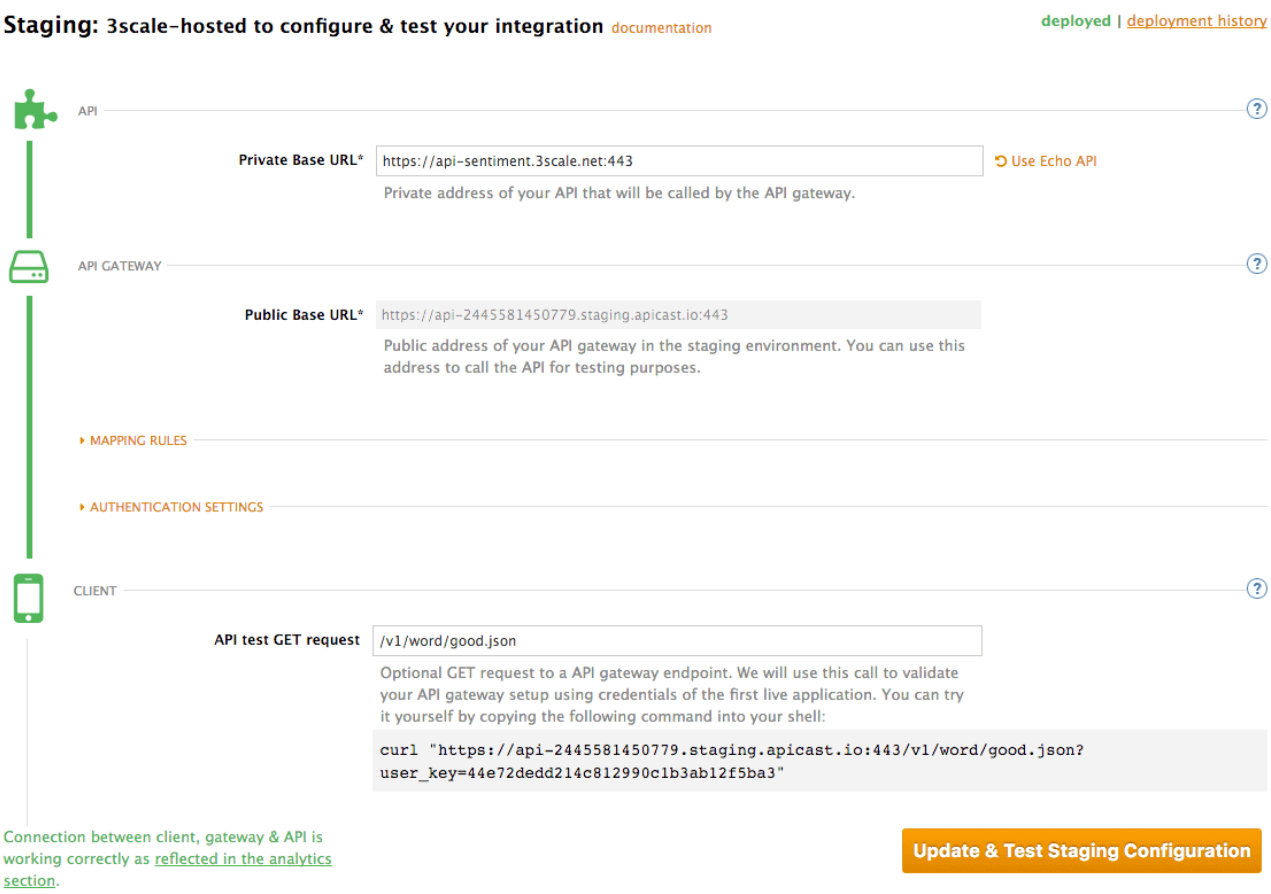

7.2.3. Step 3

The next step is to set up the API methods that you want to monitor and rate limit. To do that go to **API > Definition** and click on 'New method'.

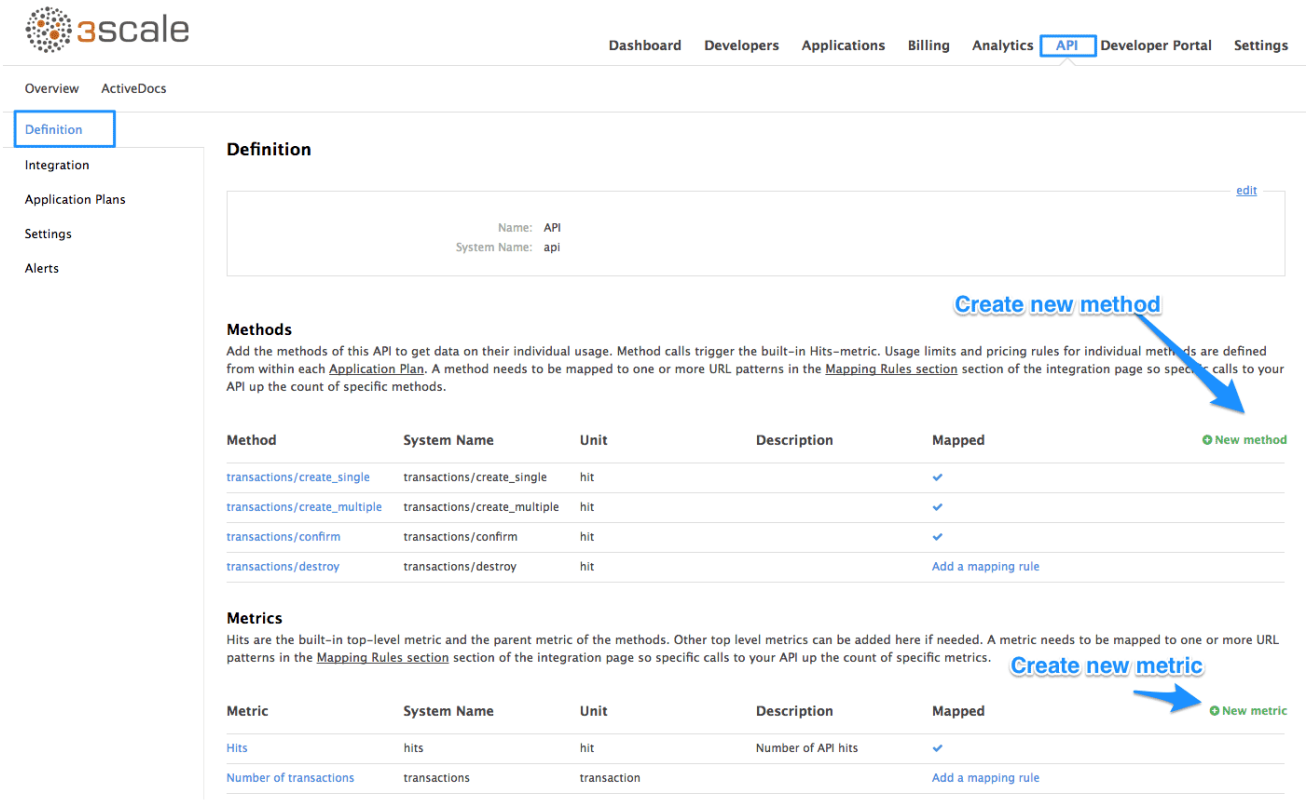

For more details on creating methods, visit our API definition tutorial.

7.2.4. Step 4

Once you have all of the methods that you want to monitor and control set up under the application plan, you'll need to map these to actual HTTP methods on endpoints of your API. Go back to the integration page and expand the "mapping rules" section.

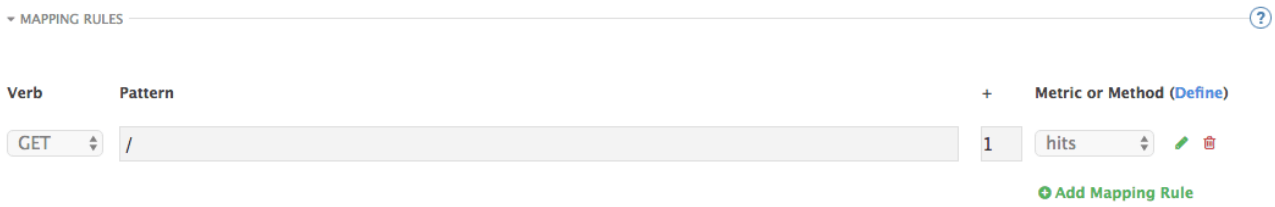

Create mapping rules for each of the methods you created under the application plan.

Once you have done that, your mapping rules will look something like this:

For more details on mapping rules, visit our tutorial about mapping rules.

7.2.5. Step 5

Once you've clicked "update and test" to save and test your configuration, you are ready to download the set of configuration files that will allow you to configure your API gateway on AWS. For the API gateway, you should use a high-performance, open-source proxy called nginx. You will find the necessary configuration files for nginx on the same integration page by scrolling down to the "production" section.

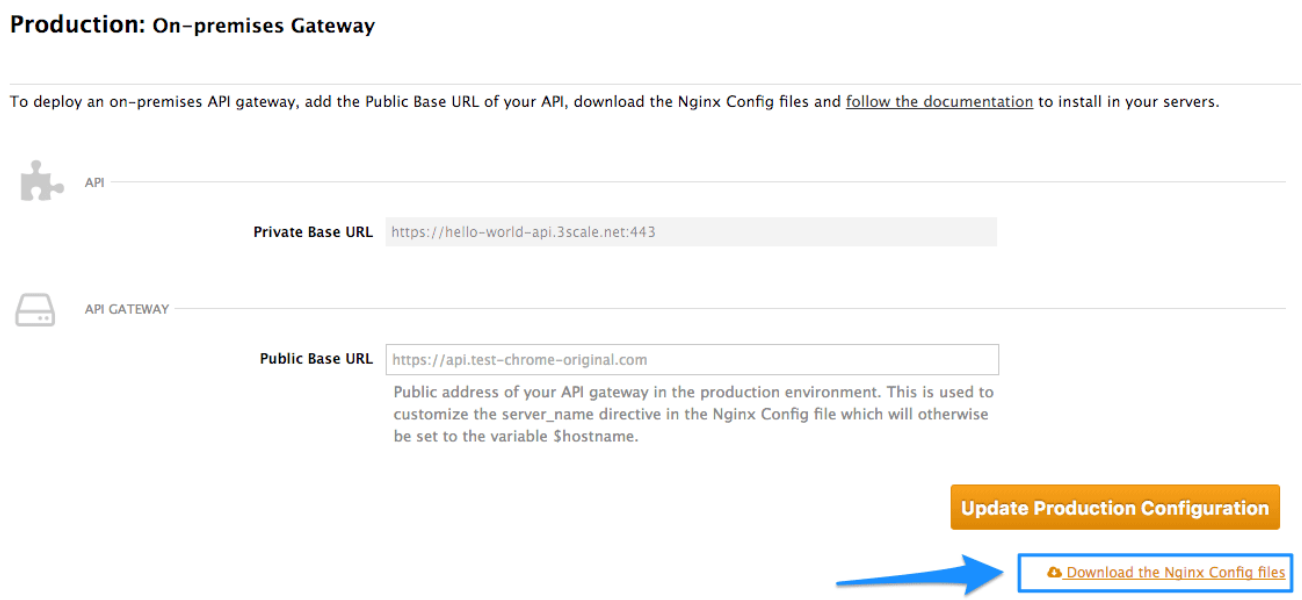

The next section will now take you through various hosting scenarios.

7.3. PART 3: INTEGRATION OF YOUR API SERVICES

There are different ways in which you can integrate your API services in 3scale. Choose the one that best fits your needs:

- APIcast hosted on AWS
- APIcast hosted
- APIcast on OpenShift

7.4. PART 4: TESTING THE API AND API MANAGEMENT

Testing the correct functioning of the API and the API Management is independent from the chosen scenario. You can use your favorite REST client and run the following commands.

7.4.1. Step 1

Retrieve the customer instance with id 123.

```
http://54.149.46.234/cxf/crm/customerservice/customers/123?
user_key=b9871b41027002e68ca061faeb2f972b
```

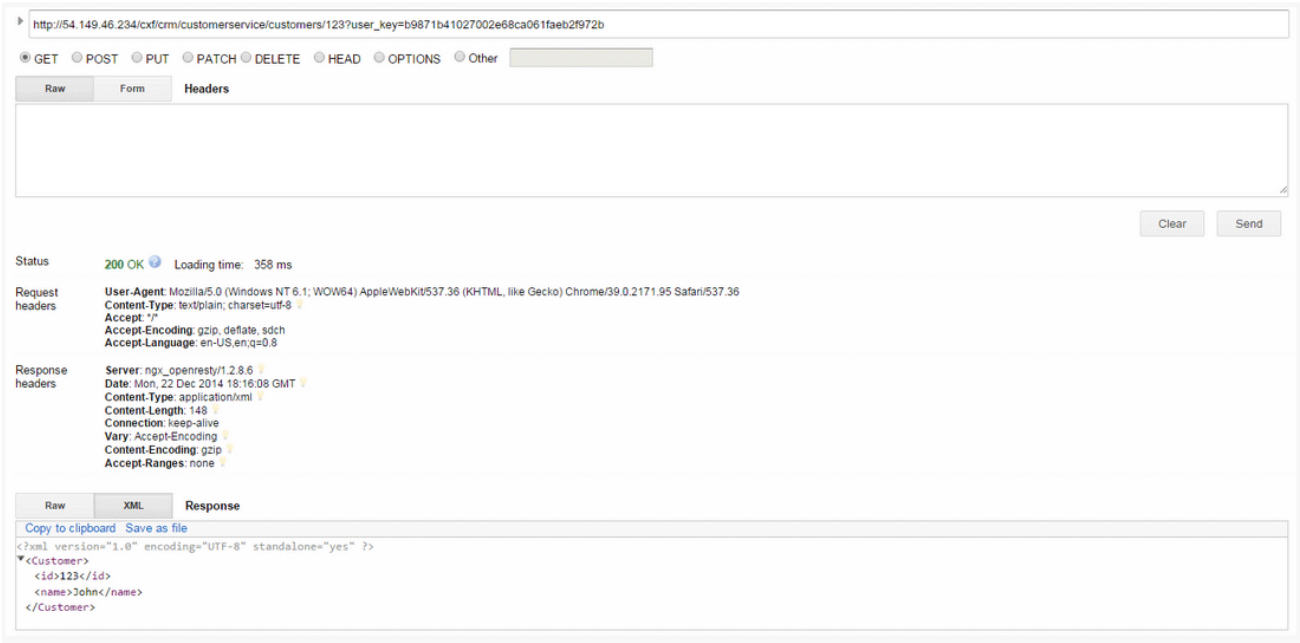

7.4.2. Step 2

Create a customer.

> http://54.149.46.234/cxf/crm/customerservice/customers?
> user_key=b9871b41027002e68ca061faeb2f972b

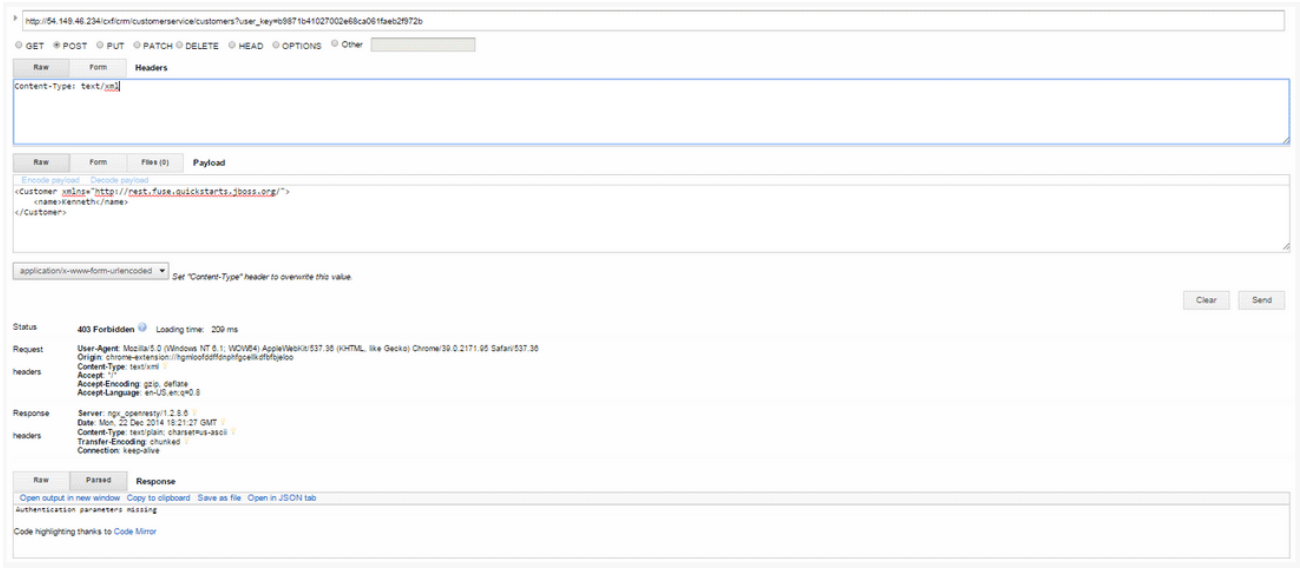

7.4.3. Step 3

Update the customer instance with id 123.

> http://54.149.46.234/cxf/crm/customerservice/customers?
> user_key=b9871b41027002e68ca061faeb2f972b

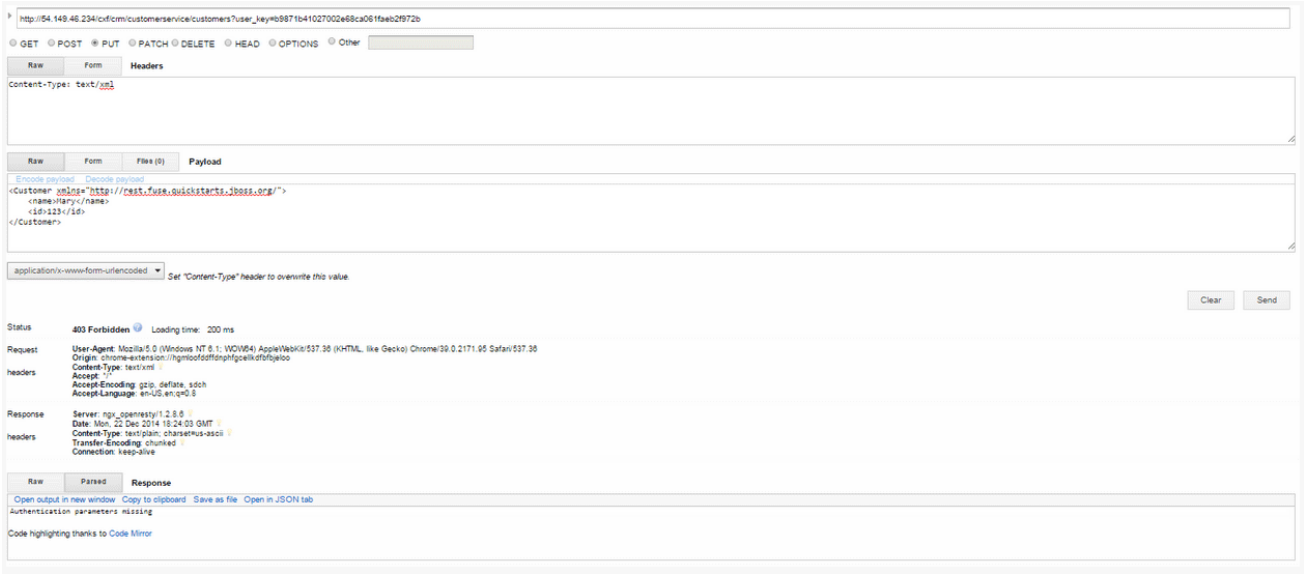

7.4.4. Step 4

Delete the customer instance with id 123.

> http://54.149.46.234/cxf/crm/customerservice/customers/123?
> user_key=b9871b41027002e68ca061faeb2f972b

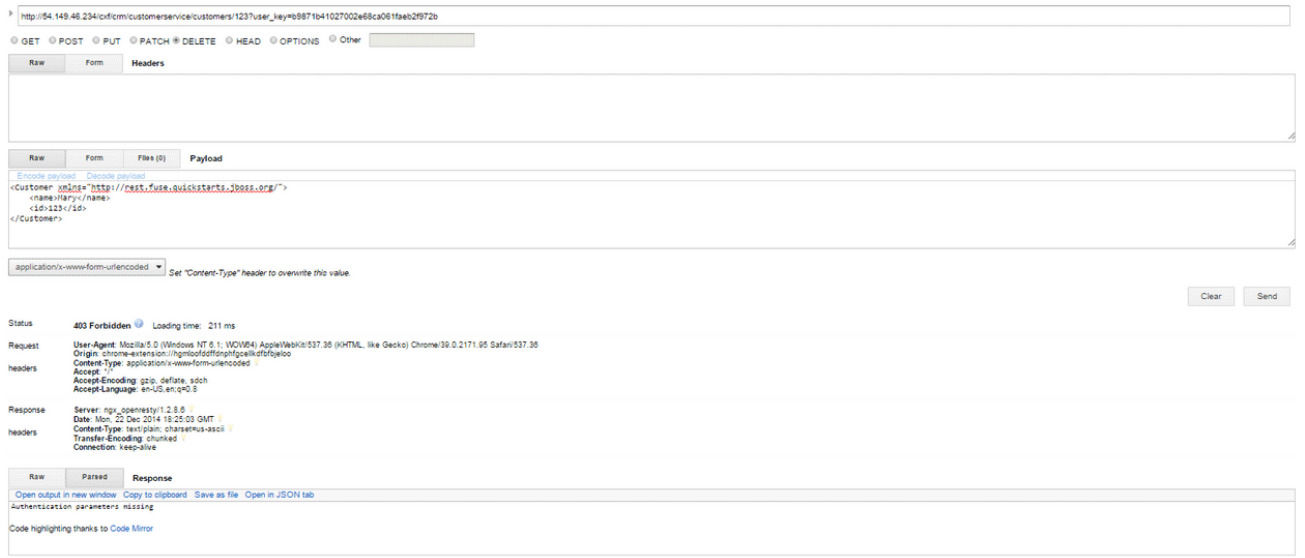

7.4.5. Step 5

Check the API Management analytics of your API.

If you now log back in to your 3scale account and go to Monitoring > Usage, you can see the various hits of the API endpoints represented as graphs.

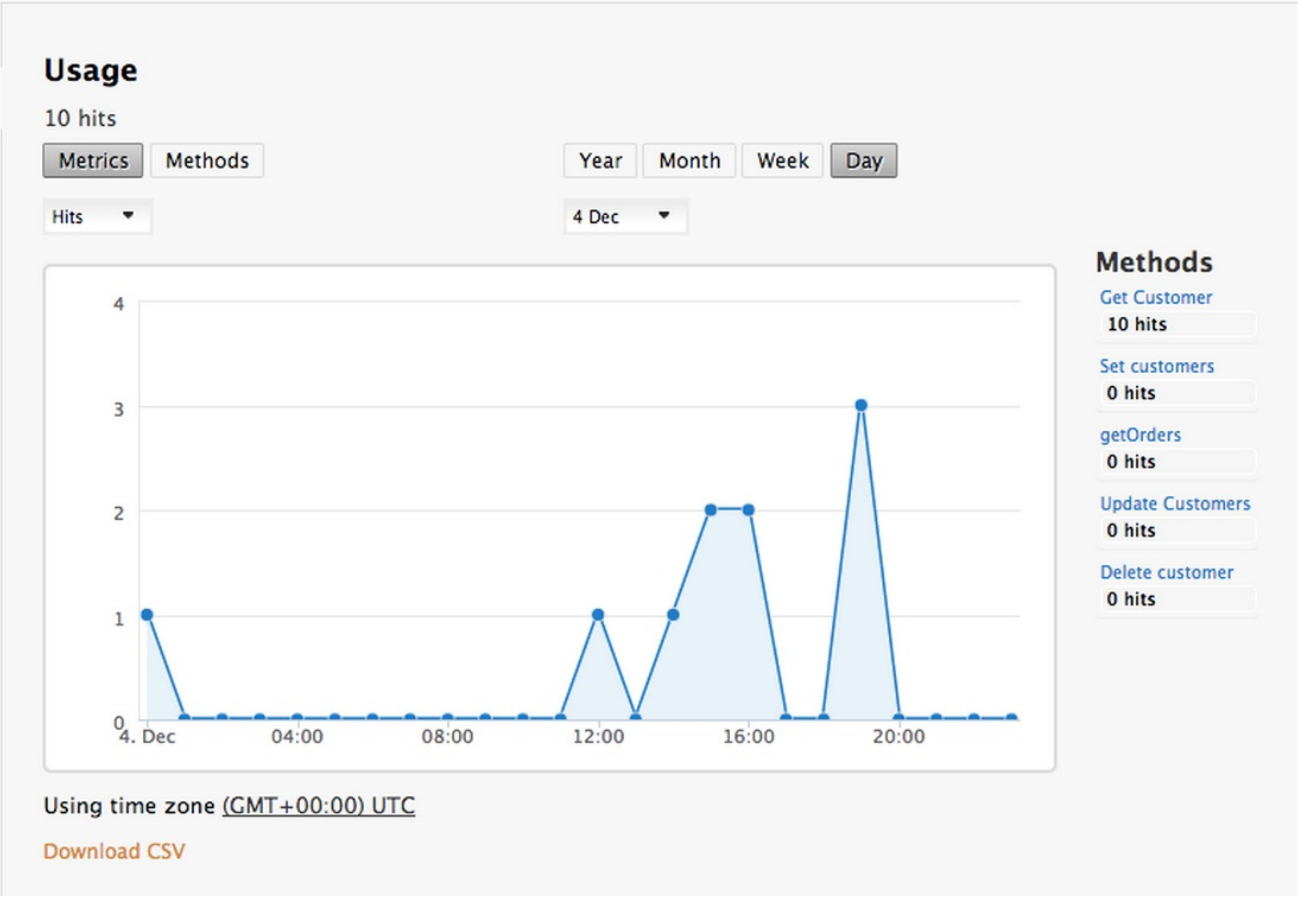

This is just one element of API Management that brings you full visibility and control over your API. Other features include:

1. Access control
2. Usage policies and rate limits
3. Reporting
4. API documentation and developer portals
5. Monetization and billing

For more details about the specific API Management features and their benefits, please refer to the 3scale API Management Platform product description.

For more details about the specific Red Hat JBoss Fuse product features and their benefits, please refer to the JBOSS FUSE Overview.

For more details about running Red Hat JBoss Fuse on OpenShift, please refer to the Getting Started with JBoss Fuse on OpenShift.