



# Federator.ai Release v4.6.1

## 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
VMWare vCenter .....	5
AWS CloudWatch.....	5
<b>Requirements and Recommended Resource Configuration.....</b>	<b>7</b>
Platforms .....	7
Data Source .....	7
Federator.ai Resource Requirements.....	7
Federator.ai Version.....	7
Datadog Agent Version(reference) .....	7
Prometheus Version(reference).....	7
Sysdig Agent Version(reference) .....	8
Persistent Volumes.....	8
Kafka.....	8
NGINX as Ingress .....	8
<b>Federator.ai Installation and Configuration.....</b>	<b>9</b>
Summary of Installation Steps.....	9
Pre-installation Check List .....	9
Before You Start .....	12
Datadog.....	12
Sysdig .....	13
VMware vCenter .....	14
AWS CloudWatch.....	14
New Installation .....	16
Upgrade from a previous version.....	20
Prerequisite.....	20
Upgrade.....	20
Installing Datadog Watermark Pod Autoscaler (WPA) .....	23
<b>Appendix .....</b>	<b>26</b>
Datadog Dashboards Overview .....	26
ProphetStor Federator.ai Cluster Overview .....	26

ProphetStor Federator.ai Application Overview .....	26
ProphetStor Federator.ai Kafka Overview.....	27
ProphetStor Federator.ai Cost Analysis Overview .....	27
Sysdig Dashboard Overview .....	28
Federator.ai Cluster Overview .....	28
Federator.ai Application Overview.....	28
Federator.ai Application Overview .....	29
Federator.ai installation/uninstallation using Helm Chart .....	29
Federator.ai installation/uninstallation using Ansible .....	31
Prerequisite.....	31
Preparation (Ansible Control Node): .....	32
Installing Federator.ai .....	32
Troubleshooting .....	34
Downgrade from v4.6.1.....	34
Datadog Integration.....	40

## Overview

### Federator.ai

ProphetStor Federator.ai is an AI-based solution that helps enterprises manage and optimize resources for applications on Kubernetes and virtual machines (VMs) in VMware or AWS EC2 clusters. Using advanced machine learning algorithms to predict application workloads, Federator.ai offers:

- AI-based workload prediction for containerized applications in Kubernetes clusters and VMs in VMware or AWS EC2 clusters
- Resource recommendations based on workload prediction, application, Kubernetes, and other related metrics
- Automatic scaling of application containers
- Multicloud cost analysis and recommendations based on workload predictions for Kubernetes clusters and VM clusters
- Actual cost and potential savings based on recommendations for clusters, Kubernetes applications, VMs, and Kubernetes namespaces

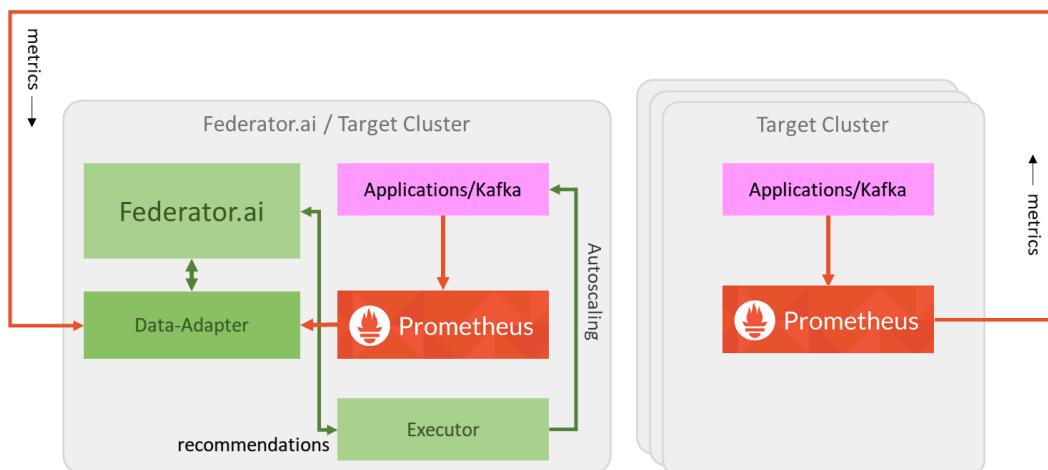
### Supported Metrics Data Sources

There are five different types of metrics data sources supported in release 4.6.1: Prometheus, Datadog, Sysdig, VMWare vCenter and AWS CloudWatch.

#### 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.

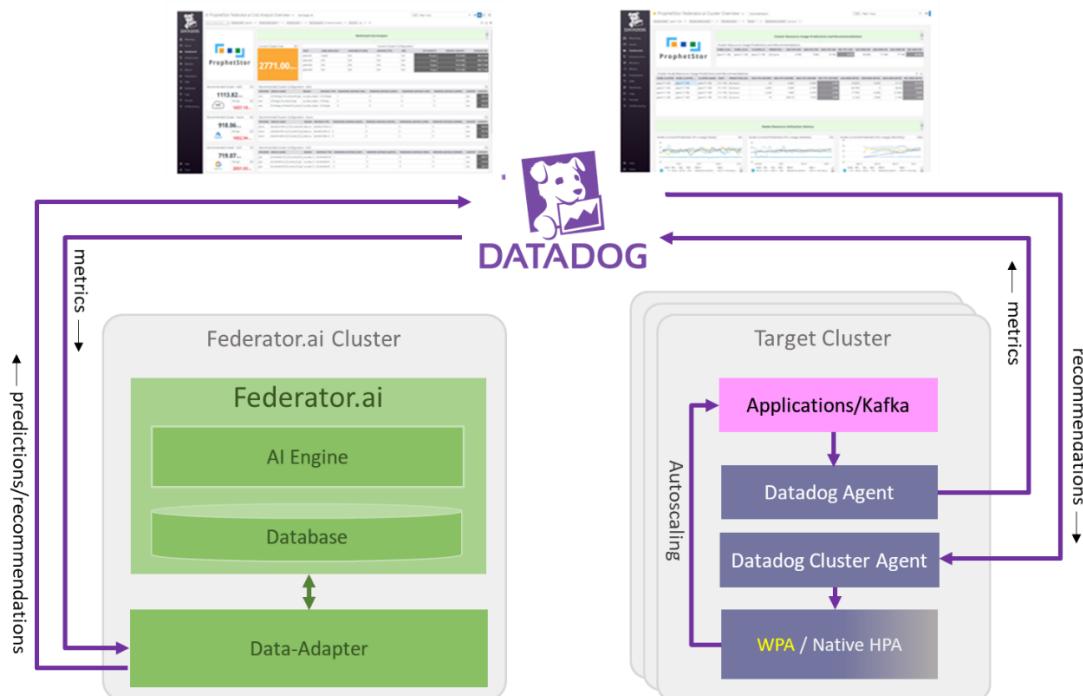
The following diagram shows how the metrics are collected from Prometheus by Federator.ai in a Kubernetes environment.



## Datadog

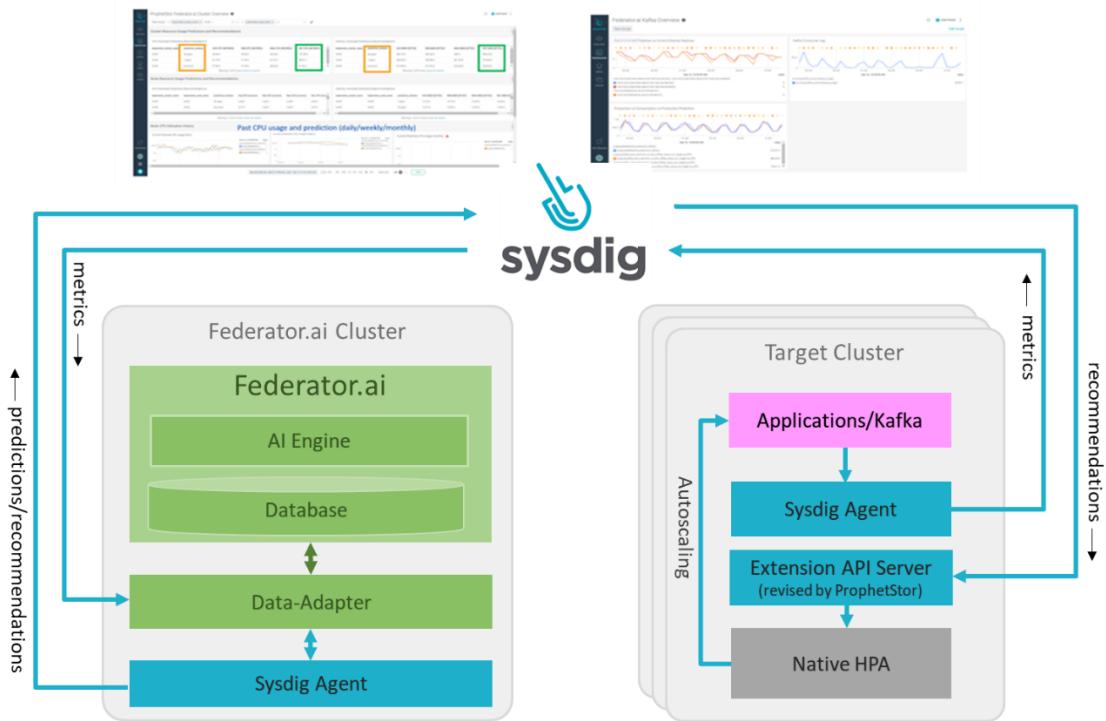
Fedederator.ai has integrated with Datadog and utilizes the metrics collected by Datadog Agent for workload predictions. The following diagram shows how application metrics are used by Federator.ai to predict workload and to automatically scale applications for better performance. Specifically,

- 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



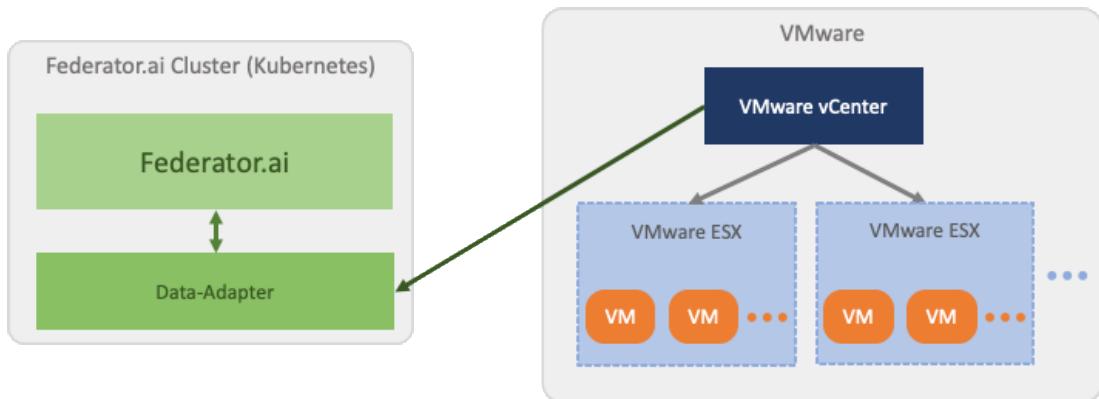
## Sysdig

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.



## VMWare vCenter

VMware vCenter Server provides integrated management of all hosts and virtual machines in the data center from a single console, allowing IT administrators to improve control, simplify daily work, and reduce the complexity and cost of managing the IT environment. Federator.ai data adapter connects to VMware vCenter servers via VMware SDK to retrieve all of VMs workload metrics data for predictions, recommendations and cost analysis for VM clusters.



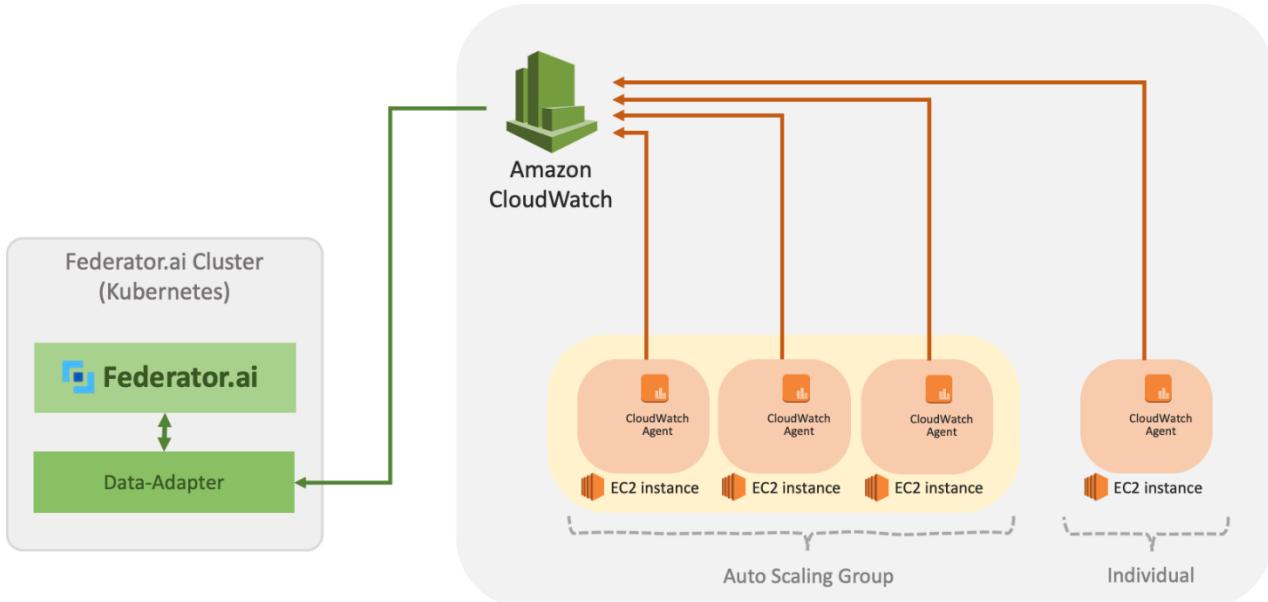
## AWS CloudWatch

AWS CloudWatch is a monitoring service for AWS cloud resources and the applications running on AWS cloud. It provides visibility into resource utilization, operational performance, and overall demand patterns.

The metrics collected by CloudWatch by default do not include memory usage of EC2 instances. CloudWatch agent is required for Federator.ai to collect memory usage metrics.

Federator.ai supports two types of AWS VM clusters:

- Auto Scaling Group
- Individual VM



# Requirements and Recommended Resource Configuration

## Platforms

- OpenShift : 4.x
- Kubernetes : 1.11 ~ 1.20
- Rancher v2.4.8, v2.5.8
- EKS/AKS/GKE

## Data Source

- Datadog
- Sysdig
- Prometheus(Kubernetes, Rancher, OpenShift)
- VMWare vCenter 5.5/6.0/6.5/6.7/7.0
- AWS CloudWatch

## Federator.ai Resource Requirements

- Total Resource Requirements
  - Request : 5.1 CPU cores (Limit :22 cores)
  - Request : 5.0 GB Memory (Limit : 28GB)
  - StorageClass: 168GB (require ReadWriteOnce access mode)
- Resource requirements for AI Engine
  - There must be at least one worker node with at least 2 CPU(Limit : 8 cores) cores and 1 GB of memory available
  - The 2 CPU cores and 1 GB memory are included in the total 5.1 CPU cores and 5.0 GB memory requirements

## Federator.ai Version

- Version: Release v4.6.1
- Tag : v4.6.1-ga

## Datadog Agent Version(reference)

- Datadog Agent helm chart version: v2.4.24, v2.13.0
- Datadog Agent version: v7.21.1, v7.27.0
- Datadog Cluster Agent version: v1.8.0, v1.12.0
- Datadog Watermark Pod Autoscaler version: v0.1.0

## Prometheus Version(reference)

- OpenShift
  - Default installed Prometheus

- Rancher
  - Default bundled 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-12.3.0/12.5.0/15.4.6

## Sysdig Agent Version(reference)

- Sysdig agent: 11.2.0+

## 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. Any data on ephemeral storage will be lost after Federator.ai pods are restarted.

For Federator.ai's application-aware resource/performance optimization feature, the following versions of applications are supported:

## Kafka

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

## NGINX as Ingress

- “nginx-ingress-controller” 0.23.0+ (0.23.0 release on 2019 Feb 28)  
<https://github.com/kubernetes/ingress-nginx/blob/master/Changelog.md#0230>
- NGINX Plus 1.11.3+  
<https://artifacthub.io/packages/helm/nginx/nginx-ingress>

# 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.
- For VMware vCenter, obtain administrator login credential and vCenter IP or FQDN.
- For AWS CloudWatch, obtain Access Key ID and Secret Access Key of AWS account.

### 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?	1.11~1.20	Use the command below to get the Kubernetes version:  \$ kubectl version ... 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 168GB.	The minimum storage size for Federator.ai Release v4.6.1 is 168GB, including database, data, and logs.
4	Kubernetes cluster CPU/memory requirement	Minimum CPU/mem/storage: <ul style="list-style-type: none"> <li>- CPU: 5,100 (mcores)</li> <li>- Memory: 5.0 (GB)</li> <li>- Storage Class Capacity: 168GB</li> </ul> At least one worker node with <ul style="list-style-type: none"> <li>- CPU: 2 Cores</li> <li>- Memory: 1GB</li> </ul>	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 5.1 CPU Cores and 5 GB memory requirements.
5	Is this Kubernetes cluster allowed for NodePort configuration?	Federator.ai creates two NodePorts for GUI and REST API by default <ul style="list-style-type: none"> <li>- REST API - https://&lt;server&gt;:31011</li> <li>- GUI - https://&lt;server&gt;:31012</li> </ul>	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 <ul style="list-style-type: none"> <li>- CPU: 5.1 Cores</li> <li>- Memory: 5.0 GB</li> </ul>	The CPU/memory required for Federator.ai depends on the number of clusters and applications being monitored/managed.  Suggestion for initial namespace quota is <ul style="list-style-type: none"> <li>- CPU 8 cores</li> <li>- Memory 12G</li> </ul> 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 prometheus</b> prometheus-adapter monitoring 1 2020-03-13 15:35:05.28963154 +0800 CST deployed prometheus-adapter-2.1.3 v0.6.0

		<pre> 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].im age  grep -i prometheus directxman12/k8s-prometheus-adapter-amd64:v0.6.0 quay.io/coreos/prometheus-operator:v0.37.0 </pre>
--	--	---

## 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	<p>Use the command to confirm Kafka integration is enabled</p> <pre>\$ kubectl exec &lt;datadog-agent-pod&gt; -n &lt;datadog-agent-namespace&gt; -- agent integration show datadog-kafka-consumer</pre> <p>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</p>
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 configurated 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 agernt 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> <b>enabled: false true</b> 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> <b>clusterName: &lt;cluster-name&gt;</b> - enable cluster agent ...

		<pre> clusterAgent:   enabled: false true ... - \$helm upgrade ... - Check "DD_Cluster_Name" \$kubectl get daemonset &lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml  - name: DD_CLUSTER_NAME   value: &lt;cluster-name&gt;  3.&gt;In Datadog Agent and Cluster Agent installed environment, with no cluster_name setting Update Datadog Agent by "helm upgrade" - assign a cluster name datadog:   clusterName: &lt;cluster-name&gt; - \$helm upgrade ... - Check "DD_Cluster_Name" \$kubectl get daemonset &lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml As: - name: DD_CLUSTER_NAME   value: &lt;cluster-name&gt;  4.&gt; In Datadog Agent and Cluster Agent installed environment, with cluster_name setting Use the command below to confirm DD_Cluster_Name - \$kubectl get daemonset &lt;datadog_agent_daemonset_name&gt; -n &lt;datadog_agnet_namespace&gt; -o yaml As: - name: DD_CLUSTER_NAME   value: &lt;cluster-name&gt; </pre>
--	--	--

## Before You Start

### Datadog

- The admin role for installing Fedeator.ai is "Cluster Admin."
- Datadog agent must be ready if Fedeator.ai runs in the same Kubernetes cluster that is being monitored.
- Obtain Datadog account API Key, Application Key.
  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 navigation links 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 contains 'APIs' (selected), Agent, and Embeds.

## Copy the API Key and Application Key for Federator.ai metrics data source configuration

The screenshot shows two sections of the Datadog API Keys page. The top section, 'API Keys', lists a single key (ID 1) with details: Name '8a94db95ec0f9e3fca9711b89e8a63c', Key 'redacted', Created by 'kyv745.chen@gmail.com', and Created at (UTC) '2020-05-21 08:31:24'. A 'Revoke' button is present. The bottom section, 'Application Keys', lists a single key (ID 2) with details: Name 'pod', Key 'redacted', Created by 'kyv745.chen@gmail.com', and a 'Revoke' button. Both sections have a 'Create API Key' or 'Create Application Key' button at the bottom.

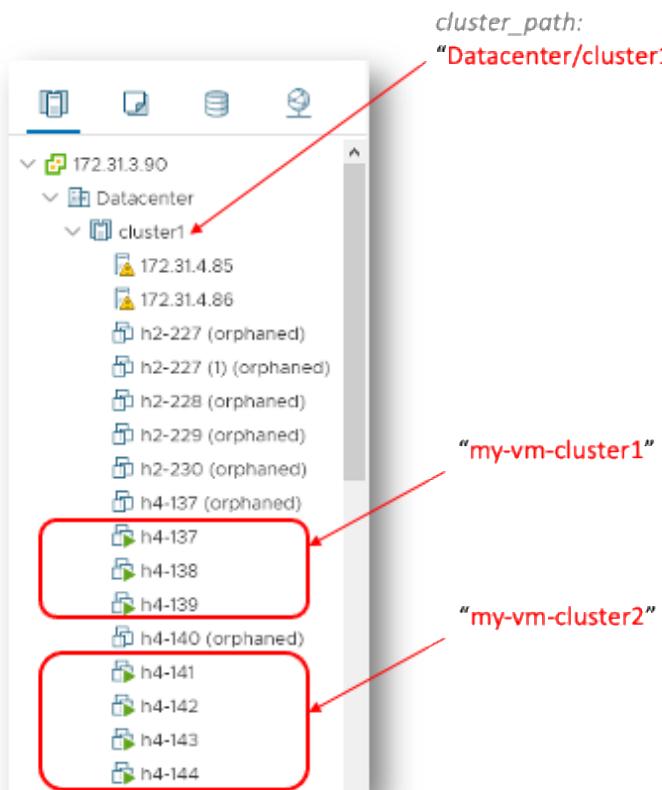
## Sysdig

- Different Sysdig API URL is needed for different regions :
  - For US East, Sysdig API URL is <https://app.sysdigcloud.com>
  - For US West, Sysdig API URL is <https://us2.app.sysdig.com>
  - For European Union, Sysdig API URL is <https://eu1.app.sysdig.com>
- Copy Sysdig Monitor API Token for Federator.ai metrics data source configuration

The screenshot shows the 'User Profile' section of the Sysdig Cloud settings. It includes fields for Account Name (stella.chen@prophetstor.com), Role (Admin), and Current Team (Monitor Operations). Under 'Admin Settings', there is a toggle for 'Hide Agent Install'. Below this is the 'Sysdig Monitor API' section, which contains a note about the REST API and tokens. A red box highlights the 'Sysdig Monitor API Token' input field, and a red button labeled 'RESET TOKEN' is visible.

## VMware vCenter

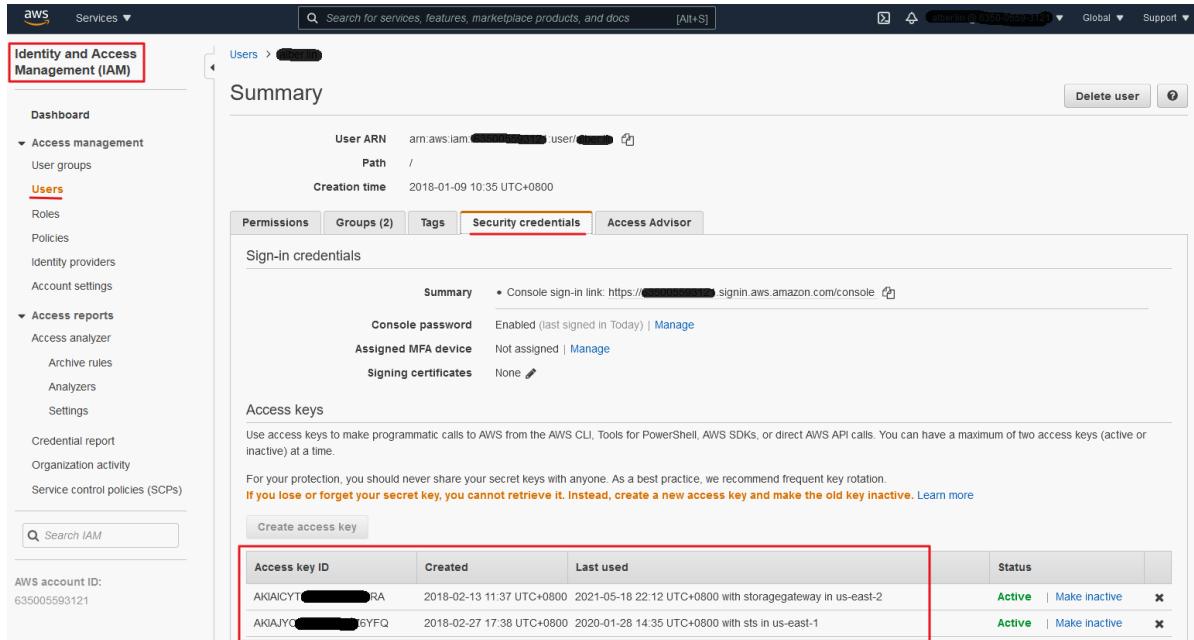
- You can define a VM cluster from any VMs under the same cluster path. See below for an example of cluster path on vCenter.



## AWS CloudWatch

- Obtain CloudWatch Account Key ID and Secret Access Key.

1. Use your AWS account ID or account alias, your IAM user name, and your password to sign in to the [IAM console](#).
2. Go to “Access management > Users > Security credentials” to get Access Key ID and Secret Access Key



The screenshot shows the AWS Identity and Access Management (IAM) service interface. On the left, a navigation pane lists various IAM management options like Dashboard, Access management, Roles, Policies, and Access reports. The 'Users' option is selected. The main area displays a 'Summary' for a specific user, including their User ARN, Path, and Creation time. Below this, tabs for 'Permissions', 'Groups (2)', 'Tags', and 'Security credentials' are present, with 'Security credentials' being the active tab. Under 'Sign-in credentials', there's a summary link and details for a console password (Enabled). Under 'Access keys', a note about key rotation is shown, followed by a warning: 'If you lose or forget your secret key, you cannot retrieve it. Instead, create a new access key and make the old key inactive.' A 'Create access key' button is available. A table lists two access keys:

Access key ID	Created	Last used	Status
AKIAICYT[REDACTED]RA	2018-02-13 11:37 UTC+0800	2021-05-18 22:12 UTC+0800 with storagegateway in us-east-2	Active   Make inactive ×
AKIAJYC[REDACTED]6YFQ	2018-02-27 17:38 UTC+0800	2020-01-28 14:35 UTC+0800 with sts in us-east-1	Active   Make inactive ×

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_access-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)

## New Installation

(For upgrading 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    Left  Speed
100 17101  100 17101     0      0  30118       0 --:--:-- --:--:-- --:--:-- 30107
Please enter Federator.ai version tag [default: latest]: v4.6.1-ga
Federator.ai version = v4.6.1-ga
Please enter the path of Federator.ai directory [default: /opt]:

Downloading v4.6.1-ga tgz file ...
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 = v4.6.1-ga
install_namespace = federatorai
-----

Is the above information correct? [default: y]:

Downloading v4.6.1-ga tgz file ...
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-service created
Applying 09-secret.yaml...
secret/federatorai-operator-service-cert 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-669566b7c-rmphp 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 v4.6.1-ga successfully

Downloading Federator.ai alamedaservice sample file ...
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., 2 for 2GB, default: 2]:
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 = 2 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(v4.6.1-ga) pod to appear ...

datahub pod is present.

Checking pods...
Waiting for pod alameda-ai-66f5c7b6b4-rx87j in namespace federatorai to be ready. phase:
[Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-66f5c7b6b4-rx87j in namespace federatorai to be ready. phase:
[Pending]
Waiting for pods in namespace federatorai to be ready...
```

```

Waiting for pod alameda-ai-66f5c7b6b4-rx87j in namespace federatorai to be ready. phase:
[Pending]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-66f5c7b6b4-rx87j in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-operator-7ff69f4bb5-v22ws 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 v4.6.1-ga successfully

Downloaded YAML files are located under /opt/federatorai/installation

Downloaded files are located under /opt/federatorai/repo/v4.6.1-ga

```

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

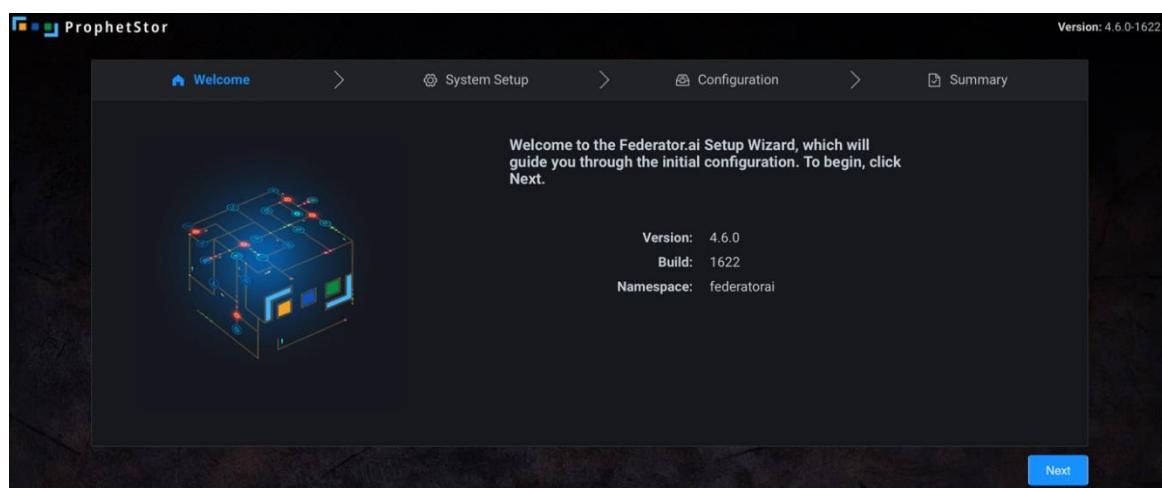
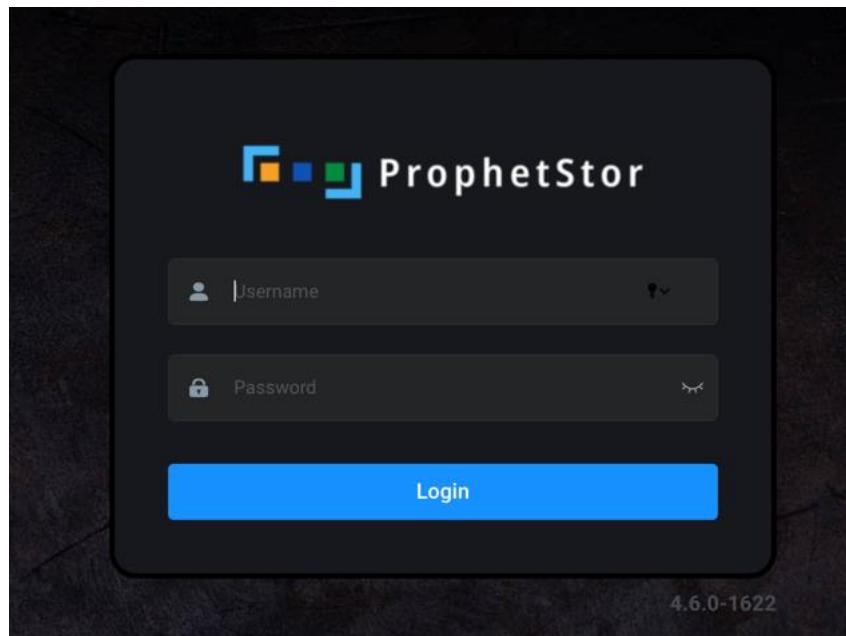
NAME	READY	STATUS	RESTARTS	AGE
alameda-ai-66f5c7b6b4-rx87j	1/1	Running	0	16m
alameda-ai-dispatcher-78d8556bd5-c4h2l	1/1	Running	0	16m
alameda-analyzer-668566d588-6cprq	1/1	Running	0	16m
alameda-datahub-79f85cd56-ff8j6	1/1	Running	1	16m
alameda-executor-7f9d899578-85jxz	1/1	Running	3	16m
alameda-influxdb-0	1/1	Running	0	16m
alameda-notifier-8bb7cc889-c6b8x	1/1	Running	2	16m
alameda-operator-7ff69f4bb5-v22ws	1/1	Running	5	16m
alameda-rabbitmq-d5868cd97-n6zgt	1/1	Running	0	16m
fedemeter-api-5cb94ff9c9-dbnxw	1/1	Running	0	16m
fedemeter-influxdb-0	1/1	Running	0	16m
federatorai-agent-7cccf55f84-jjjb4	1/1	Running	0	16m
federatorai-dashboard-backend-5fd697fd86-kks5n	1/1	Running	0	16m
federatorai-dashboard-frontend-59595cc866-6578j	1/1	Running	0	16m
federatorai-data-adapter-76bc6ff8b7-xrdrk	1/1	Running	0	16m
federatorai-operator-669566b7c-rmphp	1/1	Running	0	18m
federatorai-recommender-dispatcher-77b974bbd6-cbvlt	1/1	Running	0	16m
federatorai-recommender-worker-554c7f8694-w9w5j	1/1	Running	0	16m
federatorai-rest-75bf7dff54-xsmh8	1/1	Running	1	16m

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 4.6.1 User Guide.

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

Login ID: admin

Password: admin



## Upgrade from a previous version

Federator.ai v4.6.1 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 roll back to the previous version is needed.

### Prerequisite

1. Federator.ai version is 4.4.0 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 enter Federator.ai version tag [default: latest]: v.4.6.1-ga
Please input Federator.ai files save path [default: /opt]:
```

```
Downloading v4.6.1-ga tgz file ...
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
Previous build with tag v4.4.1-b1462 detected in namespace federatorai
```

```
-----
Upgrade:
tag_number = v4.6.1-ga
install_namespace = federatorai
```

```
-----
Is the above information correct? [default: y]:
Do you want to backup your configuration before upgrading Federator.ai? [default: y]: y
Please input path for storing backup configuration: [default:
/opt/federatorai/configuration_backup]
Backup configuration...
backup yaml saved to folder /opt/federatorai/configuration_backup/federatorai-backup-
1616551135
Done.
```

```
Downloading v4.6.1-ga tgz file ...
Done
```

```
Updating InfluxDB owner...
Done

Applying Federator.ai operator yaml files...
deployment.apps "federatorai-operator" deleted
Applying 00-namespace.yaml...
namespace/federatorai unchanged
Applying 01-serviceaccount.yaml...
serviceaccount/federatorai-operator unchanged
Applying 02-alamedaservice.crd.yaml...
customresourcedefinition.apiextensions.k8s.io/alamedaservices.federatorai.containers.ai
configured
Delay applying 03-federatorai-operator.deployment.yaml
Applying 04-clusterrole.yaml...
clusterrole.rbac.authorization.k8s.io/federatorai-operator configured
clusterrole.rbac.authorization.k8s.io/alameda-gc configured
Applying 05-clusterrolebinding.yaml...
clusterrolebinding.rbac.authorization.k8s.io/federatorai-operator unchanged
Applying 06-role.yaml...
role.rbac.authorization.k8s.io/federatorai-operator configured
Applying 07-rolebinding.yaml...
rolebinding.rbac.authorization.k8s.io/federatorai-operator unchanged
Applying 08-service.yaml...
service/federatorai-operator-service created
Applying 09-secret.yaml...
secret/federatorai-operator-service-cert 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
Applying 03-federatorai-operator.deployment.yaml...
deployment.apps/federatorai-operator created

federatorai-operator pod is present.

Waiting for pod federatorai-operator-75bdc65496-tz6fx in namespace federatorai to be
ready ... phase: [Pending]
Waiting for pod federatorai-operator-75bdc65496-tz6fx in namespace federatorai to be
ready ... phase: [Pending]
Waiting for pod federatorai-operator-75bdc65496-tz6fx in namespace federatorai to be
ready ... phase: [Pending]
Waiting for pod federatorai-operator-75bdc65496-tz6fx in namespace federatorai to be
ready ... phase: [Running]

federatorai-operator pod is ready.

Downloading Federator.ai alamedaservice sample file ...
Done

Downloading Federator.ai alamedascaler sample files ...
Done
=====
Update alamedaservice...
alamedaservice.federatorai.containers.ai/my-alamedaservice patched
Done.
alamedaservice.federatorai.containers.ai/my-alamedaservice patched
```

```

alamedaservice.federatorai.containers.ai/my-alamedaservice patched
Processing...
Waiting for datahub(v4.6.1-ga) pod to appear ...

datahub pod is present.

Checking pods...
Waiting for pod alameda-ai-667df48565-99qft in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-667df48565-99qft in namespace federatorai to be ready. phase:
[Running]
Waiting for pods in namespace federatorai to be ready...
Waiting for pod alameda-ai-dispatcher-6845456b68-8kjfr 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 v4.6.1-ga successfully

Downloaded YAML files are located under /opt/federatorai/installation

Downloaded files are located under /opt/federatorai/repo/v4.6.1-ga

```

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

```

~# kubectl get pod -n federatorai
NAME                                         READY   STATUS    RESTARTS   AGE
alameda-ai-6b56d6db77-x2r9x                   1/1     Running   0          19m
alameda-ai-dispatcher-7d46f46849-nd4z2        1/1     Running   3          19m
alameda-analyzer-56bd4d4f8d-jvw6f             1/1     Running   0          19m
alameda-datahub-597fb6f964-pqmh6              1/1     Running   3          20m
alameda-executor-6b4bff9b47-d9fdw             1/1     Running   4          19m
alameda-influxdb-0                            1/1     Running   0          19m
alameda-notifier-87cf6b94c-xn5pz              1/1     Running   1          19m
alameda-operator-cf64fb6c9-ck7lm               1/1     Running   0          20m
alameda-rabbitmq-dddcc8dd7-99jv2              1/1     Running   0          19m

```

fedemeter-api-9f49d898f-v588q	1/1	Running	0	19m
fedemeter-influxdb-0	1/1	Running	0	19m
federatorai-agent-674dcfd448-mx4ln	1/1	Running	0	19m
federatorai-dashboard-backend-844db587cb-8vvkb	1/1	Running	0	19m
federatorai-dashboard-frontend-8dff898c6-rs2q5	1/1	Running	0	19m
federatorai-data-adapter-7885c78db4-chqml	1/1	Running	0	19m
federatorai-operator-5bb7d58d9d-x92q9	1/1	Running	0	22m
federatorai-recommender-dispatcher-6cbbd8896f-5wdv2	1/1	Running	0	19m
federatorai-recommender-worker-7bb5f4b94c-tt887	1/1	Running	0	19m
federatorai-rest-5685456bb7-dq6nq	1/1	Running	2	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

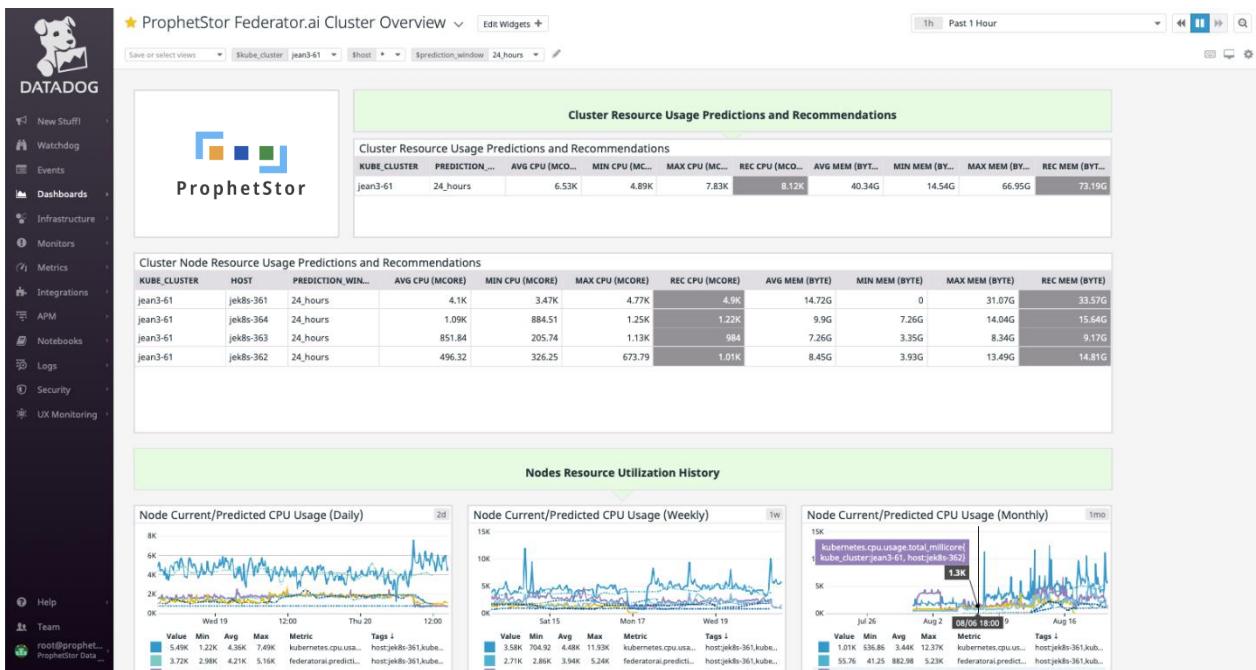
```
~# kubectl apply -f wpa.yaml
```

## Appendix

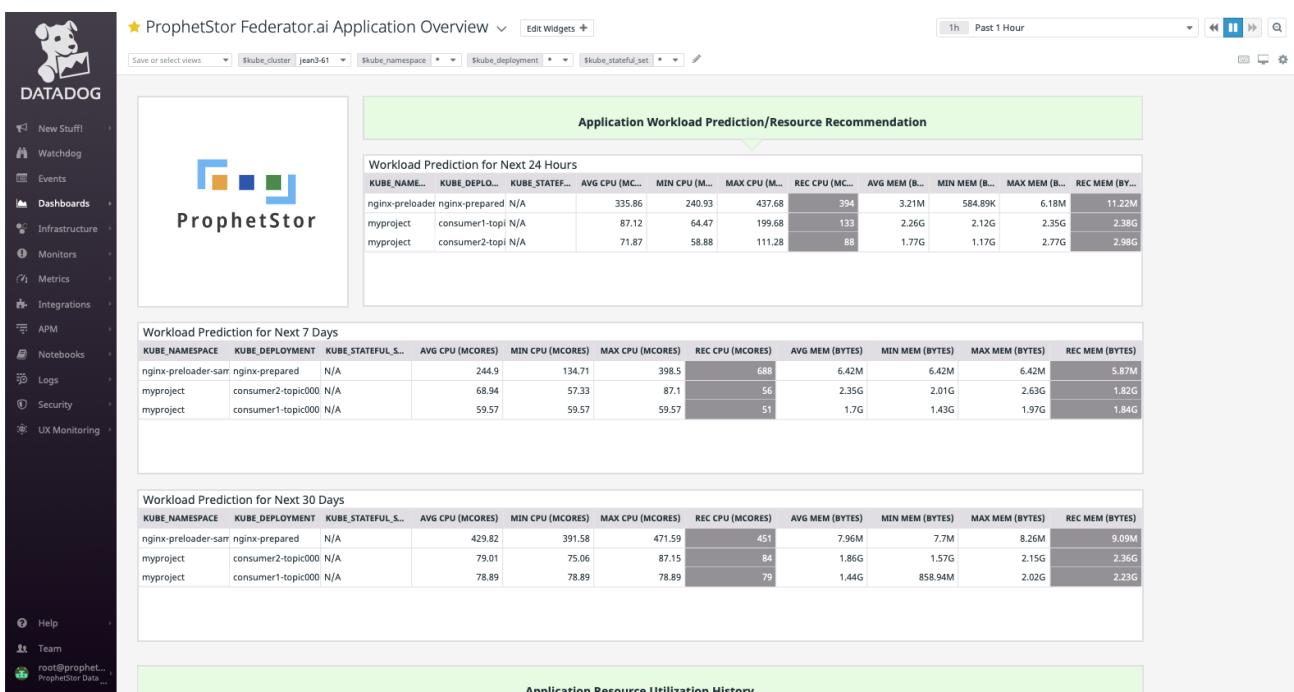
### Datadog Dashboards Overview

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

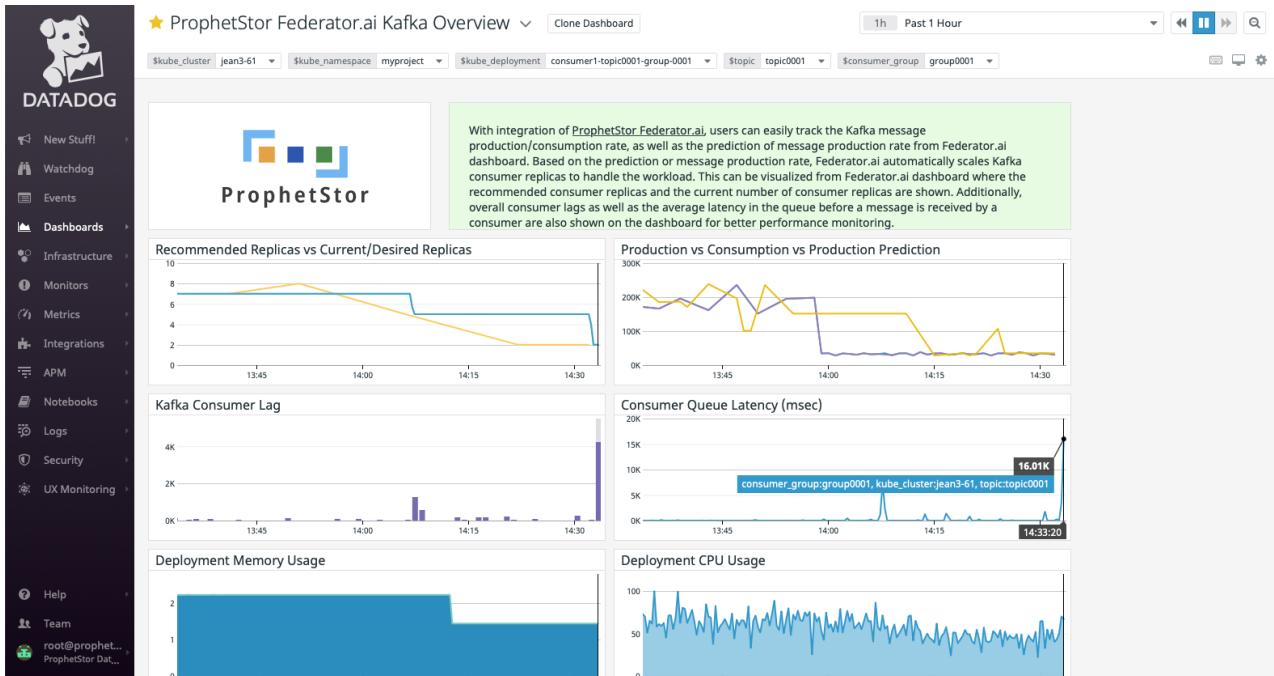
#### ProphetStor Federator.ai Cluster Overview



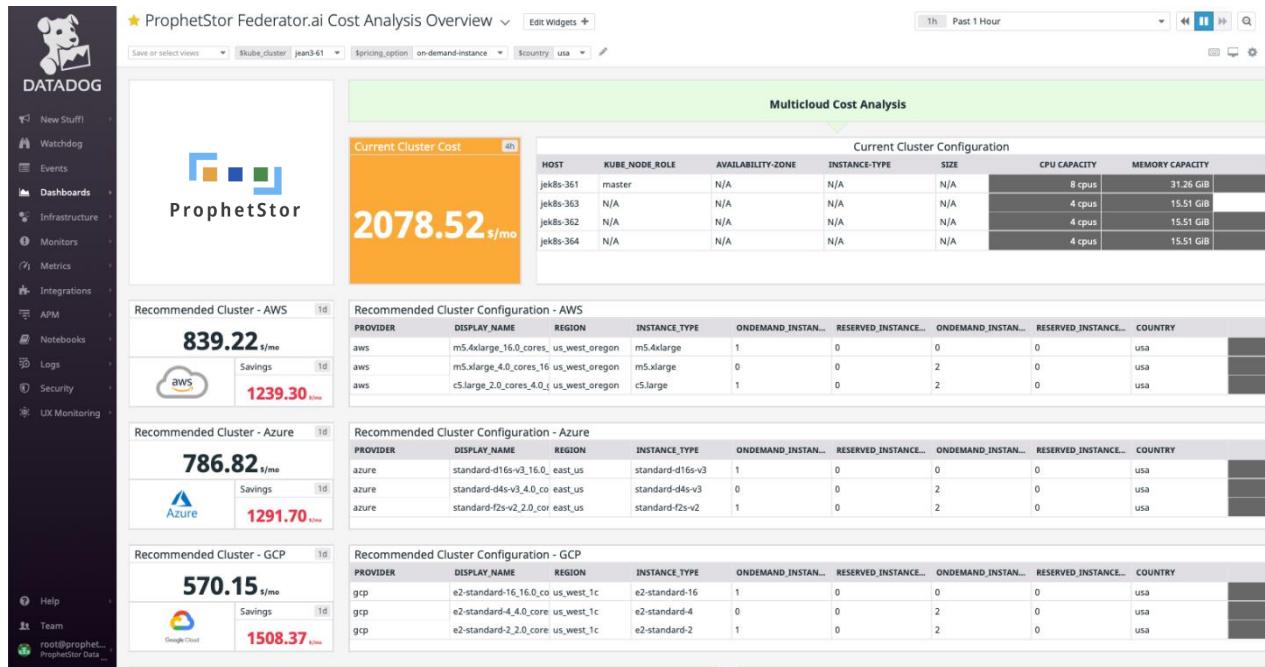
#### ProphetStor Federator.ai Application Overview



## ProphetStor Federator.ai Kafka Overview



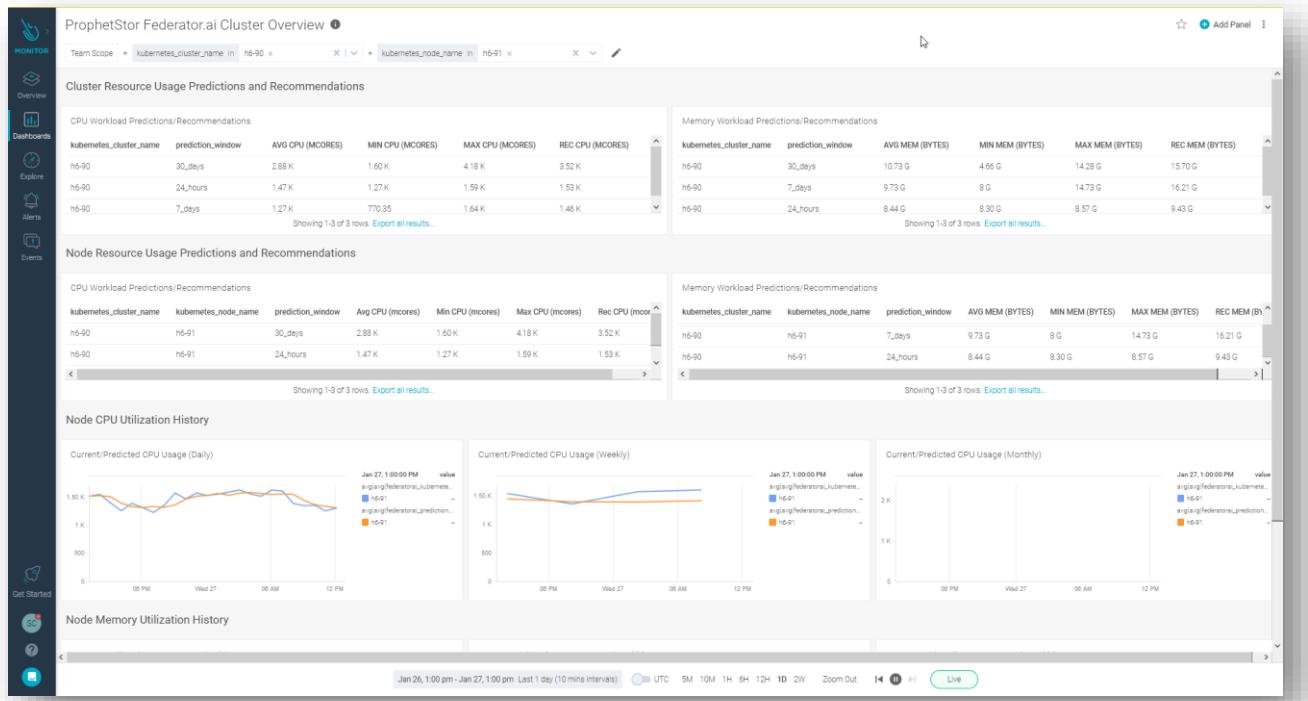
## ProphetStor Federator.ai Cost Analysis Overview



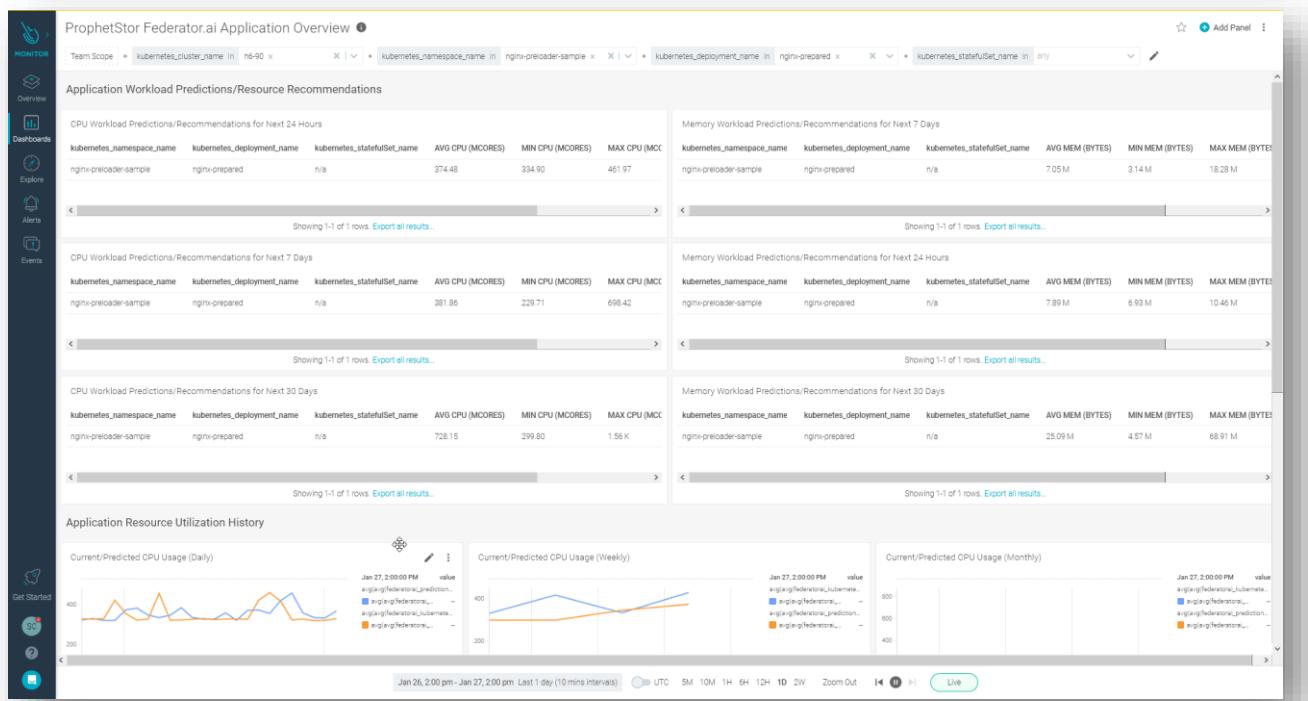
# Sysdig Dashboard Overview

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

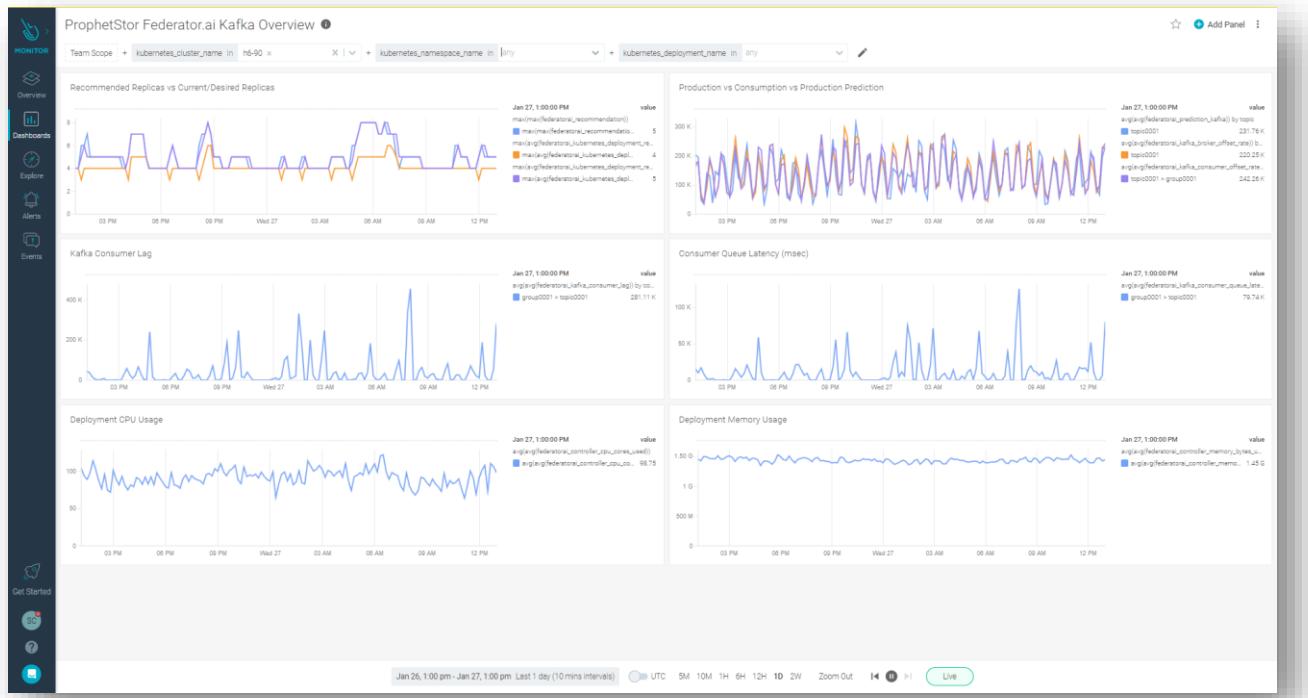
## Federator.ai Cluster Overview



## Federator.ai Application Overview



## Federator.ai Application Overview



## Federator.ai installation/uninstallation using Helm Chart

### Prerequisites

- [Kubernetes](#) version 1.18 or later
- [OpenShift](#) version 4.x.x or later
- [Helm](#) version is 3.x.x or later

### Add Helm chart repository

```
~# helm repo add prophetstor https://prophetstor-ai.github.io/federatorai-operator-helm/
```

### Test the Helm chart repository

```
~# helm search repo federatorai
```

### Installing with the release name my-name:

```
~# helm install `my-name` prophetstor/federatorai --namespace=federatorai --create-namespace
```

### To uninstall/delete the my-name deployment:

```
~# helm ls --all-namespaces  
helm delete `my-name` --namespace=federatorai
```

## Configuration

The following table lists the configurable parameters of the chart and their default values are specified in values.yaml.

Parameter	Description
image.pullPolicy	Container pull policy
image.repository	Image for Federator.ai operator
image.tag	Image Tag for Federator.ai operator
federatorai.imageLocation	Image Location for services containers
federatorai.version	Image Tag for services containers
federatorai.persistence.enabled	Enable persistent volumes
federatorai.persistence.storageClass	Storage Class Name of persistent volumes
federatorai.persistence.storages.logStorage.size	Log volume size
federatorai.persistence.aiCore.dataStorage.size	AICore data volume size
federatorai.persistence.influxdb.dataStorage.size	Influxdb data volume size
federatorai.persistence.fedemeterInfluxdb.dataStorage.size	Fedemeter influxdb data volume size
services.dashboardFrontend.nodePort	Port of the Dashboard service

Specify each parameter using the --set key=value[,key=value] argument to helm install

Tip: You can use the default values.yaml

Sample :values.yaml

```
## Default values for Federator.ai
## This is a YAML-formatted file.
## Declare variables to be passed into your templates.
##
image:
  pullPolicy: IfNotPresent
  repository: quay.io/prophetstor/federatorai-operator-ubi
  tag: v4.6.1-ga

## Set default values
##
federatorai:
  imageLocation: quay.io/prophetstor
  version: v4.6.1-ga
```

```

## If the persistence is enabled, a default StorageClass
## is needed in the k8s cluster to provision volumes.
persistence:
  enabled: true
  storageClass: "standard"
  storages:
    logStorage:
      size: 2Gi
  aiCore:
    dataStorage:
      size: 10Gi
  influxdb:
    dataStorage:
      size: 100Gi
  fedemeterInfluxdb:
    dataStorage:
      size: 10Gi

services:
  dashboardFrontend:
    ## Specify the nodePort value for the dashboard frontend
    ## Comment out the following line to disable nodePort service
    nodePort: 31012
  rest:
    ## Specify the nodePort value for the REST service
    ## Comment out the following line to disable nodePort service
    nodePort: 31011

```

### Alternative installation with configuration file

A YAML file that specifies the values for the parameters can be provided while installing the chart. For example

```

~# helm install `my-name` prophetstor/federatorai -f values.yaml --
namespace=federatorai --create-namespace

```

## 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”  

```
~# 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  

```
~# 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
Federator.ai env	federatorai_version	v4.6.1-ga	Federator.ai version tag	Y
Storage for Federator.ai 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 Federator.ai required docker images	N
Pod resource	enable_resource_requiment	y	Add pod resource requirement (limits & requests) for every Federator.ai 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 Federator.ai 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

```
~# 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).

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

4. Run Ansible playbook

```
~# 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 Federator.ai pods	storage_type	ephemeral or persistent	Specify current Federator.ai 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 Federator.ai 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 v4.6.1

- **V4.6.1->v4.5.x**

Following v4.5.1 installation step to rollback to previous version v4.5.1 directly. v4.5.1 installation script automatically keeps existing configuration, metrics, and prediction data, which is stored on persistent volume.

- **v4.6.1 ->v4.4.x**

Following v4.4.1 installation step to rollback to previous version v4.4.1 directly. v4.4.1 installation script automatically keeps existing configuration, metrics, and prediction data, which is stored on persistent volume.

- **v4.6.1-> v4.2**

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

1. Run v4.6.1/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 `/opt/federatorai/repo/v4.6.1-ga/scripts/` directory.

```
~# cd /opt/federatorai/repo/v4.6.1-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-936acfff0d30 ...
persistentvolume/pvc-b517207b-54b6-4a42-81da-936acfff0d30 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-dccdcfb7-6f7f-46a9-9388-8a5b97e7126d ...
persistentvolume/pvc-dccdcfb7-6f7f-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: v4.6.1-ga
```

```
-----  
Your tag number = v4.6.1-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 /opt/federatorai/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 /opt/federatorai/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 ...
/opt/federatorai/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

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\tgithub.com/watermarkpodautoscaler/pkg/controller/watermarkpodautoscaler.(*ReconcileWatermarkPodAutoscaler).Reconcile\n\tgithub.com/watermarkpodautoscaler/pkg/controller/watermarkpodautoscaler/watermarkpodautoscaler_controller.go:345\n\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).reconcileHandler\n\tgithub.com/watermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:216\n\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).processNextWorkItem\n\tgithub.com/watermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:192\n\tsig.sigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).worker\n\tgithub.com/watermarkpodautoscaler/vendor/sigs.k8s.io/controller-runtime/pkg/internal/controller/controller.go:171\n\tk8s.io/apimachinery/pkg/util/wait.JitterUntil.func1\n\tgithub.com/watermarkpodautoscaler/vendor/k8s.io/apimachinery/pkg/util/wait/wait.go:152\n\tk8s.io/apimachinery/pkg/util/wait/wait.go:153\n\tk8s.io/apimachinery/pkg/util/wait.Until\n\tgithub.com/watermarkpodautoscaler/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