



Federator.ai®

Intelligence-Driven Cloud Operation Automation and Optimization



- **Leveraging the operation metadata of the leading monitoring services/solutions and the CI/CD pipeline for optimized and continuous optimization of resource allocation**
- **Computationally feasible multi-layer correlation with predictive analytics for proactive scaling for assured application resilience**
- **Removing the guesswork of how the applications use resources to facilitate planning and resource orchestration**
- **Up to 80% cost savings by eliminating over-provisioning of resources and optimized placements in a MultiCloud environment**

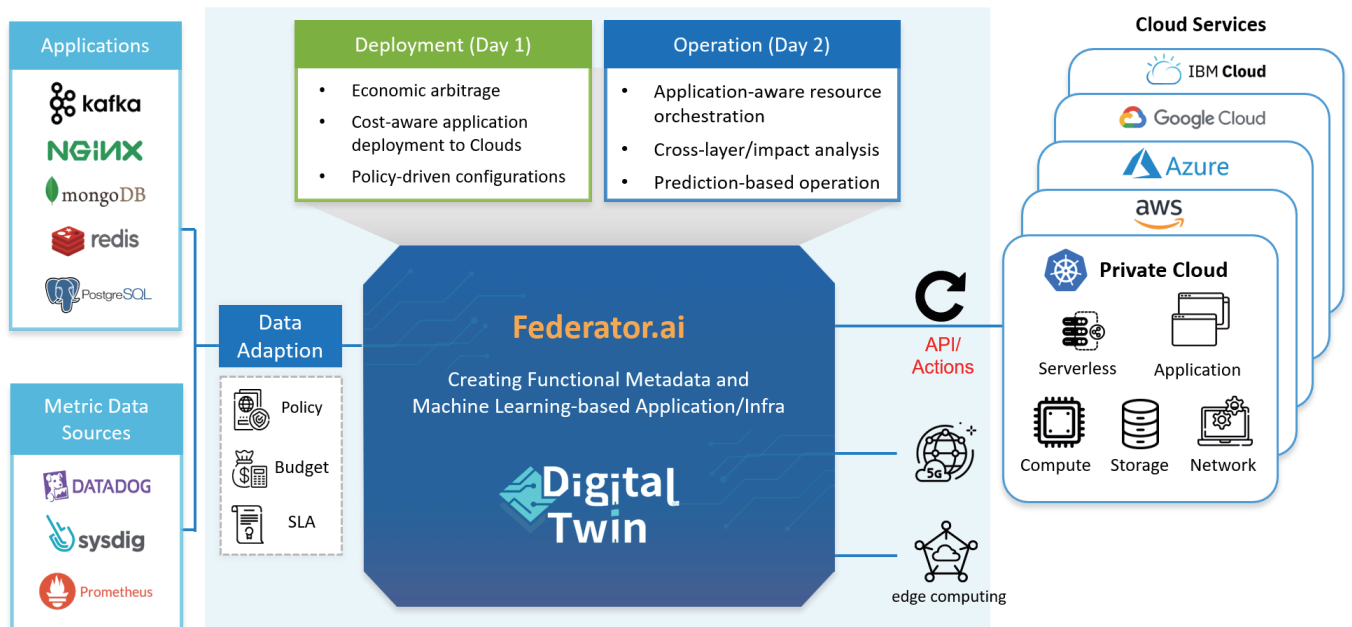
Challenges

The digital transformation journey driven by containerized applications orchestrated by Kubernetes brings agility and efficiency. However, it also brings unique challenges whether container adoption occurs on-premises, in public clouds, or both. IT administrators/architects need significant efforts in handling the increased complexity, confusing and diverse application KPIs, dynamic workloads, and many knobs to turn to fulfill the optimization objective. Often, the issues are compounded exponentially as the number of applications to manage increases, and the CIOs/CFOs are adopting the MultiCloud strategy for security, cost, and resilience purposes. Managing computing resources and optimizing costs on multiple clouds is a daunting task. Achieving automation and optimizing the applications' cost and performance are seemingly conflicting objectives. Afterthoughts are the existing application performance management, operation monitoring, and cost-savings solutions. These monitoring solutions can help relieve the pain after the problem surfaces. However, they need to have a predictive and proactive framework that can be a copilot helping them navigate and proactively plan and resolve issues before problems surface. The solution should be integrated with various operation data sources and can leverage the metadata to drive the optimization of the operation.

The Solution - Federator.ai

Federator.ai, ProphetStor's Artificial Intelligence for IT Operations (AIOps) platform, considers Cloud operations' cost and resilience. It takes the operation metadata, and understands the applications' dynamics, KPIs, and how they use the resources in the layers below. The patented analytic engine, which performs continuous but computationally feasible correlation and impact analysis, provides intelligence to orchestrate container resources on top of VMs (virtual machines) or bare metal, allowing users to operate applications without the initial complexity. Federator.ai enables micro-managing machines to optimize the resilience and cost of operations

Federator.ai addresses these problems by orchestrating resources in MultiCloud environments. The figure below shows that Federator.ai optimizes costs for both Day-1 deployment and Day-2 operations. Utilizing metrics from monitoring solutions/services, such as Datadog, Sysdig, and Prometheus, Federator.ai predicts resource consumption dynamically and recommends the right amount of resources for pods, reduces wasted resources for a typical workload, and prevents under-provisioning of resources for mission-critical workloads. Users can stack up the predicted pod resources to determine the right number and size of VMs to deploy and enable the automatic execution of these recommendations.



Key Features

After Federator.ai is deployed in any Kubernetes, OpenShift, or VMware Tanzu environment, it learns application resource usage patterns using the operation metadata. It predicts the full-stack resource usage down to the container level. Federator.ai also performs multi-layer correlation and impact analysis and provides deep insights and recommendations for resources at different layers in a Kubernetes cluster (or a VM cluster). Using Federator.ai's APIs allows automation to optimize the performance and cost of applications in a MultiCloud environment.

Multi-layer workload prediction

Federator.ai applies patented machine learning analytics to predict resource usage for applications at different levels: clusters, nodes, namespaces, applications, and controllers. The usage predictions are the basis for resource recommendations for these objects. Federator.ai supports both physical and virtual CPU, memory, and network resources.

Application-aware recommendations/execution of autoscaling/auto-provisioning

The predicted application resource demand determines the number and size of containers. Federator.ai utilizes resource usage prediction based on workload patterns to recommend the Just-in-Time Fitted containers resources, automatically scaling the containers to handle the demands.

Application correlation and impact analysis

Federator.ai provides analysis and system recommendations based on the correlation between microservices of an application, providing insights into how individual microservices are impacted by external factors that affect the primary application workload.

This correlation is crucial for preventing over-or under-provisioning of resources. In addition, customers can use the application correlation intelligence to reverse charge applications and organizations for the resource's usages.

Intelligent cost management

Federator.ai provides analysis of cost efficiency and cost trends for clusters, cluster nodes, namespaces, and applications based on expected workload. Federator.ai makes predictions and recommendations for planning and cost optimization using this information.

Policy-driven planning of CPU, memory, and networks

Federator.ai plans cluster-wide CPU and memory allocation for different types of applications according to the user-defined policy.

Continuous recommendations for optimal resource planning

Federator.ai continuously generates recommendations and learns better with more accumulated metrics data.

MultiCloud and Multi-Instances considered

Federator.ai continuously monitors the cost structures of the top cloud service providers, understands the application workloads and their nature to be placed in different instance types (On-Demand, Reserved, and Spot), considers the costs of different data centers of cloud services providers, and understands the cloud bills and service provider's reports to further optimize the cost of operation.

Enterprise-ready

Federator.ai works seamlessly with all distributions of Kubernetes, SUSE Rancher, and Red Hat OpenShift-operated environments, providing application lifecycle management based on the operator frameworks.

Benefits

Federator.ai aims to provide optimal resource planning recommendations to help enterprises make better decisions. The benefits of Federator.ai include:

Up to 80% resource savings

Federator.ai mainly reduces unnecessary spending and increases application service quality for enterprises and cloud providers. ProphetStor data scientists and engineering teams work together to build the most advanced AIOps solution to reduce resource wastage at different infrastructure layers. With the help of patented data analytic engines, Federator.ai simultaneously reduces spending and delivers necessary performance.

Increased operational efficiency

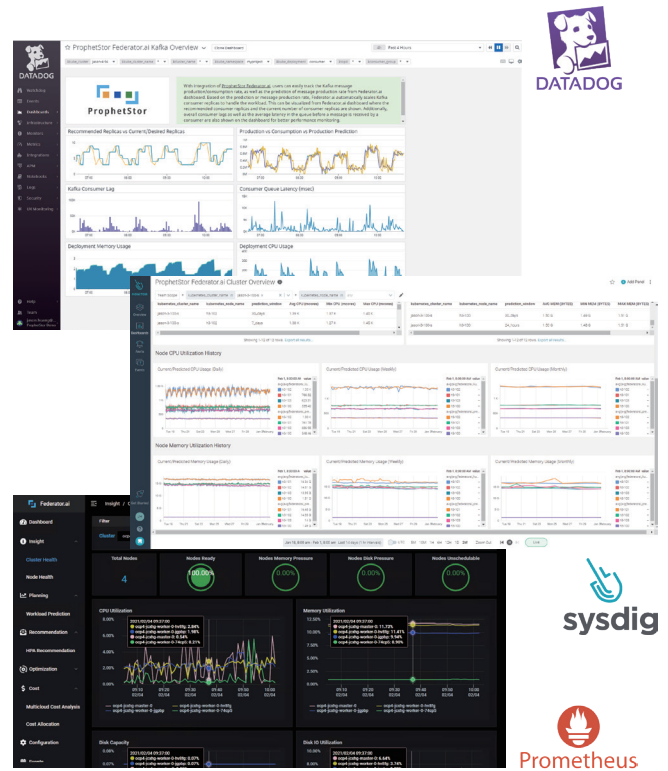
Federator.ai frees users from continuously monitoring Kubernetes / OpenShift /VMware Tanzu cluster or VM cluster utilization and cloud spending. Users also do not need to manually record usage data, calculate optimal configurations, and change configurations based on the calculations. These tasks are routinely accomplished when using Federator.ai.

Reduced manual configuration time with digital intelligence

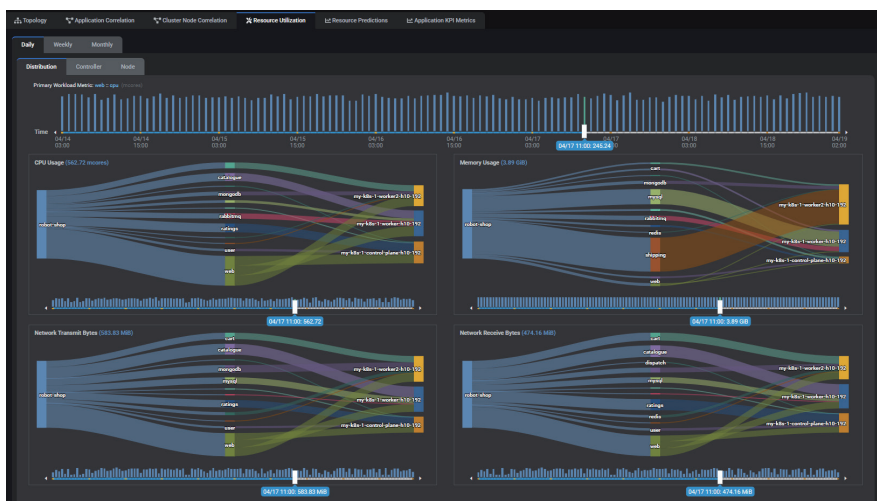
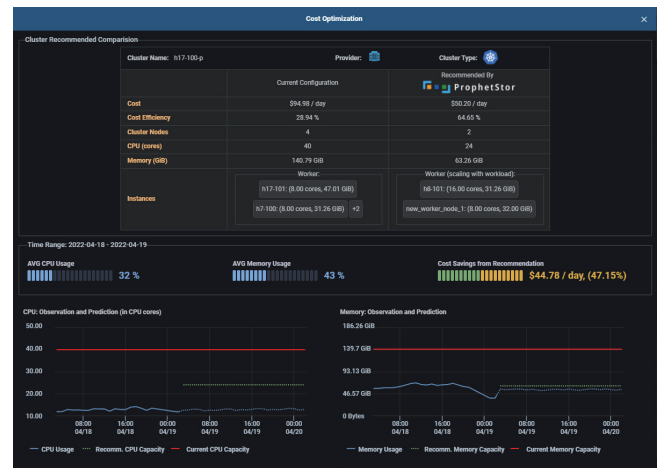
Federator.ai allows users to turn on the optimization engine and let the system do the heavy lifting, automatically provisioning resources when needed. Using the Federator.ai open API, users can re-configure pods with the right values at the right time.

Multi-Layer Correlation and Impact Analysis Facilitates Resource Planning and Cross-Layer Automatic Scaling

Fully integrated with leading cloud-monitoring services



Cost Optimization Recommendation with Dynamic Application Workloads, Workload Types, Supporting Cloud Instances, and Cloud Service Provider Cost Structures



Feature Details and Specifications

<p>AI-based multi-layer workload predictions</p>	<ul style="list-style-type: none"> • Continuous workload predictions for multi-layer Kubernetes resources: clusters, nodes, namespaces, applications, and controllers • Continuous workload predictions for individual VMs and VM clusters • Generate daily, weekly, and monthly workload predictions • Obtain near-immediate prediction results when historical metric data is available
<p>Intelligent recommendations for resource planning</p>	<ul style="list-style-type: none"> • Rightsize continuously CPU/memory resource recommendations for Kubernetes clusters, nodes, namespaces, applications, and controllers to reduce waste without compromising performance • Extend resource usage predictions and recommendations for VM nodes and VM clusters • Support application-specific metrics for controllers of Kubernetes applications, such as MongoDB, MySQL, NGINX, RabbitMQ, and Redis
<p>Application correlation and impact analysis</p>	<ul style="list-style-type: none"> • Analysis of operational metrics correlation between microservices and the application • Provide insights about how individual microservices are impacted by external factors that affect the primary workload • Intelligent system recommendations based on application workload and the impact on each microservice • Full-stack visibility of application resource utilization and performance
<p>Proactive and application-aware autoscaling</p>	<ul style="list-style-type: none"> • Intelligent autoscaling of containers based on workload predictions • Autoscale containers based on application-specific workload metrics and KPIs • Autoscale NGINX ingress upstream services • Autoscale Kafka consumers • Auto-generate KEDA config YAML files for fast integration
<p>Auto-provisioning application resources</p>	<ul style="list-style-type: none"> • User-defined provisioning policy for automatic adjusting resources allocation for application container and namespaces • Additional headroom adjustments and min/max resource allocation settings • Flexible scheduling options • Downloadable resource provisioning scripts for easy customization and 3rd party integration
<p>CI/CD integration</p>	<ul style="list-style-type: none"> • Simple integration to CI/CD pipeline for right-sized container deployments • Sample script for easy Terraform deployment integration
<p>Intelligent cost management</p>	<ul style="list-style-type: none"> • Machine-learning-based cost optimization and recommendations for clusters, nodes, namespaces, and applications • Predictive analytics for cost trends of clusters, cluster nodes, namespaces, and applications based on expected workload • Analysis of cost efficiency for better planning and optimization

Multicloud cost analysis	<ul style="list-style-type: none"> • Recommend the most cost-effective cluster configuration for AWS, Google, Azure, and on-premises • Recommendations based on On-Demand, Reserved, Spot, and Spot+Reserved instances • Workload prediction-based recommendations for individual VM instances • Custom price book for cost calculation of on-premises deployments • Automatic update of price books from public cloud service
Multiple metrics data sources	<ul style="list-style-type: none"> • Support metrics collected from Prometheus, Datadog, and Sysdig monitoring services for Kubernetes clusters, as well as VMware/vCenter, AWS Cloud-Watch, Azure Monitor, and Google Cloud's Operations Suite for VM clusters
Alert Management	<ul style="list-style-type: none"> • Custom-defined monitors on resource shortage or overage based on resource predictions • Auto-alert cancellation when the condition recovers • Automatic email notification of new alerts
Autodiscovery of cluster resources	<ul style="list-style-type: none"> • Automatic discovery of cluster nodes, namespaces, containers/pods in namespaces • Auto-classification of common containers/pods
Config DB backup and restore	<ul style="list-style-type: none"> • Auto backup of configuration during the upgrade procedure, and auto restore backup config during downgrade • Backup/restore the configuration DB to/from an external system • Backup DB is encrypted and password protected
Open REST API	<ul style="list-style-type: none"> • Open REST API for resource predictions and recommendations • Open REST API for cost management and cluster configuration recommendations
Setup wizard for easy installation	<ul style="list-style-type: none"> • Easy installation through Operator framework on all Kubernetes distributions, RedHat OpenShift, SUSE Rancher, and VMware Tanzu • Support Ansible Playbook installation • Support installation by Helm charts
Usage-based licensing	<ul style="list-style-type: none"> • License based on number of resource objects monitored (cluster nodes, namespaces, and controllers) • Free software usage for monitoring up to 10 resource objects
Easy-to-use UI	<ul style="list-style-type: none"> • Visualization of resource usages and predictions for multi-layer of Kubernetes resources • Support monitoring for multiple clusters • User-defined application with controllers from multiple namespaces

<p>Integration with third-party monitoring services</p>	<ul style="list-style-type: none"> • Single pane of glass management from Datadog and Sysdig monitoring portal with custom dashboards • Preconfigured Datadog Monitor for under-provisioned resource alerts
<p>Supported platforms</p>	<ul style="list-style-type: none"> • Kubernetes v1.11x – v1.22x • Red Hat OpenShift v4.6-v4.9 • Amazon EKS • Google GKE • Microsoft AKS • SUSE/Rancher v2.4.8, v2.5.8, v2.5.9, v2.6.3 • VMware vCenter 5.5, 6.0, 6.5, 6.7, 7.0 • VMware Tanzu Kubernetes Grid v1.6.0 • IBM Cloud/IKS • AliCloud/ACK

ProphetStor Data Services, Inc.

Headquarters

830 Hillview Court, Suite 100
Milpitas, CA 95035
+1 408 508 6255
www.prophetstor.com

Paris Office

2 place de Touraine
78000 Versailles
France
+33 1 7029 0866

Tokyo Office

7F, Wakamatsu Bldg., 3-3-6
Nihonbashihoncho, Chuo-Ku
Tokyo 103-0023, Japan
+81 3 3249 6378

Taipei Office

16F, No. 182, Sec. 2, Dunhua S. Rd.
Da'an Dist., Taipei City
Taiwan 10669
+886 2 8219 2814

Taichung Office

13F, No. 219, Minquan Rd.
West Dist., Taichung City
Taiwan 40341
+886 4 2305 1816



Visit us at www.prophetstor.com to find out more, email us at info@prophetstor.com or contact your local ProphetStor office.

Copyright © 2017-2024 ProphetStor Data Services, Inc. All rights reserved. ProphetStor Data Services and Federator.ai are trademarks or registered trademarks of ProphetStor Data Services, Inc. in the USA and other countries. All other company and product names contained herein are or may be trademarks of the respective holders.