

# VMware vRealize AI Cloud

An intelligent, self-tuning service that adapts to the changing needs of your dynamic workloads to optimize infrastructure performance.

## AT A GLANCE

vRealize AI Cloud is a secure SaaS service that uses Reinforcement Learning to continuously learn, adapt, and optimize your infrastructure KPIs to deliver consistent and optimal performance for your application workloads

## SELF-TUNING:

- Data Collection – Collects, stores and analyzes vSAN and vSphere metrics in the SaaS data lake
- Continuous Monitoring – Learns resource utilization, demands, and patterns with rapid sampling of vSAN and vSphere metrics
- Automated Actions – Dynamically adjusts vSAN and vSphere tunables to the changing to needs of application requirements

## STORAGE POLICY GENIE:

- Analyzes infrastructure and provides recommendations for improving performance or capacity and factors and cost of changes on a net benefit score

## BENEFITS:

- Improved infrastructure performance – up to 60%
- Reduced monitoring effort – up to 15 min per day
- Reduced support tickets/escalations – up to 5 tickets per month

## HOW TO BUY:

- vRealize AI Cloud can be enabled through any version of the vRealize Cloud Universal subscription via vRealize Operations Cloud;
- Also now available on a self-explainable stand-alone UI via vRealize Cloud Universal subscription

## Overview

Today's IT datacenters are complex and dynamic. IT teams continuously add new technologies while having to maintain legacy deployments. But no matter the environment, the goal is clear - maintaining optimal performance and constant balancing of system resources. Workloads are deployed on physical servers, virtual machines, containers, Kubernetes, or cloud-native applications. These workloads constantly change, move and share resources within the same hosts. Clusters that may have been initially configured to support legacy applications are now hosting modern workloads which no longer follow best practice considerations. These different deployments create IT complexity, which can hurt performance, efficiency and drift from SLA requirements.

No matter where companies host workloads - on-prem, private, hybrid, or multi-cloud environments - they need optimal performance across any platform to meet customer requirements and SLA agreements.

## Why vRealize AI Cloud?

VMware vRealize AI Cloud is an intelligent, self-tuning service that uses reinforcement learning to continuously adapt to the changing needs of your application workloads. Through data collection and machine learning, it analyzes performance KPI's to not only predict performance improvements but also dynamically self-tunes VMware SDDC and vSAN parameters to give you the best performance results.

vRealize AI Cloud uses machine learning and reinforcement learning techniques to intelligently and continuously optimize VMware infrastructure operations, starting with the first release of vRealize AI Cloud to optimize vSAN performance. The service learns about your operating environment and adapts to changing dynamics, ensuring optimization per the stated KPI.

SaaS datalake and on-prem data collections enable real-time and historical observability. IT teams can reduce or eliminate repetitive manual processes with guardrails in place to limit negative impacted performance. And companies improve application workload performance by evaluating all possible VMware SDDC tunables to select the optimal configuration.



Optimize Performance



Cost Efficient



Optimize Capacity

The latest version includes a preview of the Storage Genie Policy engine that analyzes and predicts performance, capacity and evaluates the complexity (cost) of that change to trigger recommendations based on a net performance or capacity improvements benefit score. The recommendations suggested by the Storage Genie Policy engine help:

- Suggest changes to default storage policy
- Review multiple alternatives with differing trade-offs and explanations

vRealize AI Cloud

Use Case: Continuous vSAN Optimization



Self-explainable UI

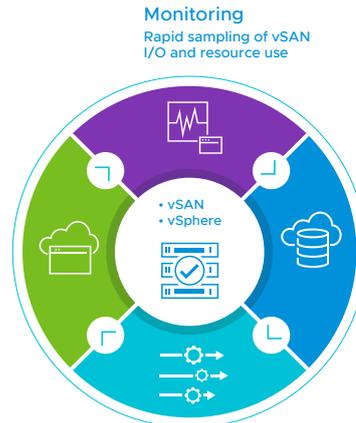
- Turn-on Service
- View performance improvements

Continuous self-tuning



vRealize AI Cloud

- Analyzes performance and determines actions
- Continuous learning based on results



Actions  
Dynamic changes to vSAN, vSphere tunables

Self-tuning benefits of vRealize AI Cloud

vRealize AI Cloud's is a performance optimization engine that uses machine learning and reinforcement learning techniques to continuously learn, adapt, and optimize your storage KPIs. This helps deliver consistent and optimal performance for your mixed application workloads. As workload applications scale out or migrate to different datacenters or clusters, vRealize AI Cloud will dynamically adjust vSAN tunables to continuously optimize KPIs such as I/O read and write throughput or reducing network latency.

Other 3rd party analytic solutions do not collect and analyze the data, but instead take decisions on mere utilization and available resources to optimize performance. vRealize AI Cloud takes a proactive approach to optimize infrastructure and application performance with its self-tuning of VMware SDDC tunables.

Continuous Value with vRealize AI Cloud



Continuous Monitoring

Move managing environments from incident based to SLO management



Continuous Optimization

Ensure the Cost, Availability, Performance & Security are being continually managed



Continually refine the policies

Refine the policies and desired state by which it is governed as we continuously learn the environment



Continually adapt the infrastructure

App KPIs and SLOs drive the configuration of the infrastructure