

Federatora.ai Release 5.0

Release Notes

January 2022

Contents

- Overview 2**

- Version 5.0..... 3**
 - Supported Metrics Data Sources..... 3
 - Supported Platforms 3
 - Enhancements in Release 5.0..... 3
 - Known Issues and Limitations in Release 5.0..... 3

- Version 4.7.2..... 4**
 - Supported Metrics Data Sources..... 4
 - Supported Platforms 4
 - Enhancements in Release 4.7.2..... 4
 - Resolved Issues in Release 4.7.2 4

- Version 4.7..... 5**
 - Supported Metrics Data Sources..... 5
 - Supported Platforms 5
 - Enhancements in Release 4.7..... 5
 - Known Issues and Limitations in Release 4.7..... 5

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 5.0, including information about new features and enhancements, as well as known issues. It also includes release note information from previous releases.

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.

Version 4.7.2

Supported Metrics Data Sources

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

Supported Platforms

- Kubernetes v1.11.x – 1.20.x
- Red Hat OpenShift v4.6-4.8
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, 2.5.8
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

Enhancements in Release 4.7.2

- Provide heatmap charts and interactive utilization analysis charts to compare actual CPU and memory usage to your goals for Kubernetes clusters, nodes, applications, and controllers, and for VM clusters and VMs.
- Provide cost efficiency charts to compare the actual cost of resource usage to the cost of allocated resources for Kubernetes and VM clusters.
- Support historical data collection for VM clusters with an AWS CloudWatch metrics data source, enabling weekly and monthly predictions, recommendations, and cost analysis for newly added clusters without waiting to collect weeks' or months' worth of data.

Resolved Issues in Release 4.7.2

- Unable to collect all metrics from OpenShift 4.8.
- Inconsistent error messages in the auto provisioning profile.
- Incorrect status of historical data collection for an application that was deleted and re-added.
- Incorrect error message in the event log when there was a connection failure with the Rancher Prometheus service.
- Added an event log warning when a price book update fails because the Internet is not reachable.
- Invalid actual savings displayed on the Application Cost Analysis page when switching between the daily and weekly view.
- Savings were not displayed in the detailed information for the daily view on the Multicloud Cost Analysis page.

Version 4.7

Supported Metrics Data Sources

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

Supported Platforms

- Kubernetes v1.11.x – 1.20.x
- Red Hat OpenShift v3.11, 4.x
- Amazon AWS/EKS
- Google GCP/GKE
- Microsoft Azure/AKS
- Rancher v2.4.8, 2.5.8
- VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0

Enhancements in Release 4.7

- 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.