



# Federator.ai Release v4.4.1

## Installation Guide

## Contents

<b>Overview .....</b>	<b>3</b>
Federator.ai.....	3
Supported Metrics Data Sources.....	3
Prometheus (Kubernetes, Rancher, RedHat OpenShift) .....	3
Datadog.....	4
Sysdig .....	4
<b>Requirements and Recommended Resource Configuration .....</b>	<b>6</b>
Platforms .....	6
Federator.ai Resource Requirements.....	6
Federator.ai Version.....	6
Datadog Agent Version(reference) .....	6
Prometheus Version.....	6
Sysdig Agent Version(reference) .....	6
Persistent Volumes.....	7
Kafka.....	7
<b>Federator.ai Installation and Configuration.....</b>	<b>8</b>
Summary of Installation Steps.....	8
Pre-installation Check List .....	8
Before You Start (Datadog) .....	11
Before You Start (Sysdig) .....	12
New Installation .....	13
Upgrade from previous version.....	18
Prerequisite.....	18
Upgrade.....	18
Installing Datadog Watermark Pod Autoscaler (WPA) .....	21
<b>Appendix .....</b>	<b>25</b>
Datadog Dashboards Overview .....	25
ProphetStor Federator.ai Cluster Overview .....	25
ProphetStor Federator.ai Application Overview .....	25
ProphetStor Federator.ai Kafka Overview.....	26
ProphetStor Federator.ai Cost Analysis Overview .....	26
Sysdig Dashboard Overview .....	27
Federator.ai Cluster Overview .....	27
Federator.ai Application Overview .....	27
Federator.ai Application Overview .....	28

Federator.ai installation/uninstallation using Ansible .....	28
Prerequisite.....	28
Preparation (Ansible Control Node): .....	29
Installing Federator.ai.....	29
Troubleshooting .....	31
Downgrade from v4.4.1.....	31
Datadog Integration.....	37

## Overview

### Federator.ai

ProphetStor Federator.ai is an AI-based solution that helps enterprise clients to manage, optimize, auto-scale resources for any applications on Kubernetes/OpenShift. Using advanced machine learning algorithms to predict applications workload, Federator.ai scales the right amount of resources at the right time for optimized application performance.

- AI-based workload prediction for Kafka or other applications
- Resource recommendation based on workload prediction, application, Kubernetes, and other related metrics
- Automatic scaling of application containers through Datadog Watermark Pod Autoscaler (WPA), Federator.ai HPA, Native HPA with Datadog Cluster Agent.


### Supported Metrics Data Sources

There are three different types of metrics data sources supported in release 4.4.1: Prometheus, Datadog, and Sysdig.

#### Prometheus (Kubernetes, Rancher, RedHat OpenShift)

Prometheus is a free and open-source event monitoring tool for containers or microservices. It uses the principle of scraping to collect numerical data based on time series. Metrics are collected in regular timestamps and stored locally. Federator.ai supports using Prometheus gathering Kubernetes cluster metrics, and leverage collected data for workload predictions, recommendations for resource planning, autoscaling containers/pods, and cost analysis for clusters deployed in a multicloud environment.


The following diagram shows how the metrics are collected from Prometheus by Federator.ai in a Kubernetes environment.



## Datadog


Fedederator.ai has integrated with Datadog and utilizes the metrics collected by Datadog Agent for workload predictions. The following diagram shows how application metrics are used by Federator.ai to predict workload and to automatically scale applications for better performance. Specifically,

- Datadog Agent sends cluster/applications metrics to Datadog Services
- Federator.ai's Data-adapter queries cluster/applications metrics from Datadog Services and forwards to Federator.ai AI engine
- Data-adapter posts the prediction/recommendation/plan created by Federator.ai to Datadog Services
- Datadog Cluster Agent gets prediction/recommendation/plan from Datadog Services
- WPA applies plans and auto-scales applications
- Datadog Dashboard displays cluster/applications metrics and prediction/recommendation/plan by Federator.ai



## Sysdig

Fedederator.ai has integrated with Sysdig and utilizes the metrics collected by Sysdig Agent for workload predictions. The following diagram shows how application metrics are used by Federator.ai to predict workload and to automatically scale applications for better performance.



# Requirements and Recommended Resource Configuration

## Platforms

- OpenShift : 3.11, 4.x
- Kubernetes : 1.11 ~ 1.19.x
- Rancher v2.4.8

## Federator.ai Resource Requirements

- Total Resource Requirements
  - 5 CPU cores
  - 4.9 GB Memory
  - StorageClass: 166GB (require ReadWriteOnce access mode)
- Resource requirements for AI Engine
  - There must be at least one worker node with at least 2 CPU cores and 1 GB of memory available
  - The 2 CPU cores and 1 GB memory are included in the total 5 CPU cores and 4.9 GB memory requirements

## Federator.ai Version

- Version: Release v4.4.1-ga

## Datadog Agent Version(reference)

- Datadog Agent helm chart version: v2.4.24
- Datadog Agent version: v7.21.1
- Datadog Cluster Agent version: v1.8.0
- Datadog Watermark Pod Autoscaler version: v0.1.0

## Prometheus Version

- OpenShift
  - Default installed Prometheus
- Kubernetes
  - prometheus-operator-8.5.11
  - Rancher v2.4.8 kube-prometheus-stack-12.3.0
  - kube-prometheus-release-0.6
  - kube-prometheus-stack-12.5.0

## Sysdig Agent Version(reference)

- Sysdig agent: 10.8.0

## Persistent Volumes

- The StorageClass that provides the persistent volumes must support RWO (ReadWriteOnce) access mode.
- It is recommended to use persistent volumes instead of using ephemeral storage to store the data in the production environment.

## Kafka

- For Federator.ai's application-aware Kafka consumer resource/performance optimization feature, the following version of Kafka is supported :

Kafka operator version(Reference) : Strimzi/kafka:0.17.0-kafka-2.4.0

# Federator.ai Installation and Configuration

## Summary of Installation Steps

Step 0: Review pre-installation checklist items, make sure the environment and required information are ready.

### Step 1: Preparation

- For Datadog, obtain API Key, Application Key of Datadog Cloud Service account. Instructions are provided below.
- For Prometheus, obtain Prometheus service URL ( ex : [http://<prometheus\\_svc\\_name>.<namesapce>:9090](http://<prometheus_svc_name>.<namesapce>:9090))
- For Sysdig, obtain Sysdig API URL and Token.

### Step 2:

- For Datadog, install and configure Datadog Agent/Cluster Agent if they have not been installed. Please follow the Datadog documentation on how to install Datadog Agent and Cluster Agent.
- For Sysdig, install and configure Sysdig Agent. Please follow Sysdig documentation on how to install Sysdig Agent.

### Step 3: Install Federator.ai.

Step 4: Configure Federator.ai Data Adapter for the external metrics data source via Federator.ai Initial Setup Wizard.

Step 5: Optionally install Datadog WPA and apply WPA autoscaling CR if using Datadog WPA for autoscaling.

Step 6: Review installation result on Datadog/Sysdig Cloud Dashboard.

## Pre-installation Check List

Kubernetes:

#	Checklist Item	Requirement	Details
1	What is the Kubernetes version?	1.11~1.19.x	Use the command below to get the Kubernetes version:  \$ <b>kubectl version</b> ... Server Version: version.Info{Major:"1", Minor:"17", GitVersion:"v1.17.2", GitCommit:"59603c6e503c87169aea6106f57b9f242f64df89", GitTreeState:"clean", BuildDate:"2020-01-18T23:22:30Z", GoVersion:"go1.13.5", Compiler:"gc", Platform:"linux/amd64"}
2	Does installation on this Kubernetes cluster require a private image repository?	If a private image repository is required, the following information is needed during installation <ul style="list-style-type: none"><li>- Private image repository URL</li><li>- Credential of the private image repository</li></ul>	Input the URL and credential when the Federator.ai installation script asks for the information.

3	StorageClass and Persistent Volumes requirement	StorageClass supports ReadWriteOnce access mode. Available storage size is larger than 166GB.	The minimum storage size for Federator.ai Release v4.4.1 is 166GB, including database, data, and logs.
4	Kubernetes cluster CPU/memory requirement	Minimum CPU/mem/storage: - CPU: 5,000 (mcores) - Memory: 4.9 (GB) - Storage Class Capacity: 166GB  At least one worker node with - CPU: 2 Cores - Memory: 1GB	To be able to run the AI Engine pod, there must be at least one worker node that has more than 2 CPU cores and 1 GB of memory available.  2 CPU Cores and 1GB for AI Engine are included in the total 5 CPU Cores and 4.9 GB memory requirements.
5	Is this Kubernetes cluster allowed for NodePort configuration?	Federator.ai creates two NodePorts for GUI and REST API by default - REST API - https://<server>:31011 - GUI - https://<server>:31012	If NodePort is not allowed, answer 'N' when the installation script prompts for creating NodePorts. Users need to expose Federator.ai GUI and REST API service manually.
6	Will there be a resource quota imposed for the namespace where Federator.ai is installed?	CPU/mem request quota should be more than the minimum resource requirement - CPU: 5 Cores - Memory: 4.9 GB	The CPU/memory required for Federator.ai depends on the number of clusters and applications being monitored/managed.  Suggestion for initial namespace quota is - CPU 8 cores - Memory 12G  The quota could be adjusted if the number of managed clusters/applications increases.  Use the command to get namespace resource quota <b>\$ kubectl get resourcequota --all-namespaces</b>
7	Does this deployment requires resource request/limit specified?	By default, Federator.ai deployments do not specify resource requests/limits. It can be done by setting up an environment variable before installation starts.	To turn on resource request/limit settings for all Federator.ai deployments, manually export environment variable before running 'federatorai-launcher.sh'  <b>\$ export ENABLE_RESOURCE_REQUIREMENT=y</b> <b>\$ ./federatorai-launcher.sh</b>

## Prometheus:

#	Checklist Item	Requirement	Details
1	What is the Prometheus version? (for Kubernetes)	Recommended version-Prometheus operator helm chart version: 8.5.11- Prometheus operator version: 0.34.0 -Prometheus server version: 2.13.1	Use the command below to get Prometheus version: <b>\$ helm ls -A  grep -i prometheusprometheus-adapter monitoring</b> 1 2020-03-13 15:35:05.28963154 +0800 CST deployed prometheus-adapter-2.1.3 v0.6.0 prometheus-operator monitoring 1 2020-03-13 14:34:16.132479221 +0800 CST deployed prometheus-operator-8.12.1 0.37.0 ~# kubectl get deployment -A -o custom-columns=IMAGE:.spec.template.spec.containers[0].image  grep -i prometheusdirectxman12/k8s-prometheus-adapter-amd64:v0.6.0 quay.io/coreos/prometheus-operator:v0.37.0

## DataDog Agent:

#	Checklist Item	Requirement	Details
1	Is DataDog Agent installed?	DataDog Agent is mandatory	Kubernetes resources and workload metrics are collected by DataDog Agent.
2	Is DataDog Cluster Agent installed?	Cluster Agent is mandatory for the HPA autoscaling feature	Cluster Agent provides metrics to HPA Autoscaler for autoscaling.
3	Is DataDog WPA controller installed?	DataDog WPA is required if auto-scaling is done by WPA	DataDog WPA is the HPA Autoscaler developed by DataDog. Users can use DataDog WPA or Kubernetes native HPA to do autoscaling.
4	DataDog Kafka Consumer integration is enabled?	DataDog Kafka Consumer integration is mandatory if user wants to use Kafka optimization feature	Use the command to confirm Kafka integration is enabled <code>\$ kubectl exec &lt;datadog-agent-pod&gt; -n &lt;datadog-agent-namespace&gt; -- agent integration show datadog-kafka-consumer</code>  Refer to <a href="https://www.datadoghq.com/blog/monitor-kafka-with-datadog/">https://www.datadoghq.com/blog/monitor-kafka-with-datadog/</a> for Kafka Consumer integration installation
5	DataDog account API key	An API key is mandatory for connecting DataDog Service	Follow the steps described in the “Before You Start” session to obtain the API key.
6	DataDog account Application key	An application key is mandatory for connecting DataDog Service	Follow the steps described in the “Before You Start” session to obtain the Application key.
7	Is one of cluster name is configured for the DataDog agent/cluster agent? 1.>DD_TAGS with value ="kube_cluster:<cluster_name>" in values.yaml or 2.>"cluster_name" in values.yaml, or 3.>"DD_CLUSTER_NAME" in DataDog cluster agent deployment	"kube_cluster","cluster_name","kub e_cluster_name(DD_CLSUTER_NAM E)" one of them is required for Federator.ai to identify Kubernetes cluster.	Case 1.>New DataDog Agent installation: Install DataDog agent and cluster agent by "helm install -f values.yaml", in values.yaml. ... <b>clusterName: &lt;cluster-name&gt;</b> ... <b>clusterAgent:</b> enabled: <b>false</b> true Case 2.> In DataDog Agent installed environment, with no Cluster Agent and no cluster_name setting Update DataDog Agent to enable Cluster agent by "helm upgrade -f values.yaml ", in values.yaml - assign a cluster name ... <b>datadog:</b> clusterName: <cluster-name> - enable cluster agent ... <b>clusterAgent:</b> enabled: <b>false</b> true ... - \$helm upgrade ... - Check "DD_Cluster_Name" \$kubectl get daemonset <b>&lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml</b>

		<ul style="list-style-type: none"> <li>- name: DD_CLUSTER_NAME value: &lt;cluster-name&gt;</li> </ul> <p>3.&gt;In Datadog Agent and Cluster Agent installed environment, with no cluster_name setting Update Datadog Agent by “helm upgrade”</p> <ul style="list-style-type: none"> <li>- assign a cluster name datadog: clusterName: &lt;cluster-name&gt;</li> <li>- \$helm upgrade ...</li> <li>- Check “DD_Cluster_Name” <b>\$kubectl get daemonset</b> <b>&lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml</b></li> </ul> <p>As:</p> <ul style="list-style-type: none"> <li>- name: DD_CLUSTER_NAME value: &lt;cluster-name&gt;</li> </ul> <p>4.&gt; In Datadog Agent and Cluster Agent installed environment, with cluster_name setting Use the command below to confirm DD_Cluster_Name</p> <ul style="list-style-type: none"> <li>- \$kubectl get daemonset <b>&lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml</b></li> </ul> <p>As:</p> <ul style="list-style-type: none"> <li>- name: DD_CLUSTER_NAME value: &lt;cluster-name&gt;</li> </ul>
--	--	--

## Before You Start (Datadog)

- The admin role for installing Fedeator.ai is "Cluster Admin."
- Datadog agent must be ready if Federator.ai runs in the same Kubernetes cluster that is being monitored.
- Obtain Datadog account APIKey, APPKey.
  1. A Datadog account is required for connecting and using Datadog Cloud Service. If you don't have an account, visit Datadog website and sign up for a free trial account.  
<https://www.datadoghq.com/>
  2. Log in Datadog Cloud Service with your account and get an API key and Application key for using Datadog API

[https://docs.datadoghq.com/account\\_management/api-app-keys/](https://docs.datadoghq.com/account_management/api-app-keys/)

The screenshot shows the Datadog dashboard with a sidebar on the left containing various monitoring and logging tools like Watchdog, Events, Dashboards, Infrastructure, Monitors, Metrics, Integrations, APM, Notebooks, Logs, Security, and UX Monitoring. The main area displays a 'Welcome' message, host reporting statistics, and a hexagonal map of infrastructure hosts. A red box highlights the 'Integrations' section under Metrics, which is expanded to show 'APIs' (selected), 'Agent', and 'Embeds'. Below this, a message states 'You have 1 integration already installed.'

## Copy the API Key and Application Key for Federator.ai Data-Adapter configuration

The screenshot shows the 'API Keys' and 'Application Keys' sections in the Datadog API Keys management interface. The 'API Keys' section (marked with a red box and number 1) lists one key: '8a94db9a5e...'. The 'Application Keys' section (marked with a red box and number 2) lists one key: 'pod'. Both sections include columns for Name, Key, Created by, and Created at (UTC). A 'Create API Key' or 'Create Application Key' button is visible at the bottom of each section.

## Before You Start (Sysdig)

- Sysdig US east side, the Sysdig API URL is <https://app.sysdigcloud.com/>
  - For east US, Sysdig API URL is <https://app.sysdigcloud.com>
  - For west US, Sysdig API URL is <https://us2.app.sysdig.com>
  - For eu Central, Sysdig API URL is <https://eu1.app.sysdig.com>
- use Sysdig Monitor API Token as data source API key.

The screenshot shows the Sysdig Cloud Settings page at <https://app.sysdigcloud.com/#/settings/user>. The left sidebar has 'MONITOR' selected. The main area shows the 'User Profile' tab. It displays account information: Account Name (stella.chen@prophetstor.com), Role (Admin), and Current Team (Monitor Operations). Below this is the 'Admin Settings' section with a 'Hide Agent Install' toggle switch (disabled). The 'Sysdig Monitor API' section contains a note about tokens being team-specific, a 'Reset Token' link, and a red-bordered 'Sysdig Monitor API Token' field. A 'COPY' button and a 'RESET TOKEN' button are also present.

## New Installation

(Upgrade from previous Federator.ai, [please refer to the next section](#))

1. Log into Kubernetes cluster
2. Install the Federator.ai for Kubernetes by using the following command

```
$ curl https://raw.githubusercontent.com/containers-
ai/prophetstor/master/deploy/federatorai-launcher.sh | bash
```

```
~# curl https://raw.githubusercontent.com/containers-
ai/prophetstor/master/deploy/federatorai-launcher.sh | bash
      % Total    % Received % Xferd  Average Speed   Time     Time   Current
          Dload  Upload Total Spent   Spent    Left Speed
100 17260  100 17260    0      0  20210       0 --:--:-- --:--:-- 20210
Please input Federator.ai version tag (e.g., v4.4.1): v4.4.1-ga

Downloading scripts ...

Downloading Federator.ai CR yaml ...

Downloading Federator.ai operator yaml ...
Done
Do you want to use a private repository URL? [default: n]:
Do you want to launch the Federator.ai installation script? [default: y]:

Executing install.sh ...
Checking environment version...
...Passed
Enter the namespace you want to install Federator.ai [default: federatorai]:
```

```
tag_number = v4.4.1-ga
install_namespace = federatorai
-----
Is the above information correct? [default: y]:
Downloading file 00-namespace.yaml ...
Done
Downloading file 01-serviceaccount.yaml ...
Done
Downloading file 02-alamedaservice.crd.yaml ...
Done
Downloading file 03-federatorai-operator.deployment.yaml ...
Done
Downloading file 04-clusterrole.yaml ...
Done
Downloading file 05-clusterrolebinding.yaml ...
Done
Downloading file 06-role.yaml ...
Done
Downloading file 07-rolebinding.yaml ...
Done
Downloading file 08-service.yaml ...
Done
Downloading file 09-secret.yaml ...
Done
Downloading file 10-mutatingwebhook.yaml ...
Done
Downloading file 11-validatingwebhook.yaml ...
Done

Applying Federator.ai operator yaml files...
Applying 00-namespace.yaml...
namespace/federatorai created
Applying 01-serviceaccount.yaml...
serviceaccount/federatorai-operator created
Applying 02-alamedaservice.crd.yaml...
customresourcedefinition.apirextensions.k8s.io/alamedaservices.federatorai.containers.ai
created
Applying 03-federatorai-operator.deployment.yaml...
deployment.apps/federatorai-operator created
Applying 04-clusterrole.yaml...
clusterrole.rbac.authorization.k8s.io/federatorai-operator created
clusterrole.rbac.authorization.k8s.io/alameda-gc created
Applying 05-clusterrolebinding.yaml...
clusterrolebinding.rbac.authorization.k8s.io/federatorai-operator created
Applying 06-role.yaml...
role.rbac.authorization.k8s.io/federatorai-operator created
Applying 07-rolebinding.yaml...
rolebinding.rbac.authorization.k8s.io/federatorai-operator created
Applying 08-service.yaml...
service/federatorai-operator created
Applying 09-secret.yaml...
secret/federatorai-operator-admission created
Applying 10-mutatingwebhook.yaml...
mutatingwebhookconfiguration.admissionregistration.k8s.io/federatorai-operator-
servicesmutation created
Applying 11-validatingwebhook.yaml...
validatingwebhookconfiguration.admissionregistration.k8s.io/federatorai-operator-
servicesvalidation created
```

```

Checking pods...
Waiting for pod federatorai-operator-646f8446d8-hmv9p in namespace federatorai to be ready.
phase: [Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod federatorai-operator-646f8446d8-hmv9p in namespace federatorai to be ready.
phase: [Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod federatorai-operator-646f8446d8-hmv9p in namespace federatorai to be ready.
phase: [Running]
Waiting for pods in namespace federatorai to be ready...

All pods under namespace(federatorai) are ready.

Install Federator.ai operator v4.4.1-ga successfully

Downloading Federator.ai CR sample files ...
Done

Downloading Federator.ai alamedascaler sample files ...
Done
=====
Which storage type you would like to use? ephemeral or persistent?
[default: persistent]:
Specify log storage size [e.g., 10 for 10GB, default: 10]: 5
Specify AI engine storage size [e.g., 10 for 10GB, default: 10]: 5
Specify InfluxDB storage size [e.g., 100 for 100GB, default: 100]: 50
Specify storage class name: managed-nfs-storage
Do you want to expose dashboard and REST API services for external access? [default: y]:


-----
install_namespace = federatorai
storage_type = persistent
log storage size = 5 GB
AI engine storage size = 5 GB
InfluxDB storage size = 50 GB
storage class name = managed-nfs-storage
expose service = y
-----
Is the above information correct [default: y]:
Processing...
Waiting for datahub(v4.4.1-ga) pod to appear ...

datahub pod is present.

Checking pods...
Waiting for pod alameda-ai-5cc75c94b7-gvttm in namespace federatorai to be ready. phase:
[Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-5cc75c94b7-gvttm in namespace federatorai to be ready. phase:
[Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-5cc75c94b7-gvttm in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-executor-5f8ffb574-5b417 in namespace federatorai to be ready.
phase: [Running]
Waiting for pods in namespace federatorai to be ready...

```

```

Waiting for pod alameda-executor-5f8ffb574-5b4l7 in namespace federatorai to be ready.
phase: [Running]
Waiting for pods in namespace federatorai to be ready...

All pods under namespace(federatorai) are ready.

The default alamedaorganization under namespace federatorai is ready.

=====
You can now access GUI through https://<YOUR IP>:31012
The default login credential is admin/admin

Also, you can start to apply alamedascaler CR for the target you would like to monitor.
Review the administration guide for further details.

=====

=====
You can now access Federatorai REST API through https://<YOUR IP>:31011
The default login credential is admin/admin
The REST API online document can be found in https://<YOUR
IP>:31011/apis/v1/swagger/index.html

=====

Install Federator.ai v4.4.1-ga successfully

Downloaded YAML files are located under /tmp/install-op

Downloaded files are located under /tmp/federatorai-scripts/v4.4.1-ga

```

### 3. Verify Federator.ai pods are running properly



NAME	READY	STATUS	RESTARTS	AGE
alameda-ai-5cc75c94b7-gvttm	1/1	Running	0	19m
alameda-ai-dispatcher-6cc9cf66b4-2gcv	1/1	Running	1	19m
alameda-analyzer-67f545df8-82flk	1/1	Running	0	19m
alameda-datahub-76cb7c5446-xwx5b	1/1	Running	0	19m
alameda-executor-5f8ffb574-5b4l7	1/1	Running	5	19m
alameda-influxdb-0	1/1	Running	0	19m
alameda-notifier-6b84d9677c-sxfh9	1/1	Running	4	19m
alameda-operator-5c5b4b67df-h4rlq	1/1	Running	4	19m
alameda-rabbitmq-764db78df-q9jmh	1/1	Running	0	19m
alameda-recommender-b997975f-gwb2z	1/1	Running	0	19m
fedemeter-api-c74b9f777-4nv95	1/1	Running	0	19m
fedemeter-influxdb-0	1/1	Running	0	19m
federatorai-agent-74d5df8689-w5clw	1/1	Running	0	19m
federatorai-dashboard-backend-7bb94c8c48-w9r7f	1/1	Running	0	19m
federatorai-dashboard-frontend-797b98b5dd-h8ghv	1/1	Running	0	19m
federatorai-data-adapter-574fc947c7-wklhv	1/1	Running	0	19m
federatorai-operator-646f8446d8-hmv9p	1/1	Running	0	22m
federatorai-rest-5b95f47f45-fqvvk	1/1	Running	1	19m

### 4. Log on Federator.ai GUI and finish installation through the Initial Setup Wizard. For more information on Initial Setup Wizard, please see [Federator.ai 4.4 User Guide](#).

[https://<master\\_node\\_IP>:31012](https://<master_node_IP>:31012)

Login ID: admin

Password: admin



## Upgrade from a previous version

Federator.ai v4.4.1 supports upgrade from previous versions. The Federator.ai installation script automatically detects previously installed Federator.ai. When the installation script prompts if a backup of the previous configuration is needed, just enter yes to save a copy of the configuration if a rollback to the previous version is needed.

### Prerequisite

1. Federator.ai version is 4.2 or later.
2. Federator.ai installed and running with Persistent Volume.

### Upgrade

1. Log into Kubernetes cluster
2. Install the Federator.ai for Kubernetes by using the following command

```
$ curl https://raw.githubusercontent.com/containers-ai/prophetstor/master/deploy/federatorai-launcher.sh | bash
```

```
~# curl https://raw.githubusercontent.com/containers-ai/prophetstor/master/deploy/federatorai-launcher.sh | bash
      % Total      % Received % Xferd  Average Speed   Time     Time      Current
                                     Dload  Upload   Total   Spent    Left  Speed
100 16783  100 16783    0      0  25155      0 --:--:-- --:--:-- --:--:-- 25161
Please input Federator.ai version tag (e.g., v4.4.1): v4.4.1-ga

Downloading scripts ...

Downloading Federator.ai CR yaml ...

Downloading Federator.ai operator yaml ...
Done
Do you want to use a private repository URL? [default: n]:
Do you want to launch the Federator.ai installation script? [default: y]: y

Executing install.sh ...
Checking environment version...
...Passed
Enter the namespace you want to install Federator.ai [default: federatorai]:
-----
tag_number = v4.4.1-ga
install_namespace = federatorai
-----
Is the above information correct? [default: y]: y
Do you want to backup your configuration before upgrading Federator.ai? [default: y]: y
Please input path for storing backup configuration: [default: /tmp/configuration_backup]
Backup configuration...
backup yaml saved to folder /tmp/configuration_backup/federatorai-backup-1610932387
Done.
Downloading file 00-namespace.yaml ...
Done
Downloading file 01-serviceaccount.yaml ...
```

```
Done
Downloading file 02-alamedaservice.crd.yaml ...
Done
Downloading file 03-federatorai-operator.deployment.yaml ...
Done
Downloading file 04-clusterrole.yaml ...
Done
Downloading file 05-clusterrolebinding.yaml ...
Done
Downloading file 06-role.yaml ...
Done
Downloading file 07-rolebinding.yaml ...
Done
Downloading file 08-service.yaml ...
Done
Downloading file 09-secret.yaml ...
Done
Downloading file 10-mutatingwebhook.yaml ...
Done
Downloading file 11-validatingwebhook.yaml ...
Done

Applying Federator.ai operator yaml files...
Applying 00-namespace.yaml...
namespace/federatorai created
Applying 01-serviceaccount.yaml...
serviceaccount/federatorai-operator created
Applying 02-alamedaservice.crd.yaml...
customresourcedefinition.apiextensions.k8s.io/alamedaservices.federatorai.containers.ai
created
Applying 03-federatorai-operator.deployment.yaml...
deployment.apps/federatorai-operator created
Applying 04-clusterrole.yaml...
clusterrole.rbac.authorization.k8s.io/federatorai-operator created
clusterrole.rbac.authorization.k8s.io/alameda-gc created
Applying 05-clusterrolebinding.yaml...
clusterrolebinding.rbac.authorization.k8s.io/federatorai-operator created
Applying 06-role.yaml...
role.rbac.authorization.k8s.io/federatorai-operator created
Applying 07-rolebinding.yaml...
rolebinding.rbac.authorization.k8s.io/federatorai-operator created
Applying 08-service.yaml...
service/federatorai-operator created
Applying 09-secret.yaml...
secret/federatorai-operator-admission created
Applying 10-mutatingwebhook.yaml...
mutatingwebhookconfiguration.admissionregistration.k8s.io/federatorai-operator-
servicesmutation created
Applying 11-validatingwebhook.yaml...
validatingwebhookconfiguration.admissionregistration.k8s.io/federatorai-operator-
servicesvalidation created

Checking pods...
Waiting for pod federatorai-operator-c689896fb-1pwc8 in namespace federatorai to be ready.
phase: [Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod federatorai-operator-c689896fb-1pwc8 in namespace federatorai to be ready.
phase: [Pending]
Waiting for pods in namespace federatorai to be ready...
```

```

Waiting for pod federatorai-operator-c689896fb-1pwc8 in namespace federatorai to be ready.
phase: [Pending]
Waiting for pods in namespace federatorai to be ready...

All pods under namespace(federatorai) are ready.

Install Federator.ai operator v4.4.1-ga successfully

Downloading Federator.ai CR sample files ...
Done

Downloading Federator.ai alamedascaler sample files ...
Done
=====
Which storage type you would like to use? ephemeral or persistent?
[default: persistent]:
Specify log storage size [e.g., 10 for 10GB, default: 10]:
Specify AI engine storage size [e.g., 10 for 10GB, default: 10]:
Specify InfluxDB storage size [e.g., 100 for 100GB, default: 100]:
Specify storage class name: managed-nfs-storage
Do you want to expose dashboard and REST API services for external access? [default: y]:


-----
install_namespace = federatorai
storage_type = persistent
log storage size = 10 GB
AI engine storage size = 10 GB
InfluxDB storage size = 100 GB
storage class name = managed-nfs-storage
expose service = y
-----
Is the above information correct [default: y]:
Processing...
Waiting for datahub (v4.4.1-ga) pod to appear ...
datahub pod is present.

Checking pods...
Waiting for pod alameda-ai-568d8b44f4-96gxj in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-568d8b44f4-96gxj in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...

All pods under namespace(federatorai) are ready.

The default alamedaorganization under namespace federatorai is ready.

=====
You can now access GUI through https://<YOUR IP>:31012
The default login credential is admin/admin

Also, you can start to apply alamedascaler CR for the target you would like to monitor.
Review the administration guide for further details.
=====

=====
You can now access Federatorai REST API through https://<YOUR IP>:31011

```

```

The default login credential is admin/admin
The REST API online document can be found in https://<YOUR
IP>:31011/apis/v1/swagger/index.html
=====
Install Federator.ai v4.4.1-ga successfully

Downloaded YAML files are located under /tmp/install-op

Downloaded files are located under /tmp/federatorai-scripts/v4.4.1-ga

```

### 3. Verify Federator.ai pods are running properly

NAME	READY	STATUS	RESTARTS	AGE
alameda-ai-5cc75c94b7-gvttm	1/1	Running	0	19m
alameda-ai-dispatcher-6cc9cf66b4-2gccc	1/1	Running	1	19m
alameda-analyzer-67f545df8-82flk	1/1	Running	0	19m
alameda-datahub-76cb7c5446-xwx5b	1/1	Running	0	19m
alameda-executor-5f8ffb574-5b4l7	1/1	Running	5	19m
alameda-influxdb-0	1/1	Running	0	19m
alameda-notifier-6b84d9677c-sxfh9	1/1	Running	4	19m
alameda-operator-5c5b4b67df-h4rlq	1/1	Running	4	19m
alameda-rabbitmq-764db78df-q9jmh	1/1	Running	0	19m
alameda-recommender-b997975f-gwb2z	1/1	Running	0	19m
fedemeter-api-c74b9f777-4nv95	1/1	Running	0	19m
fedemeter-influxdb-0	1/1	Running	0	19m
federatorai-agent-74d5df8689-w5clw	1/1	Running	0	19m
federatorai-dashboard-backend-7bb94c8c48-w9r7f	1/1	Running	0	19m
federatorai-dashboard-frontend-797b98b5dd-h8ghv	1/1	Running	0	19m
federatorai-data-adapter-574fc947c7-wklhv	1/1	Running	0	19m
federatorai-operator-646f8446d8-hmv9p	1/1	Running	0	22m
federatorai-rest-5b95f47f45-fqvvk	1/1	Running	1	19m

## Installing Datadog Watermark Pod Autoscaler (WPA)

If you wish to enable HPA autoscaling via Datadog WPA for your application, please follow the instructions below to install Datadog WPA.

- Download Datadog WPA package

```

~# wget https://github.com/DataDog/watermarkpodautoscaler/archive/master.zip
~# unzip master.zip

```

- Install Watermark Pod Autoscaler controller

*WPA Helm Chart package requires using ‘helm’ to install. If you don’t have ‘helm’ installed, use the following command to install.*

```

~# curl -L https://raw.githubusercontent.com/helm/helm/master/scripts/get-helm-3| bash

```

- Set up environment variables and then use ‘helm’ command to install WPA

```
$ DD_NAMESPACE="default"
$ DD_NAMEWPA="wpacontroller"
$ helm install $DD_NAMEWPA -n $DD_NAMESPACE ./chart/watermarkpodautoscaler
```

```
~# pwd
/root/datadog_wpa/watermarkpodautoscaler
~# DD_NAMESPACE="default"
~# DD_NAMEWPA="wpacontroller"
~# helm install $DD_NAMEWPA -n $DD_NAMESPACE ./chart/watermarkpodautoscaler
~# kubectl get pods -n default
NAME                                     READY   STATUS    RESTARTS   AGE
datadog-monitoring-6lckr                  2/2     Running   0          2d19h
datadog-monitoring-cluster-agent-7d79559979-cnjhj   1/1     Running   0          2d19h
datadog-monitoring-dwq7f                  2/2     Running   0          2d19h
datadog-monitoring-hlm8x                  2/2     Running   0          2d19h
datadog-monitoring-kube-state-metrics-765978777d-b5dnq   1/1     Running   0          6d3h
nfs-client-provisioner-7cd5f68cf7-cfqqb      1/1     Running   0          6d3h
wpacontroller-watermarkpodautoscaler-68484f8dd4-zxm22   1/1     Running   18         6d3h
```

- Download WPA pod autoscaler CR yaml file

```
~# wget
https://github.com/DataDog/watermarkpodautoscaler/blob/master/deploy/crds/datadoghq.com_watermarkpodautoscalers_cr.yaml
```

- Edit datadoghq.com\_watermarkpodautoscalers\_cr.yaml

Configure WPA to auto-scale Kafka consumer group and generic application (NGINX)

```
~# mv datadoghq.com_watermarkpodautoscalers_cr.yaml wpa.yaml
~# vi wpa.yaml
apiVersion: datadoghq.com/v1alpha1
kind: WatermarkPodAutoscaler
metadata:
  name: consumer
  namespace: myproject
spec:
  # Add fields here
  # algorithm must be average
  algorithm: average
  maxReplicas: 10
  minReplicas: 1
  tolerance: 0.01
  downscaleForbiddenWindowSeconds: 300
  upscaleForbiddenWindowSeconds: 15
  scaleUpLimitFactor: 90
  scaleDownLimitFactor: 90
  scaleTargetRef:
    kind: Deployment
    apiVersion: apps/v1
    name: consumer
  readinessDelay: 10
  metrics:
    # Resource or External type supported
    # Example usage of External type
    - type: External
      external:
```

```

# do not edit highWatermark, and lowWatermark
# highWatermark and lowWatermark must be 1
highWatermark: "1"
lowWatermark: "1"
metricName: federatorai.recommendation
metricSelector:
  matchLabels:
    resource: replicas
    kube_cluster: k8s-4-205 ← see below #notes-1 for more details
    kube_deployment: consumer
    kube_namespace: myproject

# Example usage of Resource type
# - type: Resource
#   resource:
#     highWatermark: "50"
#     lowWatermark: "10"
#     name: cpu
#     metricSelector:
#       matchLabels:
#         foo: bar
---
apiVersion: datadoghq.com/v1alpha1
kind: WatermarkPodAutoscaler
metadata:
  name: nginx-sample
  namespace: nginx-sample
spec:
  # Add fields here
  # algorithm must be average
  algorithm: average
  maxReplicas: 5
  minReplicas: 1
  tolerance: 0.01
  downscaleForbiddenWindowSeconds: 300
  upscaleForbiddenWindowSeconds: 15
  scaleUpLimitFactor: 90
  scaleDownLimitFactor: 90
  scaleTargetRef:
    kind: Deployment
    apiVersion: apps/v1
    name: nginx-sample
  readinessDelay: 10
  metrics:
    # Resource or External type supported
    # Example usage of External type
    - type: External
      external:
        # do not edit highWatermark, and lowWatermark
        # highWatermark and lowWatermark must be 1
        highWatermark: "1"
        lowWatermark: "1"
        metricName: federatorai.recommendation
        metricSelector:
          matchLabels:
            resource: replicas
            kube_cluster: k8s-4-205 ← see below #notes-1 for more details
            kube_deployment: nginx-sample

```

```
kube_namespace: nginx-sample
```

**#notes-1:** “`kube_cluster`” value must match with `DD_TAGS` (value=“`kube_cluster:<cluster_name>`”) configured in Datadog Agent (`datadog-values.yaml`)

- Deploy WPA and confirm the status


```
~# kubectl apply -f wpa.yaml
```

## Appendix


### Datadog Dashboards Overview

The following Custom Datadog Dashboards are available after Federator.ai is installed.


#### ProphetStor Federator.ai Cluster Overview




#### ProphetStor Federator.ai Application Overview



## ProphetStor Federator.ai Kafka Overview




## ProphetStor Federator.ai Cost Analysis Overview




# Sysdig Dashboard Overview

The following Custom Sysdig Dashboards are available after Federator.ai is installed.


## Federator.ai Cluster Overview



## Federator.ai Application Overview



## Federator.ai Application Overview



## Federator.ai installation/uninstallation using Ansible

Only support Federator.ai since v4.4.0 or later

### Prerequisite

#### Ansible Control Node

Software:	Version:	Query Command:
Ansible	2.10.2 or later	ansible --version
Ansible Collection - community.kubernetes	1.1.1 or later	ansible-galaxy collection list or ansible-galaxy collection install community.kubernetes -vvv
Python	3.7 or later	python3 --version
OpenShift python client (Required by community.kubernetes collection)	0.11.2 or later	pip3 list   grep openshift
kubeconfig file (Need to copy target cluster's kubeconfig file to the Ansible Control Node)		e.g. file is put on /root/.kube/config.135

## Preparation (Ansible Control Node):

1. Install Ansible  
[https://docs.ansible.com/ansible/latest/installation\\_guide/intro\\_installation.html](https://docs.ansible.com/ansible/latest/installation_guide/intro_installation.html)
2. Install collection “community.kubernetes”  
`root@node1740:~# ansible-galaxy collection install community.kubernetes`
3. Install python & pip  
<https://www.python.org/downloads/>
4. Install OpenShift python client if you are using OpenShift clusters  
`root@node1740:~# pip3 install openshift`
5. Download Ansible playbook for [Federator.ai](#)
6. Modify user\_variable.yaml file for customizing needed info.

## Installing Federator.ai

### Variables for in user\_variable.yml

Group	Variable Name	Sample value	Description	Mandatory
<a href="#">Federator.ai</a> env	federatorai_version	v4.4.1-ga	<a href="#">Federator.ai</a> version tag	Y
Storage for <a href="#">Federator.ai</a> pods	storage_type	ephemeral or persistent	Using ephemeral persistent volume type	Y
Storage info (Only be used when storage_type is persistent)	log_storage_size	10	Log size reserved for every pod. 10 means 10GB	N
Private repo	enable_private_repo	y	Using private repo to pull the <a href="#">Federator.ai</a> required docker images	N
Pod resource	enable_resource_requirement	y	Add pod resource requirement (limits & requests) for every <a href="#">Federator.ai</a> pod	N
Expose services (Only be used when openshift_env is “n”)	expose_dashboard_and_rest_api_services	y	Expose the dashboard and API services in the Kubernetes cluster.	Y
Cluster type	openshift_env	n	Input “y” if installed cluster is OpenShift cluster	Y
	installed_namespace	federatorai	namespace where <a href="#">Federator.ai</a> will be installed	N

	image_url_prefix	""	Input the private repo URL	N
	ai_engine_size	10	Storage size reserved for Alameda AI engine.	N
	influxdb_storage_size	100	Data size reserved for InfluxDB pod.	N
	storage_class_name	scname	To specifying storage class name for provisioning persistent volumes	Y

Steps:

1. Go to Ansible playbook folder

```
root@node1740:~# cd ansible_for_federatorai
```

2. Modify user\_variable.yaml (under uninstaller folder) file for customizing needed info.

3. Export K8S\_AUTH\_KUBECONFIG to specify kubeconfig file path for Ansible collection (community.kubernetes).

```
root@node1740:~/ansible_for_federatorai# export  
K8S_AUTH_KUBECONFIG=/root/.kube/config.135
```

4. Run Ansible playbook

```
root@node1740:~/ansible_for_federatorai# ansible-playbook federtorai_installation.yaml
```

## Uninstalling Federator.ai

For Uninstallation, please use the file under ansible\_for\_federatorai/uninstaller directory.

### Variables in user\_variable.yml.

Group	Variable Name	Sample value	Description	Mandatory
Storage for <a href="#">Federator.ai</a> pods	storage_type	ephemeral or persistent	Specify current <a href="#">Federator.ai</a> storage type (ephemeral or persistent)	Y
Preserve current persistent volume (Only be used when storage_type is persistent)	preserve_pv	Y	Specify whether to preserve <a href="#">Federator.ai</a> PVs	

Steps:

1. Get to Ansible playbook uninstallation folder

```
~# cd ansible_for_federatorai/uninstaller
```

2. Modify user\_variable.yaml (under uninstaller folder) file for customizing needed info.

3. Export **K8S\_AUTH\_KUBECONFIG** to specify kubeconfig file path for Ansible collection (community.kubernetes)

```
~# export K8S_AUTH_KUBECONFIG=/root/.kube/config.135
```

4. Run Ansible playbook

```
~# ansible-playbook federatorai_uninstaller.yaml
```

## Troubleshooting

### Downgrade from v4.4.1

- **V4.4.1 ->v4.3.1**

Following v4.3.1 installation step to rollback to previous version v4.3.1 directly. Installation script keeps configuration, metrics, and prediction data, which is stored on Persistent volume, user does not need to perform a complicated process.

- **V4.4.1-> v4.2**

If Federator.ai is upgraded from v4.2, rollback could be done from the 4.2 configuration backup saved during the upgrade process. Here is the general workflow for downgrading to the 4.2 version:

1. Run v4.4.1/Uninstall.sh script.
2. Re-install v4.2.
3. Restore 4.2 backup configuration.

**Step 1:** Run uninstall.sh. The uninstall script is placed under /tmp/federatorai-scripts/v4.4.1-ga/scripts directory.

```
~#cd /tmp/federatorai-scripts/v4.4.1-ga/scripts/
~#bash uninstall.sh
Do you want to preserve your Federator.ai persistent volumes? [default: y]:
Patching pv pvc-09324a63-01cc-44d1-9d67-313d2172b41e ...
persistentvolume/pvc-09324a63-01cc-44d1-9d67-313d2172b41e patched (no change)
Done.
Patching pv pvc-0a0b1fb2-b96b-4c74-abdf-5aa1ef930f4f ...
persistentvolume/pvc-0a0b1fb2-b96b-4c74-abdf-5aa1ef930f4f patched (no change)
Done.
Patching pv pvc-0da9f9c8-9ee0-4ac1-b5dc-50e6311d5920 ...
persistentvolume/pvc-0da9f9c8-9ee0-4ac1-b5dc-50e6311d5920 patched (no change)
Done.
Patching pv pvc-0f6554ab-d0d6-46f1-a295-b3cf133ceef6 ...
persistentvolume/pvc-0f6554ab-d0d6-46f1-a295-b3cf133ceef6 patched (no change)
Done.
```

```
Patching pv pvc-15eef793-2012-44a7-9b6b-067fbba999e0 ...
persistentvolume/pvc-15eef793-2012-44a7-9b6b-067fbba999e0 patched (no change)
Done.
Patching pv pvc-29e8d506-b659-4f78-b22e-b74a0baea80e ...
persistentvolume/pvc-29e8d506-b659-4f78-b22e-b74a0baea80e patched (no change)
Done.
Patching pv pvc-33cae9a9-8b6d-4786-806d-34ac3ca2a3d5 ...
persistentvolume/pvc-33cae9a9-8b6d-4786-806d-34ac3ca2a3d5 patched (no change)
Done.
Patching pv pvc-4531b2ae-6678-4342-b83f-03e757013523 ...
persistentvolume/pvc-4531b2ae-6678-4342-b83f-03e757013523 patched (no change)
Done.
Patching pv pvc-4ad88729-6c1b-4fb7-95b9-fbc30748c2b6 ...
persistentvolume/pvc-4ad88729-6c1b-4fb7-95b9-fbc30748c2b6 patched (no change)
Done.
Patching pv pvc-5452d9fd-e471-42a5-a03c-2435c7539972 ...
persistentvolume/pvc-5452d9fd-e471-42a5-a03c-2435c7539972 patched (no change)
Done.
Patching pv pvc-570ad717-a306-4800-b6a0-cbe02a1805e3 ...
persistentvolume/pvc-570ad717-a306-4800-b6a0-cbe02a1805e3 patched (no change)
Done.
Patching pv pvc-65bb40fe-0c98-4f5c-8af0-42558f0510f1 ...
persistentvolume/pvc-65bb40fe-0c98-4f5c-8af0-42558f0510f1 patched (no change)
Done.
Patching pv pvc-6a1257b4-7582-4ab9-be66-5f7d8e85badc ...
persistentvolume/pvc-6a1257b4-7582-4ab9-be66-5f7d8e85badc patched (no change)
Done.
Patching pv pvc-6feb2a5a-7b53-421a-85e4-25491688057a ...
persistentvolume/pvc-6feb2a5a-7b53-421a-85e4-25491688057a patched (no change)
Done.
Patching pv pvc-7412750a-fe39-4a79-a78b-b47fd6f18f68 ...
persistentvolume/pvc-7412750a-fe39-4a79-a78b-b47fd6f18f68 patched (no change)
Done.
Patching pv pvc-79dfbb73-cdc7-4ac0-a73e-94b1b973f60b ...
persistentvolume/pvc-79dfbb73-cdc7-4ac0-a73e-94b1b973f60b patched (no change)
Done.
Patching pv pvc-7fdb8acb-461a-4633-815a-2eea4b8d1148 ...
persistentvolume/pvc-7fdb8acb-461a-4633-815a-2eea4b8d1148 patched (no change)
Done.
Patching pv pvc-83f71f04-9516-44fa-a083-84732e9240ed ...
persistentvolume/pvc-83f71f04-9516-44fa-a083-84732e9240ed patched (no change)
Done.
Patching pv pvc-8de6d659-d003-4243-91c3-ca7526f33c2d ...
persistentvolume/pvc-8de6d659-d003-4243-91c3-ca7526f33c2d patched (no change)
Done.
Patching pv pvc-8fe030f1-24cd-4a9e-b7e1-4e7b50d76f65 ...
persistentvolume/pvc-8fe030f1-24cd-4a9e-b7e1-4e7b50d76f65 patched (no change)
Done.
Patching pv pvc-90bd467b-5730-4300-86c6-1ec65cba9b08 ...
persistentvolume/pvc-90bd467b-5730-4300-86c6-1ec65cba9b08 patched (no change)
Done.
Patching pv pvc-9a4ecf7e-8579-45d0-92cd-655aaaf0853f9 ...
persistentvolume/pvc-9a4ecf7e-8579-45d0-92cd-655aaaf0853f9 patched (no change)
```

```
Done.
Patching pv pvc-9acd0c3d-c299-44b4-bc59-b5a9eb856521 ...
persistentvolume/pvc-9acd0c3d-c299-44b4-bc59-b5a9eb856521 patched (no change)
Done.
Patching pv pvc-9b7c2a77-0bde-4748-9eda-ca067cd6c710 ...
persistentvolume/pvc-9b7c2a77-0bde-4748-9eda-ca067cd6c710 patched (no change)
Done.
Patching pv pvc-9e7429c9-30df-4790-b706-61a1b86cbe35 ...
persistentvolume/pvc-9e7429c9-30df-4790-b706-61a1b86cbe35 patched (no change)
Done.
Patching pv pvc-b10c40d3-6485-4ddb-828e-dec8693ca31e ...
persistentvolume/pvc-b10c40d3-6485-4ddb-828e-dec8693ca31e patched (no change)
Done.
Patching pv pvc-b3b35cad-1a5b-4f6b-93de-2026e4502112 ...
persistentvolume/pvc-b3b35cad-1a5b-4f6b-93de-2026e4502112 patched (no change)
Done.
Patching pv pvc-b517207b-54b6-4a42-81da-936acf0d30 ...
persistentvolume/pvc-b517207b-54b6-4a42-81da-936acf0d30 patched (no change)
Done.
Patching pv pvc-bc70b3d2-9e14-442a-930b-3f817f312b79 ...
persistentvolume/pvc-bc70b3d2-9e14-442a-930b-3f817f312b79 patched (no change)
Done.
Patching pv pvc-bd3cf813-ec79-4649-a685-1a8fac8f375c ...
persistentvolume/pvc-bd3cf813-ec79-4649-a685-1a8fac8f375c patched (no change)
Done.
Patching pv pvc-c4e1717a-bff7-4997-ab94-d3c6e13c05a3 ...
persistentvolume/pvc-c4e1717a-bff7-4997-ab94-d3c6e13c05a3 patched (no change)
Done.
Patching pv pvc-ced6151e-962a-4bd9-854c-82083ca292e8 ...
persistentvolume/pvc-ced6151e-962a-4bd9-854c-82083ca292e8 patched (no change)
Done.
Patching pv pvc-d09eee21-f5f4-4c01-8c28-c02a9f951b7d ...
persistentvolume/pvc-d09eee21-f5f4-4c01-8c28-c02a9f951b7d patched (no change)
Done.
Patching pv pvc-d88f0ae6-e645-4980-a477-b354f1182a8e ...
persistentvolume/pvc-d88f0ae6-e645-4980-a477-b354f1182a8e patched (no change)
Done.
Patching pv pvc-dccdcbf7-6f7f-46a9-9388-8a5b97e7126d ...
persistentvolume/pvc-dccdcbf7-6f7f-46a9-9388-8a5b97e7126d patched (no change)
Done.
Patching pv pvc-e7667f8f-8e7d-4a7b-9fae-4d5f9726a59d ...
persistentvolume/pvc-e7667f8f-8e7d-4a7b-9fae-4d5f9726a59d patched (no change)
Done.
Patching pv pvc-f4f884c6-066e-4ebf-90e4-426a132417cf ...
persistentvolume/pvc-f4f884c6-066e-4ebf-90e4-426a132417cf patched (no change)
Done.
Patching pv pvc-f8257bf4-abf9-4de2-b3da-1f2daa1451ad ...
persistentvolume/pvc-f8257bf4-abf9-4de2-b3da-1f2daa1451ad patched (no change)
Done.
Patching pv pvc-fb35cb57-436a-4561-80f3-2a3e0b763c8f ...
persistentvolume/pvc-fb35cb57-436a-4561-80f3-2a3e0b763c8f patched (no change)
Done.
Patching pv pvc-fe27c9cf-80d4-4acc-b50b-94bd09d575a4 ...
```

```
persistentvolume/pvc-fe27c9cf-80d4-4acc-b50b-94bd09d575a4 patched (no change)
Done.

-----
Starting to remove the Federator.ai product
-----

Please input your Federator.ai Operator tag: v4.4.1-ga

-----
Your tag number = v4.4.1-ga
-----
Is the above information correct? [default: y]:
Downloading file 00-namespace.yaml ...
Downloading file 01-serviceaccount.yaml ...
Downloading file 02-alamedaservice.crd.yaml ...
Downloading file 03-federatorai-operator.deployment.yaml ...
Downloading file 04-clusterrole.yaml ...
Downloading file 05-clusterrolebinding.yaml ...
Downloading file 06-role.yaml ...
Downloading file 07-rolebinding.yaml ...
Downloading file 08-service.yaml ...
Downloading file 09-secret.yaml ...
Downloading file 10-mutatingwebhook.yaml ...
Downloading file 11-validatingwebhook.yaml ...

Deleting my-alamedaservice in federatorai namespace...
clusterrole.rbac.authorization.k8s.io "alameda-gc" deleted

Deleting 11-validatingwebhook.yaml ...
validatingwebhookconfiguration.admissionregistration.k8s.io "federatorai-
operator-servicesvalidation" deleted

Deleting 10-mutatingwebhook.yaml ...
mutatingwebhookconfiguration.admissionregistration.k8s.io "federatorai-operator-
servicesmutation" deleted

Deleting 09-secret.yaml ...
secret "federatorai-operator-admission" deleted

Deleting 08-service.yaml ...
service "federatorai-operator" deleted

Deleting 07-rolebinding.yaml ...
rolebinding.rbac.authorization.k8s.io "federatorai-operator" deleted

Deleting 06-role.yaml ...
role.rbac.authorization.k8s.io "federatorai-operator" deleted

Deleting 05-clusterrolebinding.yaml ...
clusterrolebinding.rbac.authorization.k8s.io "federatorai-operator" deleted

Deleting 04-clusterrole.yaml ...
```

```
clusterrole.rbac.authorization.k8s.io "federatorai-operator" deleted
Error from server (NotFound): error when deleting "04-clusterrole.yaml":
clusterroles.rbac.authorization.k8s.io "alameda-gc" not found
Error in removing 04-clusterrole.yaml

Deleting 03-federatorai-operator.deployment.yaml ...
deployment.apps "federatorai-operator" deleted

Deleting 02-alamedaservice.crd.yaml ...
customresourcedefinition.apiextensions.k8s.io
"alamedaservices.federatorai.containers.ai" deleted

Deleting 01-serviceaccount.yaml ...
serviceaccount "federatorai-operator" deleted

Deleting 00-namespace.yaml ...
namespace "federatorai" deleted

Namespace federatorai is removed successfully.
```

## Step 2: Reinstall Federator.ai 4.2.

## Step 3: Restore 4.2 backup configuration.

Follow the steps below:

- a. Go to /tmp/configuration\_backup, which is the default federator.ai configuration backup directory.
- b. Change to the directory where the 4.2 configuration backup is stored.
- c. Run backup-restore.sh script.

```
~# cd /tmp/configuration_backup
~# cd federatorai-backup-1611212333
~# bash backup-restore.sh -r
Download origin operator upstream files and apply
v4.3.1046
Downloading file 00-namespace.yaml ...
Downloading file 01-serviceaccount.yaml ...
Downloading file 02-alamedaservice.crd.yaml ...
Downloading file 03-federatorai-operator.deployment.yaml ...
Downloading file 04-clusterrole.yaml ...
Downloading file 05-clusterrolebinding.yaml ...
Downloading file 06-role.yaml ...
Downloading file 07-rolebinding.yaml ...
/tmp/configuration_backup/federatorai-backup-1611212333
namespace/federatorai created
serviceaccount/federatorai-operator created
customresourcedefinition.apiextensions.k8s.io/alamedaservices.federatorai.
containers.ai created
deployment.apps/federatorai-operator created
clusterrole.rbac.authorization.k8s.io/federatorai-operator created
clusterrole.rbac.authorization.k8s.io/alameda-gc created
```

```
clusterrolebinding.rbac.authorization.k8s.io/federatorai-operator created
role.rbac.authorization.k8s.io/federatorai-operator created
rolebinding.rbac.authorization.k8s.io/federatorai-operator created
Restore service
alamedaservice.federatorai.containers.ai/my-alamedaservice created
Patch pv if necessary
persistentvolume/pvc-09324a63-01cc-44d1-9d67-313d2172b41e patched
persistentvolume/pvc-0a0b1fb2-b96b-4c74-abdf-5aa1ef930f4f patched
persistentvolume/pvc-0da9f9c8-9ee0-4ac1-b5dc-50e6311d5920 patched
persistentvolume/pvc-0f6554ab-d0d6-46f1-a295-b3cf133ceef6 patched
persistentvolume/pvc-15eef793-2012-44a7-9b6b-067fbba999e0 patched
persistentvolume/pvc-29e8d506-b659-4f78-b22e-b74a0baea80e patched
persistentvolume/pvc-33cae9a9-8b6d-4786-806d-34ac3ca2a3d5 patched
persistentvolume/pvc-4531b2ae-6678-4342-b83f-03e757013523 patched
persistentvolume/pvc-4ad88729-6c1b-4fb7-95b9-fbc30748c2b6 patched
persistentvolume/pvc-5452d9fd-e471-42a5-a03c-2435c7539972 patched
persistentvolume/pvc-570ad717-a306-4800-b6a0-cbe02a1805e3 patched
persistentvolume/pvc-65bb40fe-0c98-4f5c-8af0-42558f0510f1 patched
persistentvolume/pvc-6a1257b4-7582-4ab9-be66-5f7d8e85badc patched
persistentvolume/pvc-6feb2a5a-7b53-421a-85e4-25491688057a patched
persistentvolume/pvc-7412750a-fe39-4a79-a78b-b47fd6f18f68 patched
persistentvolume/pvc-79dfbb73-cdc7-4ac0-a73e-94b1b973f60b patched
persistentvolume/pvc-7fdb8acb-461a-4633-815a-2eea4b8d1148 patched
persistentvolume/pvc-83f71f04-9516-44fa-a083-84732e9240ed patched
persistentvolume/pvc-8de6d659-d003-4243-91c3-ca7526f33c2d patched
persistentvolume/pvc-8fe030f1-24cd-4a9e-b7e1-4e7b50d76f65 patched
persistentvolume/pvc-90bd467b-5730-4300-86c6-1ec65cba9b08 patched
persistentvolume/pvc-9a4ecf7e-8579-45d0-92cd-655aaf0853f9 patched
persistentvolume/pvc-9acd0c3d-c299-44b4-bc59-b5a9eb856521 patched
persistentvolume/pvc-9b7c2a77-0bde-4748-9eda-ca067cd6c710 patched
persistentvolume/pvc-9e7429c9-30df-4790-b706-61a1b86cbe35 patched
persistentvolume/pvc-b10c40d3-6485-4ddb-828e-dec8693ca31e patched
persistentvolume/pvc-b3b35cad-1a5b-4f6b-93de-2026e4502112 patched
persistentvolume/pvc-b517207b-54b6-4a42-81da-936acfff0d30 patched
persistentvolume/pvc-bc70b3d2-9e14-442a-930b-3f817f312b79 patched
persistentvolume/pvc-bd3cf813-ec79-4649-a685-1a8fac8f375c patched
persistentvolume/pvc-c4e1717a-bff7-4997-ab94-d3c6e13c05a3 patched
persistentvolume/pvc-ced6151e-962a-4bd9-854c-82083ca292e8 patched
persistentvolume/pvc-d09eee21-f5f4-4c01-8c28-c02a9f951b7d patched
persistentvolume/pvc-d88f0ae6-e645-4980-a477-b354f1182a8e patched
persistentvolume/pvc-dccdcfb7-6f7f-46a9-9388-8a5b97e7126d patched
persistentvolume/pvc-e7667f8f-8e7d-4a7b-9fae-4d5f9726a59d patched
persistentvolume/pvc-f4f884c6-066e-4ebf-90e4-426a132417cf patched
persistentvolume/pvc-f8257bf4-abf9-4de2-b3da-1f2daa1451ad patched
persistentvolume/pvc-fb35cb57-436a-4561-80f3-2a3e0b763c8f patched
persistentvolume/pvc-fe27c9cf-80d4-4acc-b50b-94bd09d575a4 patched
Restore CRs
alamedanotificationchannel.notifying.containers.ai/default created
alamedanotificationtopic.notifying.containers.ai/default created
alamedaorganization.tenant.containers.ai/default created
Restore complete
```

## Datadog Integration

### 1. Datadog WPA dumps errors during autoscaling

- Error message in WPA Controller

```
~# kubectl get pod -n default
NAME                                         READY   STATUS    RESTARTS
AGE
datadog-agent-2m6kk                           1/1     Running   2
2d
datadog-agent-8kd54                           1/1     Running   0
2d
datadog-agent-94rl6                           1/1     Running   0
2d
datadog-agent-mq4mv                           1/1     Running   0
2d
datadog-cluster-agent-74f44fdd4d-82tjp       1/1     Running   0
1d
docker-registry-1-vw59s                         1/1     Running   4
324d
prometheus-adapter-799b7dfc4f-rs7zj          1/1     Running   1
6d
registry-console-2-jxfdl                      1/1     Running   2
6d
router-1-sw781                               1/1     Running   4
324d
wpacontroller-watermarkpodautoscaler-7ffbb97f9d-hcbsg 1/1     Running   0
1d

~# kubectl logs wpacontroller-watermarkpodautoscaler-7ffbb97f9d-hcbsg -n
default
```

```
{"level": "info", "ts": 1589533961.5993037, "logger": "wpa_controller", "msg": "Successful rescale", "Request.Namespace": "myproject", "Request.Name": "consumer1-topic0001-group-0001", "currentReplicas": 40, "desiredReplicas": 40, "rescaleReason": ""}

{"level": "error", "ts": 1589533961.600972, "logger": "wpa_controller", "msg": "Error during reconcileWPA", "Request.Namespace": "myproject", "Request.Name": "consumer1-topic0001-group-0001", "error": "the server could not find the requested resource (put watermarkpodautoscalers.datadoghq.com consumer1-topic0001-group-0001)", "stacktrace": "github.com/go-logr/zapr.(*zapLogger).Error\n\twatermarkpodautoscaler/vendor/github.com/go-logr/zapr/zapr.go:128\n\tgithub.com/DataDog/watermarkpodautoscaler/pkg/controller/watermarkpodautoscaler.(*ReconcileWatermarkPodAutoscaler).Reconcile\n\t\twatermarkpodautoscaler/pkg/controller/watermarkpodautoscaler/watermarkpodautoscaler_controller.go:345\n\t\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).reconcileHandler\n\t\twatermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:216\n\t\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).processNextWorkItem\n\t\twatermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:192\n\t\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).worker\n\t\twatermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:171\n\t\tk8s.io/apimachinery/pkg/util/wait.JitterUntil\n\t\tfunc1\n\t\twatermarkpodautoscaler/vendor/k8s.io/apimachinery/pkg/util/wait/wait.go:152\n\t\tk8s.io/apimachinery/pkg/util/wait.JitterUntil\n\t\twatermarkpodautoscaler/vendor/k8s.io/apimachinery/pkg/util/wait/wait.go:153\n\t\tk8s.io/apimachinery/pkg/util/wait.Until\n\t\twatermarkpodautoscaler/vendor/k8s.io/apimachinery/pkg/util/wait/wait.go:88"}
```

- Reason
  - WPA is incompatible with Kubernetes 1.11
  - Install WPA on Kubernetes 1.11 dumps errors

must only have "properties", "required" or "description" at the root if the status subresource is enabled

- Workaround
  - Comment out 'subresources' key in WatermarkPodAutoscaler CRD

```
~# cd  
datadog_wpa/watermarkpodautoscaler_for_k8s_1.11/chart/watermarkpodautoscaler/templates  
~# vi datadoghq.com_watermarkpodautoscalers_crd.yaml  
...  
...  
  shortNames:  
    - wpa  
  singular: watermarkpodautoscaler  
  scope: Namespaced  
  #subresources: ← comment out  
  # status: {} ← comment out  
  validation:  
    openAPIV3Schema:  
      description: WatermarkPodAutoscaler is the Schema for the  
watermarkpodautoscalers  
      API  
      properties:  
        apiVersion:  
          description: 'APIVersion defines the versioned schema of this  
representation  
...  
...
```

*Note: It can auto-scale monitored application, but dump some errors during update status*

- Related Datadog WPA ticket
- <https://github.com/DataDog/watermarkpodautoscaler/issues/50>

## 2. Data Adapter reports errors

- Error messages in Data Adapter logs

```
~# oc exec -it $(oc get pods|grep federatorai-data-adapter|grep Running|awk '{print $1}') -- cat /var/log/telegraf.log  
> telegraf.log  
~# cat telegraf.log | grep "E!"  
2020-05-15T09:59:33Z E! [datadog][application_aware] Failed to get kafka consumer spec replicas  
2020-05-15T09:59:33Z E! [inputs.datadog_application_aware] Error in plugin:  
[url=https://api.datadoghq.com/api/v1/query][kafka]: Failed to get consumer information.
```

- Reason

Datadog Agent does not work with ‘kube-state-metrics’ comes with OpenShift

- Solution

Install another compatible ‘kube-state-metrics’

If there is another kube-state-metrics running on openshift, rename all the clusterrole and clusterrolebinding name of kube-state-metrics to prevent kube-state-metrics clusterrole name collision

restart datadog agent and make sure agent integrate with kube-state-metrics properly.  
check all the node agent status by following command

```
~# oc exec <datadog-agent-pod-name> agent status
```