Datasheet

Federator.ai®

Intelligence-Driven Cloud Operation Automation and Optimization

ProphetStor

90%	80%	59%
Resource Savings	Streamlined Operations	Carbon Emission Reduction

- **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 delivers agility and efficiency. However, it also introduces new operational challenges across on-premises and multi-cloud deployments.

IT administrators and architects must manage increasing system complexity, diverse and unclear application KPIs, highly dynamic workloads, and numerous configuration variables required to meet optimization goals. As application scale grows and MultiClod strategies expand for cost, security, and resilience, these challenges increase significantly.

To move forward, organizations must address several key challenges:

- Increasing complexity across hybrid and multi-cloud infrastructures
- Lack of clear and unified application performance indicators
- Difficulty balancing cost optimization with performance requirements
- Limited ability to predict and proactively prevent issues

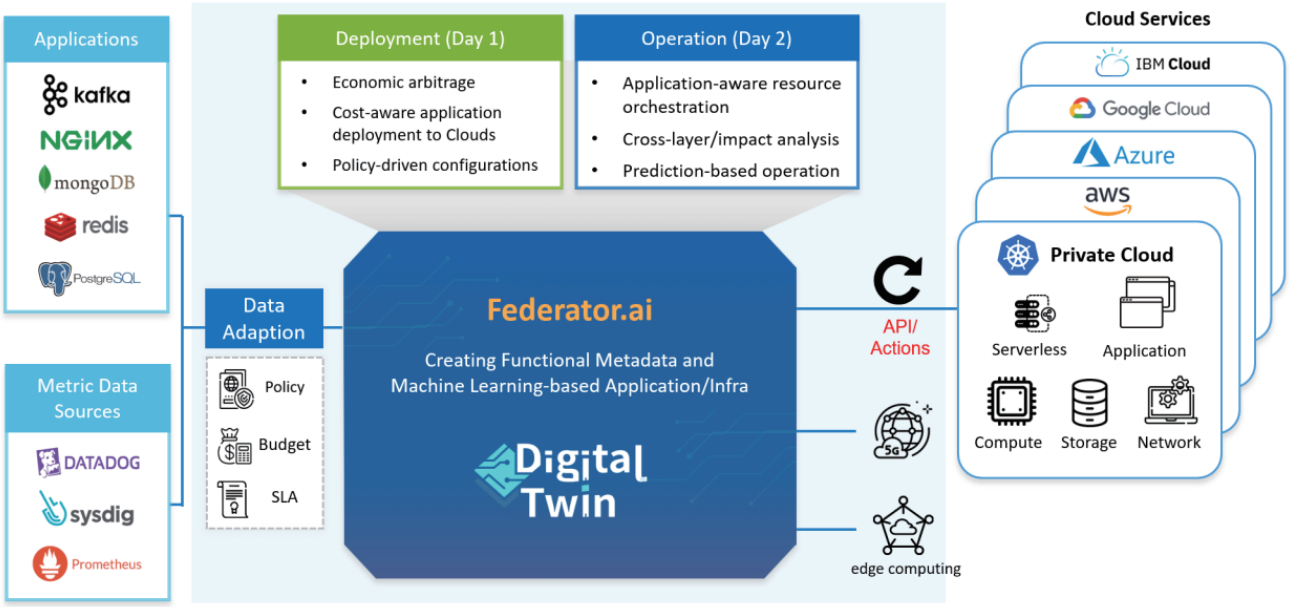
Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

The Solution - Federator.ai®

Federator.ai, ProphetStor's Artificial Intelligence for IT Operations (AIOps) platform, optimizes cloud operations for both cost and resilience. It ingests operational metadata to understand application behavior, KPIs, and how applications consume resources across underlying layers.

Its patented analytics engine performs continuous yet computationally efficient correlation and impact analysis, providing the intelligence needed to orchestrate container resources on top of virtual machines (VMs) or bare metal. This enables users to operate applications without dealing with initial infrastructure complexity. Federator.ai also enables fine-grained resource management to optimize both operational resilience and cost.

Federator.ai addresses these challenges by orchestrating resources across MultiCloud environments. As illustrated in the figure below, it optimizes both Day-1 deployment and Day-2 operations. By leveraging metrics from monitoring solutions such as Datadog, Sysdig, and Prometheus, Federator.ai dynamically predicts resource consumption and recommends the optimal allocation for pods. It reduces resource waste for typical workloads while preventing under-provisioning for mission-critical applications. Users can aggregate predicted pod requirements to determine the appropriate number and size of VMs, and optionally automate the execution of these recommendations.



Federator.ai: An AI-driven digital twin platform for application-aware, cost-optimized orchestration across MultiCloud

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Key Features

After Federator.ai is deployed in any Kubernetes, OpenShift, or VMware Tanzu environment, it learns application resource usage patterns from operational metadata. It predicts full-stack resource consumption down to the container level.

Federator.ai also performs multi-layer correlation and impact analysis, providing deep insights and resource recommendations across different layers of a Kubernetes cluster (or VM cluster). Through its APIs, it enables automation to optimize application performance and cost in multi-cloud environments.

- **Multi-layer workload prediction:** Federator.ai applies patented machine learning analytics to predict resource usage across multiple levels, including clusters, nodes, namespaces, applications, and controllers. These predictions serve as the foundation for resource recommendations at each level. Federator.ai supports both physical and virtual CPU, memory, and network resources.
- **Application-aware autoscaling and auto-provisioning:** Predicted application resource demand determines the number and size of containers. Federator.ai uses workload-based resource usage predictions to recommend Just-in-Time Fitted container resources and automatically scale containers to meet demand.
- **Application correlation and impact analysis:** Federator.ai analyzes correlations between application microservices and provides system-level recommendations. It offers insights into how individual microservices are affected by external factors impacting the overall workload, helping prevent over- or under-provisioning. In addition, this correlation intelligence can support chargeback across applications and organizations based on resource usage.
- **Intelligent cost management:** Federator.ai analyzes cost efficiency and trends across clusters, nodes, namespaces, and applications based on expected workloads. It generates predictions and recommendations for planning and cost optimization.
- **Policy-driven planning of CPU, memory, and networks:** Federator.ai plans cluster-wide CPU, memory, and network allocation for different application types based on user-defined policies.
- **Continuous recommendations for optimal resource planning:** Federator.ai continuously generates recommendations and improves accuracy as more operational data is collected.
- **MultiCloud and Muti-Instances considered:** Federator.ai monitors cost structures across major cloud providers, understands workload characteristics, and recommends optimal placement across instance types (On-Demand, Reserved, Spot). It also considers regional pricing differences and analyzes cloud billing data to further optimize operational costs.

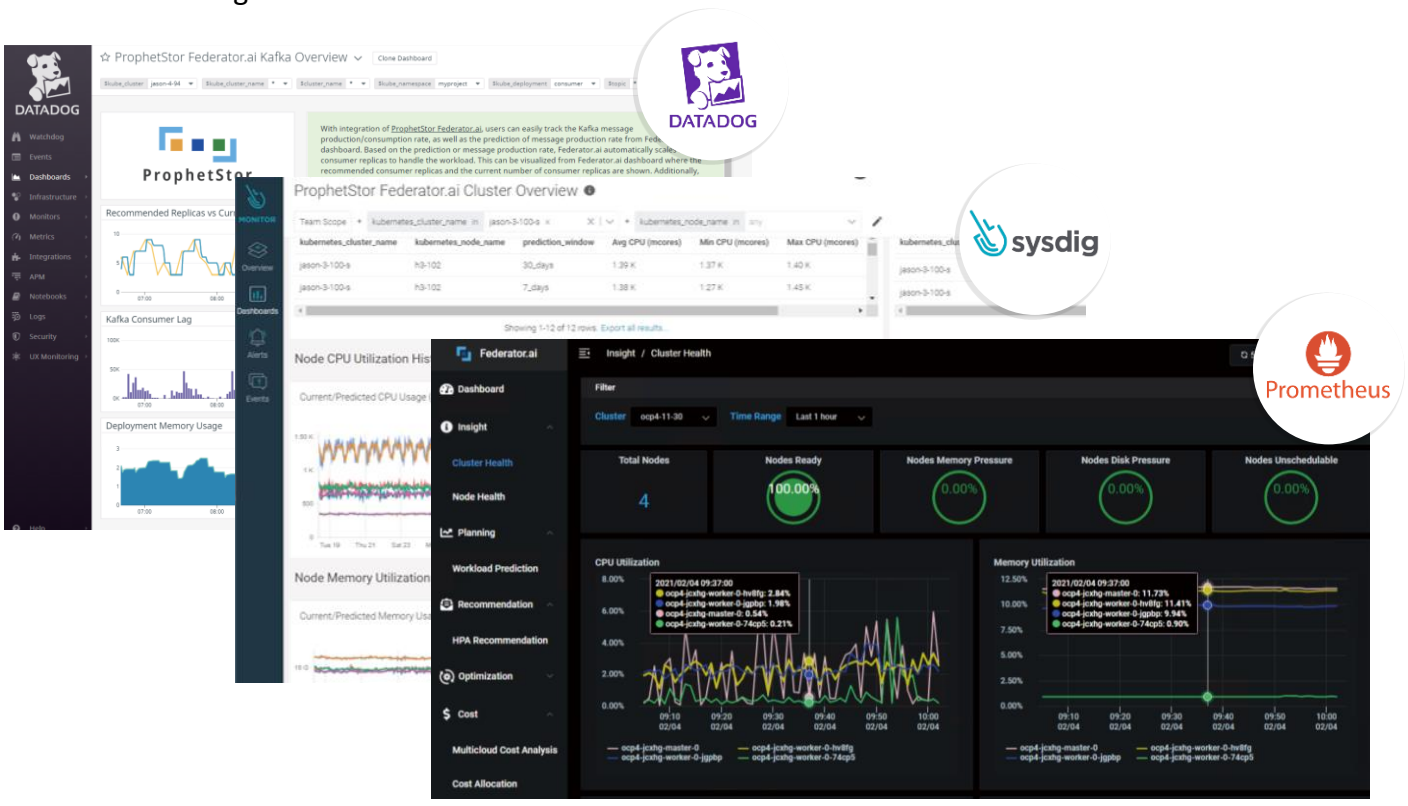
Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

- **Enterprise-ready:** Federator.ai integrates with all Kubernetes distributions, including SUSE Rancher and Red Hat OpenShift environments, and supports application lifecycle management through operator frameworks.

Benefits

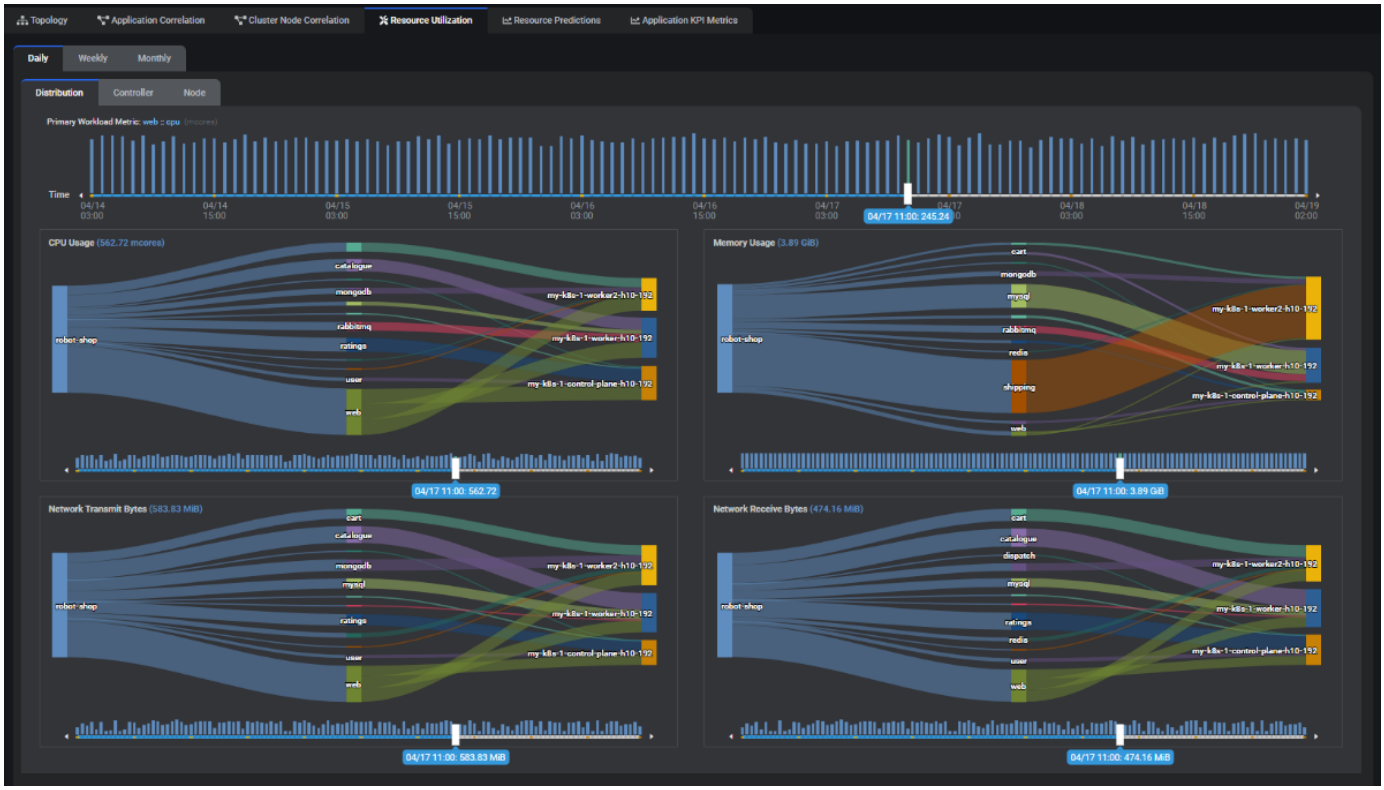
Federator.ai provides optimal resource planning recommendations to help enterprises make better operational and cost decisions. Key benefits include:

- **Up to 80% resource savings:** Federator.ai reduces unnecessary spending while improving application service quality. By leveraging patented analytics, it minimizes resource waste across infrastructure layers while maintaining required performance levels.
- **Increased operational efficiency:** Federator.ai eliminates the need for continuous manual monitoring of Kubernetes, OpenShift, VMware Tanzu, or VM clusters and cloud spending. It automates data collection, analysis, and configuration adjustments, reducing operational overhead.
- **Reduced manual configuration with intelligent automation:** Federator.ai enables users to activate its optimization engine to automatically provision and adjust resources as needed. Through its open APIs, users can dynamically configure pods with the right resource settings at the right time.



Federator.ai Collects Metadata from Leading Monitoring Tools to Optimize Performance and Resource Efficiency

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com



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



Cost optimization recommendations based on dynamic workloads, cloud instances, and provider cost structures

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Feature Details and Specifications

Feature Area	Key Capabilities/ Specifications
AI-based multi-layer workload predictions	• 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
Intelligent recommendations for resource planning	• Continuously rightsize CPU and memory resources across Kubernetes clusters, nodes, namespaces, applications, and controllers to reduce waste without compromising performance • Extend resource prediction and recommendations to VM nodes and VM clusters • Support application-specific metrics for Kubernetes controllers, including MongoDB, MySQL, NGINX, RabbitMQ, and Redis
Application correlation and impact analysis	• 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
Proactive and application-aware autoscaling	• 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

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Feature Area	Key Capabilities/ Specifications
Auto-provisioning application resources	• 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
CI/CD integration	• Simple integration to CI/CD pipeline for right-sized container deployments • Sample script for easy Terraform deployment integration
Intelligent cost management	• 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	• 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	• Support metrics collected from Prometheus, Datadog, and Sysdig monitoring services for Kubernetes clusters, as well as VMware/vCenter, AWS CloudWatch, Azure Monitor, and Google Cloud's Operations Suite for VM clusters
Alert Management	• 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

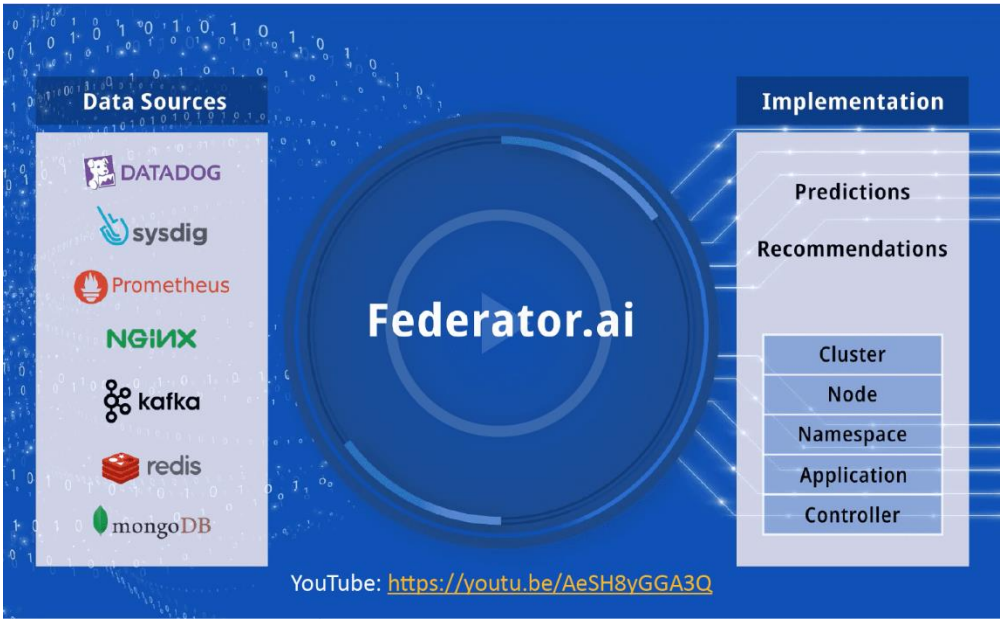
Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Feature Area	Key Capabilities/ Specifications
Auto-discovery of cluster resources	• Automatic discovery of cluster nodes, namespaces, containers/pods in namespaces • Auto-classification of common containers/pods
Config DB backup and restore	• 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	• Open REST API for resource predictions and recommendations • Open REST API for cost management and cluster configuration recommendations
Setup wizard for easy installation	• 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	• 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	• 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
Integration with third-party monitoring services	• Single pane of glass management from Datadog and Sysdig monitoring portal with custom dashboards • Preconfigured Datadog Monitor for under-provisioned resource alerts

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Feature Area	Key Capabilities/ Specifications
Supported platforms	• 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

Video | Federator.ai Feature Demo



Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com