Lower TCO for containers, AI, and VMs

with Red Hat OpenShift on Google Cloud and Google Cloud NetApp Volumes

Google Cloud | Red Hat | NetApp

Red Hat OpenShift

Application modernization

Powers the entire application lifecycle to help you modernize your systems and reach your goals.

Hybrid platform

Delivers open-source innovation in a single, seamless platform across your on-prem and cloud environments with Kubernetes at its core.

Enterprise-grade

Trusted by industry leaders, Red Hat OpenShift offers security features paired with dedicated support, freeing your teams to focus on the innovative work that organizations, users, and customers demand.

Google Cloud

Optimized infrastructure

Purpose-built optimizations for OpenShift workloads to deliver low TCO, best reliability, and security.

AI transformation

Transform your business with the industry-leading AI full-stack platform.

NetApp

Enterprise-grade storage

Fully managed, robust, scalable, and secure storage.

Hybrid platform and simple migration

Operational consistency and seamless data mobility between on-prem, self-managed, and fully managed.

High performance, low cost

No trade-off between price, performance and management.

Outline

OpenShift workloads on Google Cloud's robust and complete platform

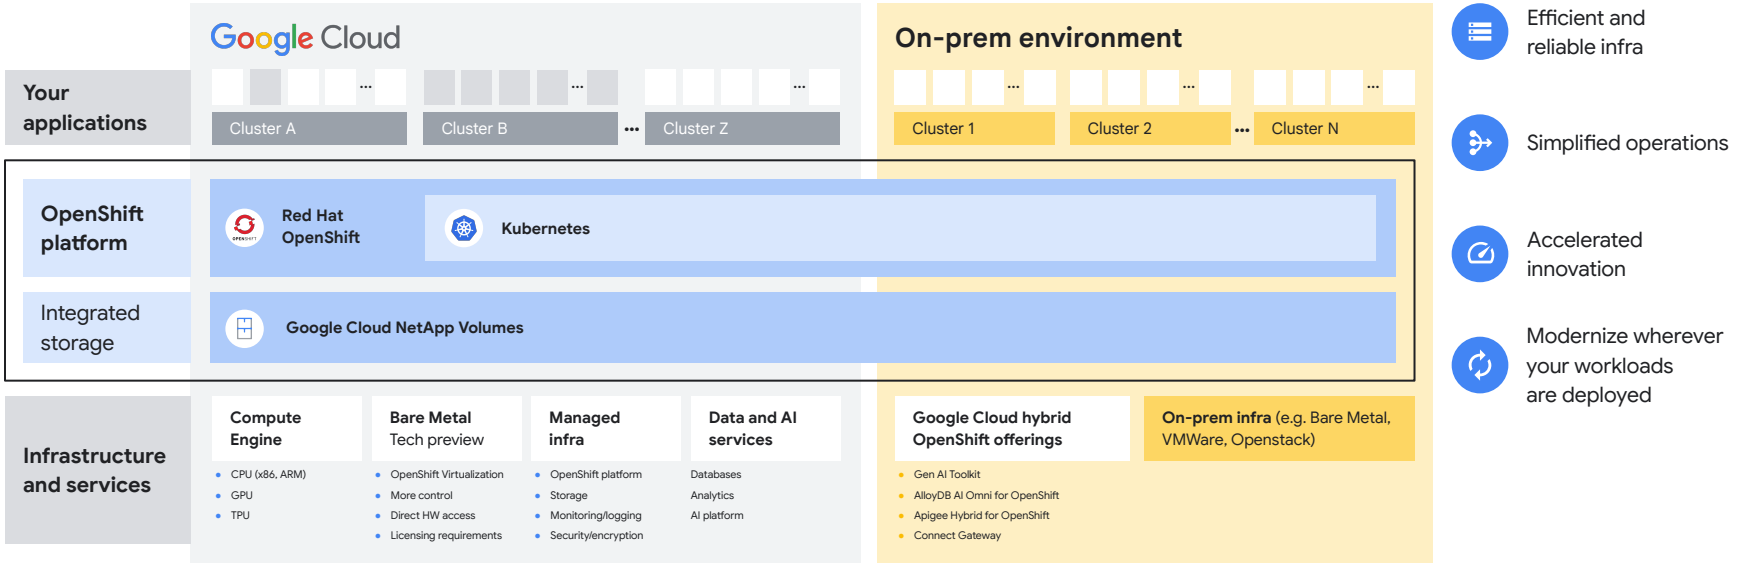

OpenShift on Google Cloud differentiators

Reduce TCO

Operating **OpenShift on-prem is expensive** as infra costs are difficult to optimize.

50%+ infra cost savings are typical when migrating OpenShift workloads to the cloud.

Google Cloud-native technologies enable **25–40% greater infra cost savings than AWS/Azure.**

Sleep well at night

Most **reliable** and **secure** cloud.

Kubernetes leader: #1 contributor, 85% of Kubernetes tech leads → **unmatched L3/L4 support for your production workloads.**

Seamlessly migrate

Flexible across fully or self-managed OpenShift solutions **enables you to 80/20 your migration with the right approach for you.**

Broadest set of services for management of Kubernetes and OpenShift clusters.

Turbocharge with AI

OpenShift/GKE-native CRDs for all Google Cloud AI services.

OpenShift/GKE-native offerings* to **integrate OpenShift-hosted APIs/DBs/apps with agentic AI workflows.**

Hybrid AI: **The only hyperscaler providing a hybrid RAG/DB solution** — AlloyDB AI Omni for OpenShift.

Google Cloud provides 25-40% more infrastructure cost savings than AWS/Azure

Traditional enterprise applications

25%+ more cost savings than AWS/Azure

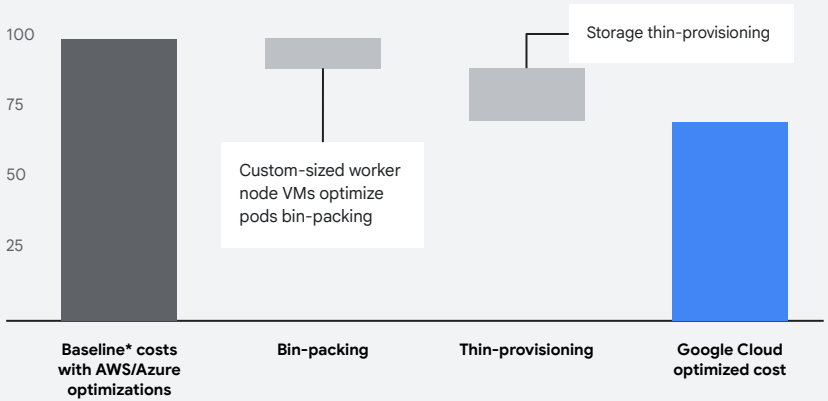

Large scale-out applications**

40%+ more cost savings than AWS***/ Azure

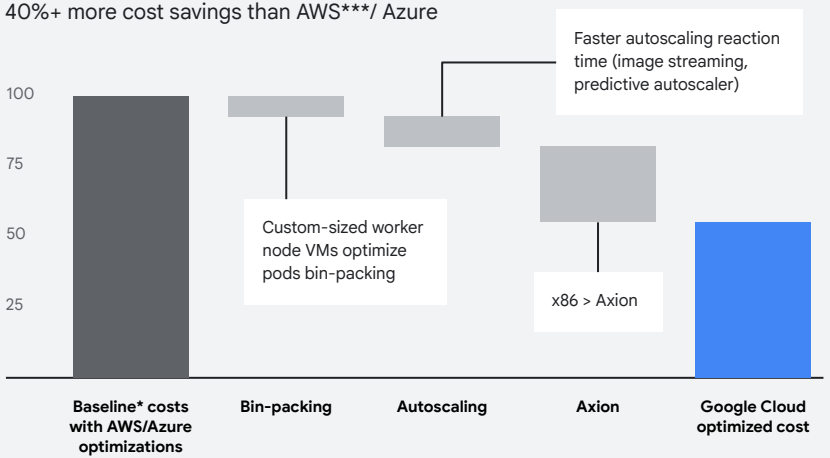

*Baseline = workloads running on AWS/Azure with all available cost optimizations applied **Such as large-scale web serving
***Cost reduction advantage vs AWS is 20% as the relevant Axion comparable is Graviton 4 which doesn't provide as good price/perf as Axion but it is still better than x86 equivalent

Optimized compute consumption via bin-packing with custom VM shapes

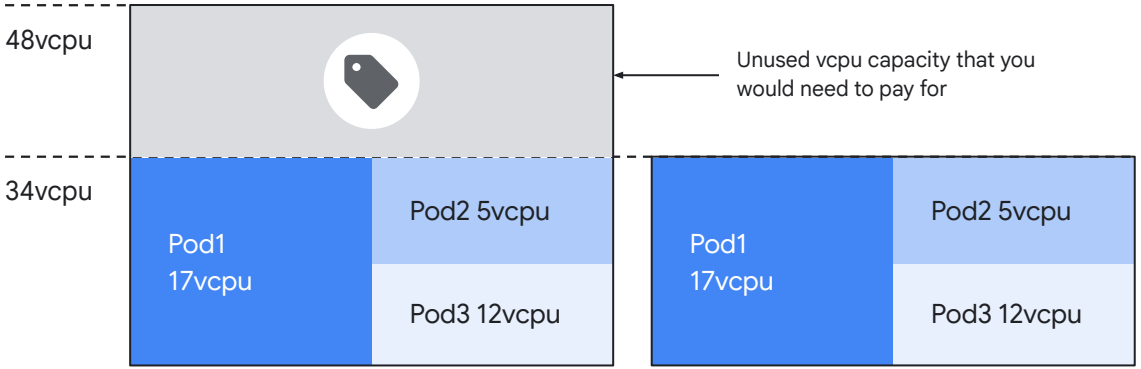

AWS/Azure

OpenShift worker node using **standard VM shapes**

Google Cloud

OpenShift worker node using **custom VM shapes**

- Standard shape VMs have predefined sizes (e.g. 4, 8, 16, 32, 48... vcpu)
- OpenShift workloads often involve odd or large-sized container images
- Google Cloud enables custom VM shapes which are not available in AWS/ Azure
- Custom VM shapes provision VMs with the right amount of vcpu and ram to optimize pods bin-packing and thereby compute spend

Enables 20% compute efficiency for typical OpenShift workloads

Cost savings of migrating OpenShift workloads to Google Cloud

Normalized costs for large FSI customers.

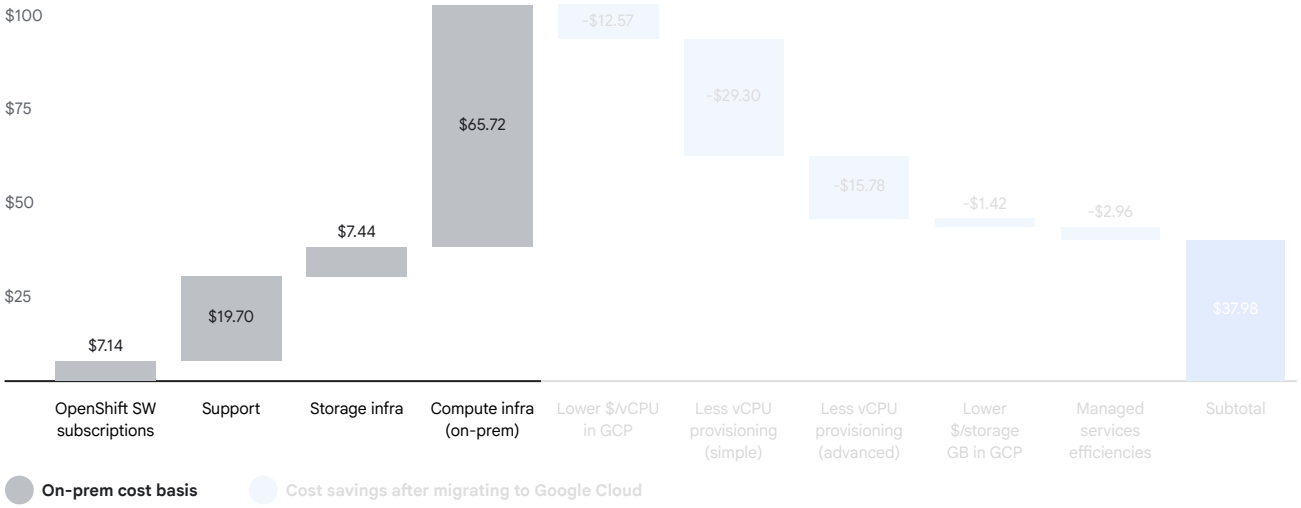

On-prem clusters

- Provisioned for peak capacity to ensure stability and performance
- Can result in lower utilization rates (8% worker nodes utilization on average)
- Costs are tied to fixed infrastructure

Cost savings of migrating OpenShift workloads to Google Cloud

Normalized costs for large FSI customers.

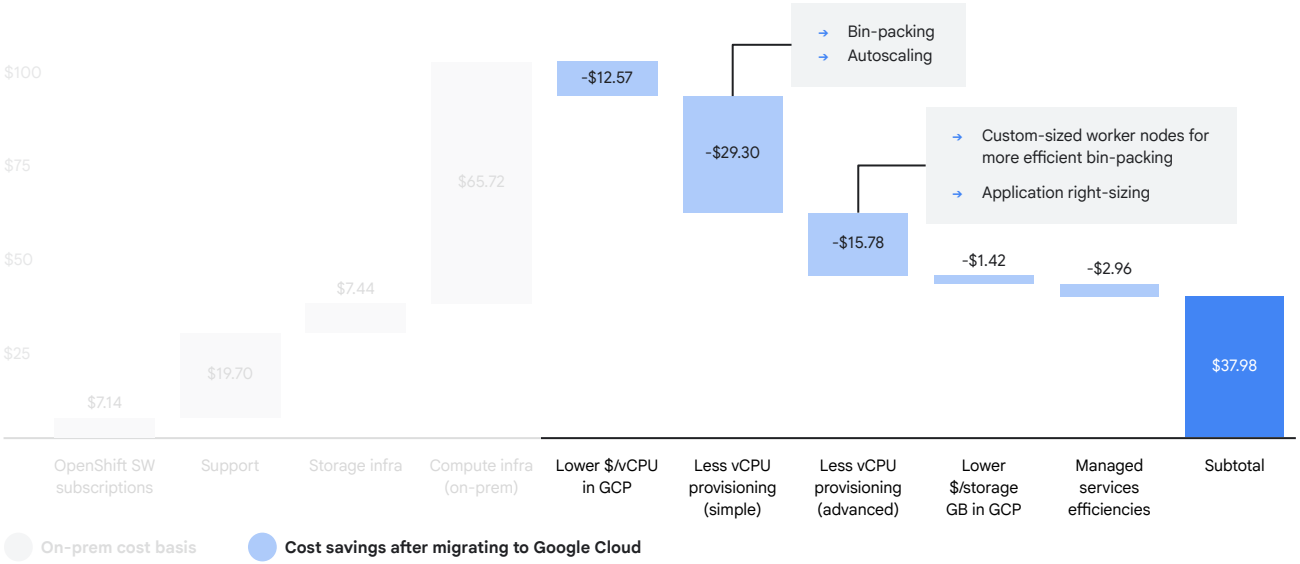

Google Cloud migration

- Lower infra unit costs
- Easy to apply relatively simple optimizations
- Access to advanced optimization features
- 50% worker nodes utilization target

~60% overall cost savings

Get a free assessment...

...of how OpenShift on Google Cloud can optimize costs while minimizing migration effort through your Google Cloud account team.

Phase	Initial BOM and cost savings opportunity assessment	Landing zone options/ migration journey optimization	Migration planning
Data request	• For all OpenShift clusters, export snapshot of metrics from all core Kubernetes System Components • Specifically in default OpenShift/OCP configurations, get snapshot of data from "/metrics" endpoints captured from each worker node's node-exporter and kubelet and each master node kube-scheduler • Absolute bare minimum data ask is this Prometheus query: **"100 - (avg by (instance) (rate (node_cpu_seconds_total {mode="idle"}[5m])) * 100)"** for all OpenShift clusters and with as much timespan as possible	• For all OpenShift clusters, export snapshot of metrics associated with Kubernetes Object States, the OpenShift Core Operators, and presence of CRDs for plug-in/add-ons such as IngressController CRDs (from K8s API) • Absolute bare minimum data ask (to get started) is this Prometheus query: **service="kube_state_metrics"** for all OpenShift clusters and with as much timespan as possible	• Comprehensive application/system/services inventory (e.g. CMDB, ITAM, ...) • Application dependencies, including with other applications and other systems or services • Costs and cost allocation information for all applications, systems, ... • Categorization along business criticality, data regulations; other requirements with regards to cloud migration and interviews w/ BU/ application teams...
Analyses	• Analysis of current level of optimization of compute infrastructure given workloads demands • Assessment of potential future compute/infra cost optimization targets on Google Cloud	• Analysis of in-OpenShift dependencies for the clusters and applications (along providers/solutions for networking, storage, observability, ...) • Estimate of migration effort for different options (keep dependencies/L&S, move to managed services, ...)	• Migration waves and prioritization • Capacity and regional planning; including architectural blueprints high level network architecture

Salling Group cut capacity time to market amid retail surges

Situation

Salling Group needed to migrate from their on-premise OKD in order to scale more rapidly during peak retail periods.

Solution

Looking for a more dynamic, robust and flexible environment, Salling Group adopted Red Hat OpenShift 4 on Google Cloud.

Impact

- 20x less downtime during critical peak periods
- Freed up five person-days each quarter
- Reduced time to market for new nodes and clusters from 45 minutes to 3 minutes

“With Red Hat OpenShift 4 on Google Cloud, we now have a stable environment that we spend almost zero time maintaining.”

Lead Engineer
Sailing Group

Learn more about
Red Hat OpenShift 4, Google Cloud

UPS delivered flexibility with Red Hat OpenShift Dedicated on Google Cloud

Situation

UPS needed a highly responsive and accurate dynamic pricing platform to provide greater flexibility for their customers.

Solution

Red Hat consulted with UPS to implement Red Hat OpenShift Dedicated on Google Cloud.

Impact

Improved performance, reliability, and speed to market, providing greater flexibility and optionality for customers — all while prioritizing security.

We were looking for a platform that could provide performance, reliability and ensure speed to market. The Red Hat team listened and, most importantly, they took action."

Software Engineering Director
UPS Technology Group

Learn more about
Red Hat OpenShift Dedicated, Google Cloud

Optimizing OpenShift storage with Google Cloud Netapp Volumes

Situation

A large FSI needed a cost-effective, native file storage solution smaller than 1TB for OpenShift workloads.

Solution

Google Cloud NetApp Volumes delivered a fully managed, first-party storage service meeting sizing and performance needs for flexible and efficient Red Hat OpenShift deployments.

Impact

Reduced costs and elevated operational efficiency significantly by precisely allocating high-performance storage to the specific application needs.

By using Google Cloud Netapp Volumes we succeeded in gaining the flexible store requirements for our Red Hat OpenShift clusters on Google Cloud while using a first-party, fully managed storage solution."

Technical leader
Large EMEA FSI Company

Outline

Google Cloud infrastructure is optimized for OpenShift 01

Google Cloud NetApp Volumes: Managed storage for OpenShift **02**

Modernize virtualization and AI with an open platform 03

Why Google Cloud NetApp Volumes?

Enterprise storage
Fast, secure, and built for OpenShift.

Move your data fast and save big
with seamless hybrid cloud agility.

Maximize operational efficiency
with fully managed, cloud-native instances.

Ready for AI
Trusted, compliant, and always protected.

Google Cloud NetApp Volumes

Simplified storage, seamless integration, and unified access.

	Red Hat OpenShift **Container Platform** GA on Google Marketplace globally	Red Hat OpenShift **Dedicated** GA on Google Marketplace globally	Red Hat OpenShift **Virtualization** Self-managed tech preview in Google Cloud
Kubernetes control plane	**Self-managed** OpenShift control plane	**Managed OpenShift** control plane (OSD)	**Self-managed OpenShift** control plane
Trident CSI Driver	Storage: **Google Cloud NetApp Volumes**		

- Fully managed storage — no infrastructure cost or maintenance — that works the same across all OpenShift configurations.
- Space efficient, low-cost and high-performance storage with simple, unified and robust data management backed by robust NetApp technology
- NetApp Trident (open-source CSI) **certified and optimized for Red Hat OpenShift** to support for seamless management of persistent storage in containerized applications
- In-built data protection and disaster recovery supported by NetApp Volumes. Optional CSI configuration using NetApp Trident Protect

Demo:

Google Cloud NetApp Volumes

Google Cloud NetApp Volumes for Openshift Dedicated

Google Cloud

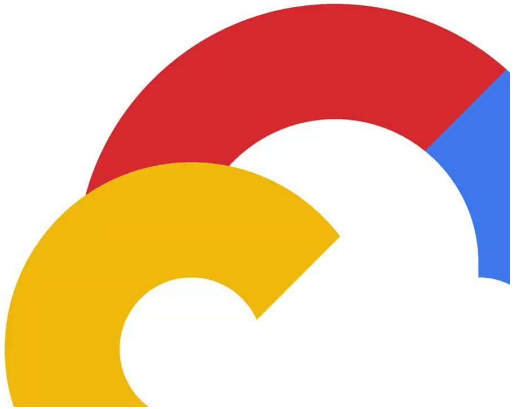

Outline

Google Cloud infrastructure is optimized for OpenShift 01

Google Cloud NetApp Volumes: Managed storage for OpenShift 02

Modernize virtualization and AI with an open platform **03**

Coming soon!

Tech preview

Red Hat OpenShift Virtualization
on Google Cloud bare metal

Coming soon

Red Hat AI Inference Server on
Google Cloud Marketplace

Efficiently run traditional virtual machines alongside containers on OpenShift on Google Cloud managed as native Kubernetes objects

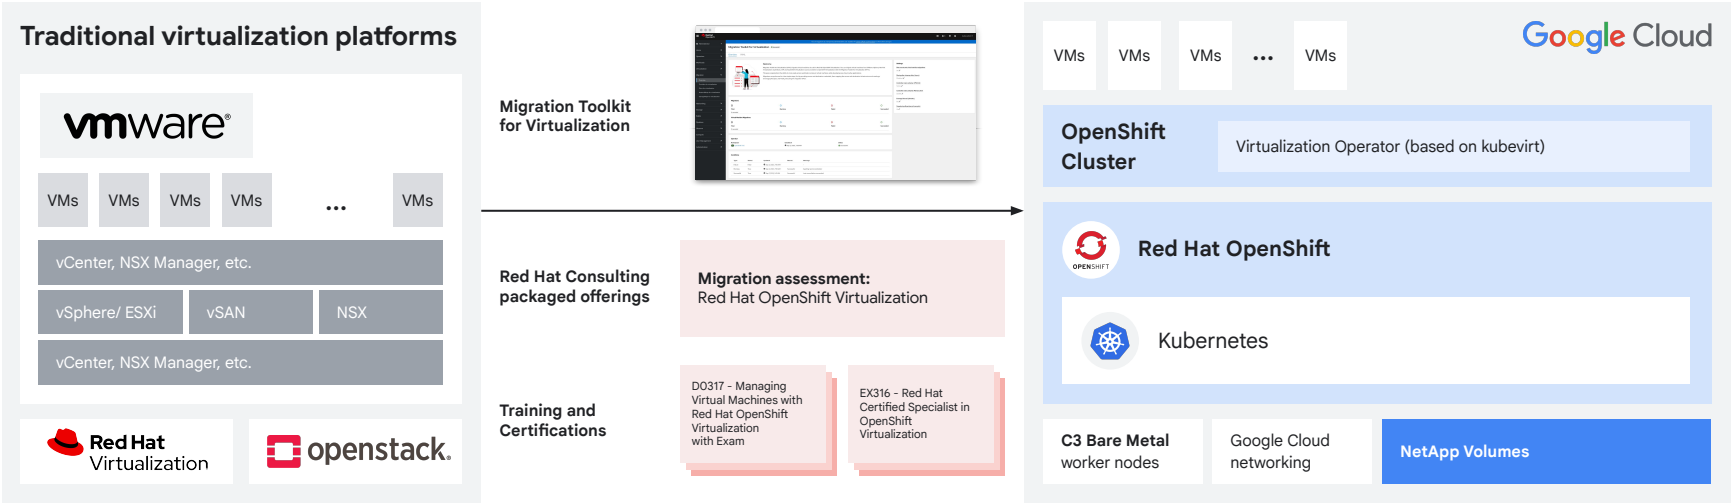

Red Hat AI Inference Server

Red Hat AI Inference Server is included in OpenShift AI and RHEL AI

Gain consistent, fast, and cost-effective inference at scale.

Optimized for Google Cloud TPUs

Leverage leading price/performance of TPUs with Day0 support for all supported models

Inference runtime (vLLM) for the hybrid cloud

Run your models of choice across any accelerator and any environment.

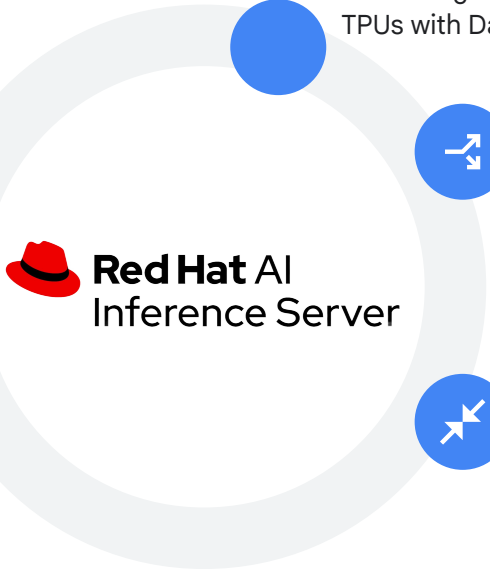

Certified for all Red Hat products

Deployable across non-Red Hat Linux and Kubernetes platforms.

Red Hat AI Hugging Face repository

Access a collection of third-party validated and optimized models ready for inference.

Compress models

Reduce compute and costs while preserving accuracy.

Red Hat AI Inference Server

Red Hat AI Inference Server is included in OpenShift AI and RHEL AI

Gain consistent, fast, and cost-effective inference at scale.

Compress models

Reduce compute and costs while preserving accuracy.

Inference runtime (vLLM) for the hybrid cloud

Run your models of choice across any accelerator and any environment.

Certified for all Red Hat products

Deployable across non-Red Hat Linux and Kubernetes platforms.

Red Hat AI Hugging Face repository

Access a collection of third-party validated and optimized models ready for inference.

Optimized for Google Cloud TPUs

Leverage leading price/performance of TPUs with Day0 support for all supported models

Get started

Get a free OpenShift migration assessment

Get in touch with a Google Cloud expert for a free, no-obligation assessment of your existing Red Hat OpenShift footprint to better understand your cost saving potential.

Try Google Cloud Netapp Volumes yourself

Explore NetApp Volumes in Google Cloud Skills Boost or let us set up a customer Hands-on Lab workshop!

Resources

Read: Save storage costs simply by enabling NetApp Volumes in your environment today

Read: Deploy and configure the Red Hat OpenShift Container Platform on Google Cloud with NetApp

Watch: Google Cloud NetApp Volumes for Openshift Dedicated

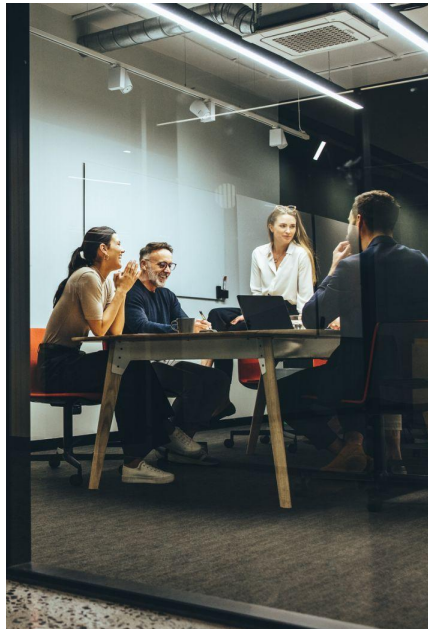

Thank you

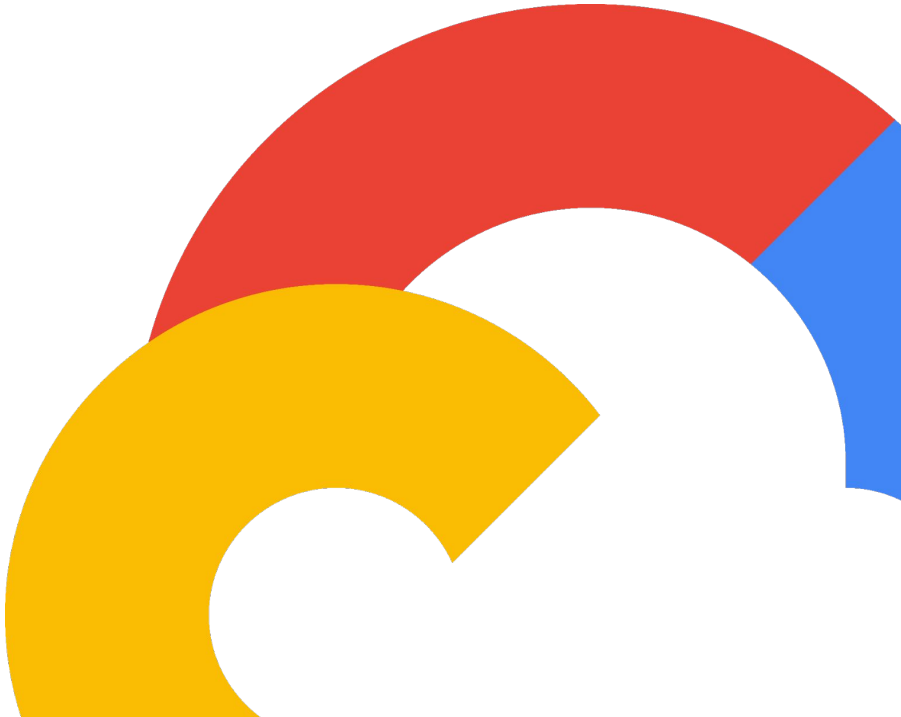

Google Cloud | Red Hat | NetApp