



Federatora.ai Release 5.1

Release Notes

May 2022

Contents

Overview	2
Version 5.1.....	3
Supported Metrics Data Sources.....	3
Supported Platforms	3
Enhancements in Release 5.1.....	3
Version 5.0.....	4
Supported Metrics Data Sources.....	4
Supported Platforms	4
Enhancements in Release 5.0.....	4
Known Issues and Limitations in Release 5.0.....	4

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
- Correlation and causality analysis of microservices/controllers of Kubernetes applications
- 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
- Statistical analysis and predictions based on the correlation between resource usage and application workload

This document contains the release notes for Federator.ai Release 5.1, including information about new features and enhancements, as well as known issues. It also includes release note information from previous releases.

Version 5.1

Supported Metrics Data Sources

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

Supported Platforms

- Kubernetes v1.11.x – 1.22x
- Red Hat OpenShift v4.6-4.9
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, 2.5.8, 2.5.9, 2.6.3
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0
- IBM Cloud/IKS
- Alicloud

Enhancements in Release 5.1

- Alert notifications monitor clusters, nodes, namespaces, applications, and controllers for a variety of conditions, based on predicted usage. Alerts can be emailed when triggered in addition to being viewed in the portal.
- Automatically discover all namespaces and controllers in a cluster to simplify the process of adding an application.
- Provide the ability to back up and restore the system configuration.
- Define custom price books to calculate hourly operating costs for CPU, memory, and storage for on-premises clusters.
- Support Spot instance recommendations for local VM clusters and AWS clusters without AWS Auto Scaling groups.
- KEDA integration enables Horizontal Pod Autoscaling (HPA) in remote Kubernetes clusters to scale the number of replicas of a container based on the recommendations of Federator.ai.
- Display the CPU, memory, network bytes received, and number of network bytes transmitted resource distribution among different controllers and cluster nodes.
- Various UI enhancements including the addition of tabs on the *Predictions and Planning* pages to select the resource level.

Version 5.0

Supported Metrics Data Sources

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

Supported Platforms

- Kubernetes v1.11.x – 1.22x
- Red Hat OpenShift v4.6-4.9
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, 2.5.8, 2.5.9, 2.6.3
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

Enhancements in Release 5.0

- Application Analysis section offers statistical analysis and predictions based on the correlation between resource usage and application workload for Kubernetes.
- Enhanced cost analysis, cost trends, and cost optimization for resources at different levels, including Kubernetes and VM clusters and nodes, as well as Kubernetes namespaces and applications.
- Support application-specific metrics for controllers of Kubernetes applications.
- Provide a wizard to simplify the process of adding an application with multiple controllers, as well as configuring application-aware autoscaling of Kafka consumers or upstream HTTP services of an NGINX Ingress controller.
- Various UI enhancements including ability to zoom into a section of a chart to improve granularity for a specified time frame.

Known Issues and Limitations in Release 5.0

- The Sysdig monitoring service does not support untyped metrics, such as the following application-specific metrics: `mysql_global_status_questions`, `mysql_global_status_threads_connected`, and `mysql_global_status_slow_queries`. Federator.ai will not show the metric values for these untyped metrics if Sysdig is used as a metric data source.