

# Federatora.ai Release 4.7

## Release Notes

August 2021

# Contents

- Overview ..... 3**
  
- Version 4.7 ..... 4**
  - Supported Metrics Data Sources ..... 4
  - Supported Platforms ..... 4
  - Enhancements in Release 4.7 ..... 4
    - Management ..... 4
  
- Version 4.6.1 ..... 5**
  - Supported Metrics Data Sources ..... 5
  - Supported Platforms ..... 5
  - Enhancements in Release 4.6.1 ..... 5
    - Integration ..... 5
  
- Version 4.6 ..... 5**
  - Supported Metrics Data Sources ..... 5
  - Supported Platforms ..... 5
  - Enhancements in Release 4.6 ..... 6
    - Integration ..... 6
    - Management ..... 6
  - Known Issues and Limitations in Release 4.6 ..... 6
  
- Version 4.5.x ..... 7**
  - Supported Metrics Data Sources ..... 7
  - Supported Platforms ..... 7
  - Enhancements in Release 4.5.1 ..... 7
    - Management ..... 7
    - Installation and Configuration ..... 7
  - Enhancements in Release 4.5 ..... 7

Integration .....	7
Management .....	7
UI Enhancements.....	8

**Version 4.4.x ..... 8**

Supported Metrics Data Sources.....	8
Supported Platforms .....	8
Enhancements in Release 4.4.1.....	8
UI Enhancement .....	8
Resolved Issues in Release 4.4.1 .....	8
Enhancements in Release 4.4.....	8
Integration .....	8
Multi-Cloud Cost Analysis.....	9
Installation and Configuration .....	9
UI Enhancements.....	9

## Overview

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 clusters.

Using advanced machine learning algorithms to predict application workloads, Federator.ai offers:

- AI-based workload prediction for containerized applications in Kubernetes clusters as well as VMs in VMware clusters and Amazon Web Services (AWS) Elastic Compute Cloud (EC2)
- Resource recommendations based on workload prediction, application, Kubernetes, and other related metrics
- Automatic provisioning of CPU/memory for generic Kubernetes application controllers/namespaces
- Automatic scaling of Kubernetes application containers, Kafka consumer groups, and Ingress upstream services
- 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

This document contains the release notes for Federator.ai Release 4.7, including information about new features and enhancements, as well as known issues. It also includes release note information from previous releases.

## Version 4.7

### Supported Metrics Data Sources

- Prometheus
- Datadog
- Sysdig
- VMware vCenter
- AWS CloudWatch

### Supported Platforms

- Kubernetes v1.11.x – v1.22.x
- Red Hat OpenShift – v4.x
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, v2.5.9
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

### Enhancements in Release 4.7

#### Management

- Support historical data collection for Kubernetes clusters and generic applications, enabling weekly and monthly predictions, recommendations, and cost analysis for newly added clusters/applications without waiting to collect weeks' or months' worth of data.
- Provide APIs for cost management, including cost analysis for clusters, namespaces, and applications, as well as cluster configuration recommendations for public cloud services with the lowest cost.
- License Federator.ai based on resource usage instead of the number of licensed CPU cores to more accurately reflect customer usage and control costs based on actual needs.

### Known Issues and Limitations in Release 4.7

- Federator.ai uses the APIs provided by your metrics data source (Datadog, Sysdig) to access historical data. The data source imposes limits on how many calls can be made to their service per hour. If the rate limit is too low, the queries for historical data may exceed the limit and the API will return an error. You will need to contact your metrics data service provider to raise your API rate limit.

## Version 4.6.1

### Supported Metrics Data Sources

- Prometheus
- Datadog
- Sysdig
- VMware vCenter
- AWS CloudWatch

### Supported Platforms

- Kubernetes v1.11.x – v1.20.x
- Red Hat OpenShift – v3.11, v4.x
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, v2.5.8
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

### Enhancements in Release 4.6.1

#### Integration

- Support automatic scaling for HTTP services with NGINX Plus in Kubernetes clusters that use Prometheus as the metrics data source.

## Version 4.6

### Supported Metrics Data Sources

- Prometheus
- Datadog
- Sysdig
- VMware vCenter
- AWS CloudWatch

### Supported Platforms

- Kubernetes v1.11.x – v1.20.x
- Red Hat OpenShift – v3.11, v4.x
- Amazon AWS/EKS

- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, v2.5.8
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

## Enhancements in Release 4.6

### Integration

- Support VM clusters on AWS EC2.
- Support AWS CloudWatch as a data source for metrics in VM clusters.
- Support Federation for Prometheus.
- Support Continuous Integration/Continuous Delivery (CI/CD) integration with Terraform.

### Management

- Support automatic resource provisioning of Kubernetes namespaces and controllers for generic applications.
- Support automatic scaling for HTTP services with NGINX Ingress in Kubernetes clusters.
- Provide the ability to email administrators when system errors and fatal issues occur.
- Enhance the Application Cost Analysis to show actual cost savings for automatic resource provisioning.

## Known Issues and Limitations in Release 4.6

- Sysdig does not support metrics for Disk IO Utilization and Pod Running Status Count. Therefore, no data will be displayed on the *Cluster Health* and *Node Health* pages in the web portal.
- Autoscaling with the Datadog Watermark Pod Autoscaler (WPA) is not supported for *DeploymentConfig* controllers in OpenShift.
- Auto provisioning of resources is not supported for Kubernetes applications that are under operator control.
- When using Prometheus Federation, every cluster needs to be specified with a *target\_label*, including the cluster where the Prometheus master is installed.

## Version 4.5.x

### Supported Metrics Data Sources

- Prometheus
- Datadog
- Sysdig
- VMware vCenter

### Supported Platforms

- Kubernetes v1.11.x – v1.20.x
- Red Hat OpenShift – v3.11, v4.x
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

### Enhancements in Release 4.5.1

#### Management

- Provide the ability to specify different metric data source for each cluster.
- Provide the ability to manage Kubernetes clusters with different metric data sources and VM clusters at the same time.

#### Installation and Configuration

- Support Federator.ai installation through Helm.

### Enhancements in Release 4.5

#### Integration

- Integrate with VMware systems for VM workload predictions and recommendations.

#### Management

- Analyze Kubernetes application costs and provide potential cost/savings based on recommendations.
- Analyze VM costs and provide potential cost/savings based on recommendations.



- Provide the ability to define a custom price book to define your operating costs for local clusters.

### UI Enhancements

- Provide the ability to display CPU and memory capacity for clusters and nodes/VMs on Workload Prediction charts.
- Add the ability to quickly select the number of hours, days, weeks, or months to display data.

## Version 4.4.x

### Supported Metrics Data Sources

- Prometheus
- Datadog
- Sysdig

### Supported Platforms

- Kubernetes v1.11.x – v1.19.x
- Red Hat OpenShift – v3.11, v4.x
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS

### Enhancements in Release 4.4.1

#### UI Enhancement

- Provide a live connection test for external metric data source to immediately determine if a configured metric data source (e.g., Prometheus, Datadog, or Sysdig) is reachable.

### Resolved Issues in Release 4.4.1

- Federator.ai installation fails on block storage persistent volume.
- Failed to retrieve CPU/Memory metrics from Prometheus on Kubernetes v1.12.

### Enhancements in Release 4.4

#### Integration

- Query metrics from the Sysdig monitoring platform for application workload predictions and recommendations.
- Integrate with the Prometheus open-source monitoring system for application workload predictions and recommendations.

## Multi-Cloud Cost Analysis

- Support Spot Instance pricing for multi-cloud cost analysis.
- Provide a cost analysis time series chart that recommends the number and type of instances for future workloads at specific times, which can help minimize resources and costs.

## Installation and Configuration

- Provide integration with Ansible playbook to simplify Federator.ai installation.
- Provide a setup wizard to simplify initial system configuration after software installation.
- Provide graphical configuration of applications, clusters, and system settings.
- Support automatic or manual update of cloud service provider price books.

## UI Enhancements

- Display cluster and application workload predictions and recommendations on the Dashboard.
- Provide the ability to manage clusters:
  - Add/edit/remove clusters.
  - Start/pause monitoring and prediction for all namespaces or a specific namespace.
  - Stop/start collecting metrics and making predictions for all namespaces or a specific namespace.
- Provide the ability to manage applications:
  - Add/edit/remove applications.
  - Add/edit/remove controllers and consumer groups.
- Provide the ability to add/manage the system license.
- Provide a separate Events page for all clusters.
- Display a new Workload Prediction chart with easy-to-read average/minimum/maximum CPU and memory usage, as well as recommendations.