VEEAM OpenShift Backup – Day 2 Operations

Supporting Red Hat OpenShift Virtualization natively in Veeam Backup & Replication & Kasten

Veeam Kasten *for RedHat OpenShift*



DATA BACKUP

APPLICATION-CENTRIC

Policy-based data protection and management for the entire application.

STATEFUL AND STATELESS

Point-in-time capture of stateful data, whether inside or outside the cluster.

GITOPS INTEGRATED

Integrated into GitOps pipelines for the holistic protection of application configurations and data.

KUBERNETES NATIVE

Cloud-native, Kubernetes-aware protection for nodes, pods, namespaces, and more.



DATA RECOVERY

COMPLETE APPLICATION RECOVERY

Recover entire applications, including all data and artifacts, either in place or to a new namespace.

GRANULAR RESTORE

Full control over what artifacts and data to restore, allowing you to quickly recover exactly what you need.

AUTOMATED DR

Failover to any Kubernetes distribution and transform your application, including Veeam Kasten itself.

TEST AND DEVELOPMENT

Leverage production-fresh data in backups for business continuity planning, DR testing, and other use cases.

DATA SECURITY

COMPLIANCE READY

Multiple industry certifications and validations, including FIPS 140-3 and ISO/IEC 27001.

IMMUTABLE BACKUP

Object locking that safeguards backups from disaster scenarios and security incidents.

SIEM INTEGRATED

Direct ingestion of Veeam Kasten events into multiple SIEM systems for activity correlation and analysis.

POLICY-AS-CODE

Code-defined and managed rules and conditions to ensure all Kubernetes applications are protected.



DATA FREEDOM

AGNOSTIC MOBILITY

Move entire applications to and from any distribution across cloud and on-premises environments.

TRANSFORMS ENGINE

Upon restore and migration, modify Kubernetes resources to fit your new target environment.

VMS ON KUBERNETES

Unify the protection of modern VMs on Kubernetes that run alongside cloud-native applications.

VEEAM DATA PLATFORM INTEGRATED

Manage backup and restore tasks for holistic management of traditional and modern applications.



DATA INTELLIGENCE

APPLICATION DISCOVERY

Automatically discover all stateful and stateless resources that make up your entire application.

BLUEPRINT BINDINGS

Ensure data consistency across critical services during snapshot, export, and restore operations.

GITOPS PIPELINE

Ensure that the state of the entire application is protected dynamically between version upgrades.

DATALABS AND SUREBACKUP

Leverage cloning technologies to unlock intelligence with production-fresh data in backups.

ALIGNING TO RED HAT'S RESPONSE TO THE MARKET

Red Hat created a new OpenShift offer to lower the barrier and help customers migrate VM workloads to OpenShift.

OPENSHIFT VIRTUALIZATION ENGINE (OVE), A VM ONLY OPENSHIFT SOLUTION.

MAINTAIN

vShere Foudation
Stay as-is

- Slow evolution
- Increasing costs
- Low developer productivity

MIGRATE

OpenShift Virtualization Engine
Replace VMware
Utilize automation and services

- Dedicated VM management
- Reduced operating cost
- Increased IT efficiency and reliability

MODERNIZE

OpenShift
Application Modernization
Automation for Day 2 Ops

- Integrate containers and VMs
- Faster time to market
- Higher developer output

OpenShift AI
New application

- AI/ML Management
- Increase revenue
- Reduce Skill sprawl

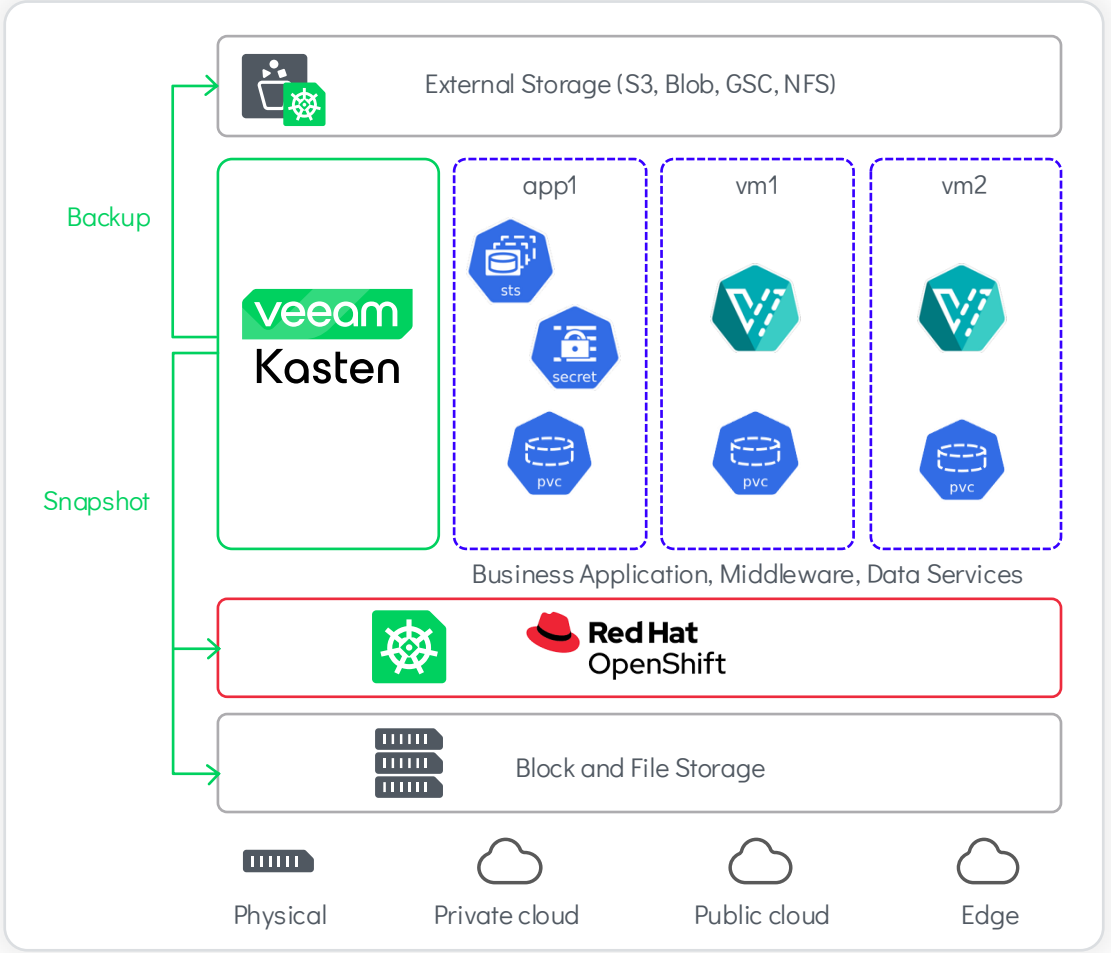
Veeam created an offer to align to Red Hat's OVE:
Veeam Kasten for Containerized VMs

VM WORKLOADS MANAGED AND PROTECTED WITH

RED HAT OPENSHIFT VIRTUALIZATION ENGINE AND VEEAM KASTEN FOR CONTAINERIZED VMS

Veeam now has two Veeam Kasten offerings:

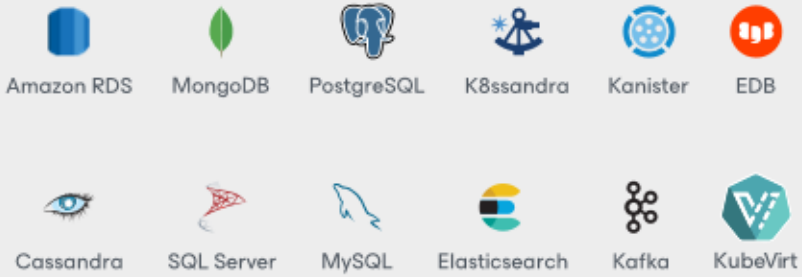
- **Veeam Kasten for Kubernetes:** Protects all applications and data including VMs and containerized applications in OpenShift (or any Kubernetes distribution).
- **Veeam Kasten for Containerized VMs:** Protect only **VMs** in OpenShift or any Kubernetes distribution that supports KubeVirt (i.e. OpenShift or SUSE Virtualization).



Veeam Kasten

Kubernetes Data Protection Ecosystem

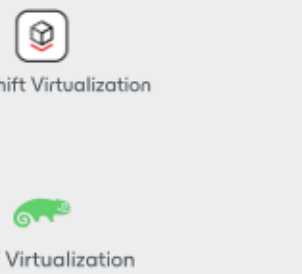
Applications



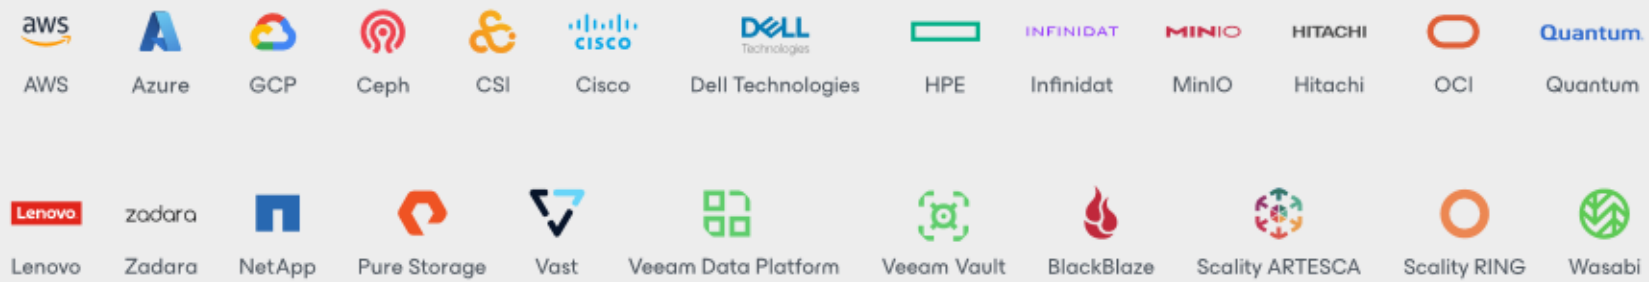
Kubernetes Distributions



Kubernetes Virtualization



Storage Infrastructure



Security Services



RED HAT OPENSHIFT VIRTUALIZATION with Veeam Backup & Replication

What Is Red Hat OpenShift Virtualization?



Built on KubeVirt

OpenShift Virtualization is a feature of Red Hat OpenShift Container Platform, built on the open-source KubeVirt project.



VMs as Native Workloads

Traditional virtual machines run as native, managed Kubernetes workloads alongside containers on the same cluster.



Declarative Management

VMs are managed via Kubernetes-style Custom Resources (CRs), using the same OpenShift console, oc CLI, and virtctl tooling as containers.



Deployed via OLM

OpenShift Virtualization is deployed and managed via the Operator Lifecycle Manager (OLM).

veeam

VMS AS FIRST-CLASS KUBERNETES OBJECTS

FULL LIFECYCLE MANAGEMENT

virt-controller manages the full VM lifecycle: power-on, shutdown, reboot, and deletion.

STANDARD KUBERNETES STORAGE

Storage is PVC/CSI-backed, like any other Kubernetes workload.

RECENT CAPABILITY HIGHLIGHTS

CROSS-CLUSTER LIVE MIGRATION

Zero-downtime VM migration between OpenShift clusters — across regions, clouds, or datacenters — eliminating maintenance windows.

USER-DEFINED NETWORKS

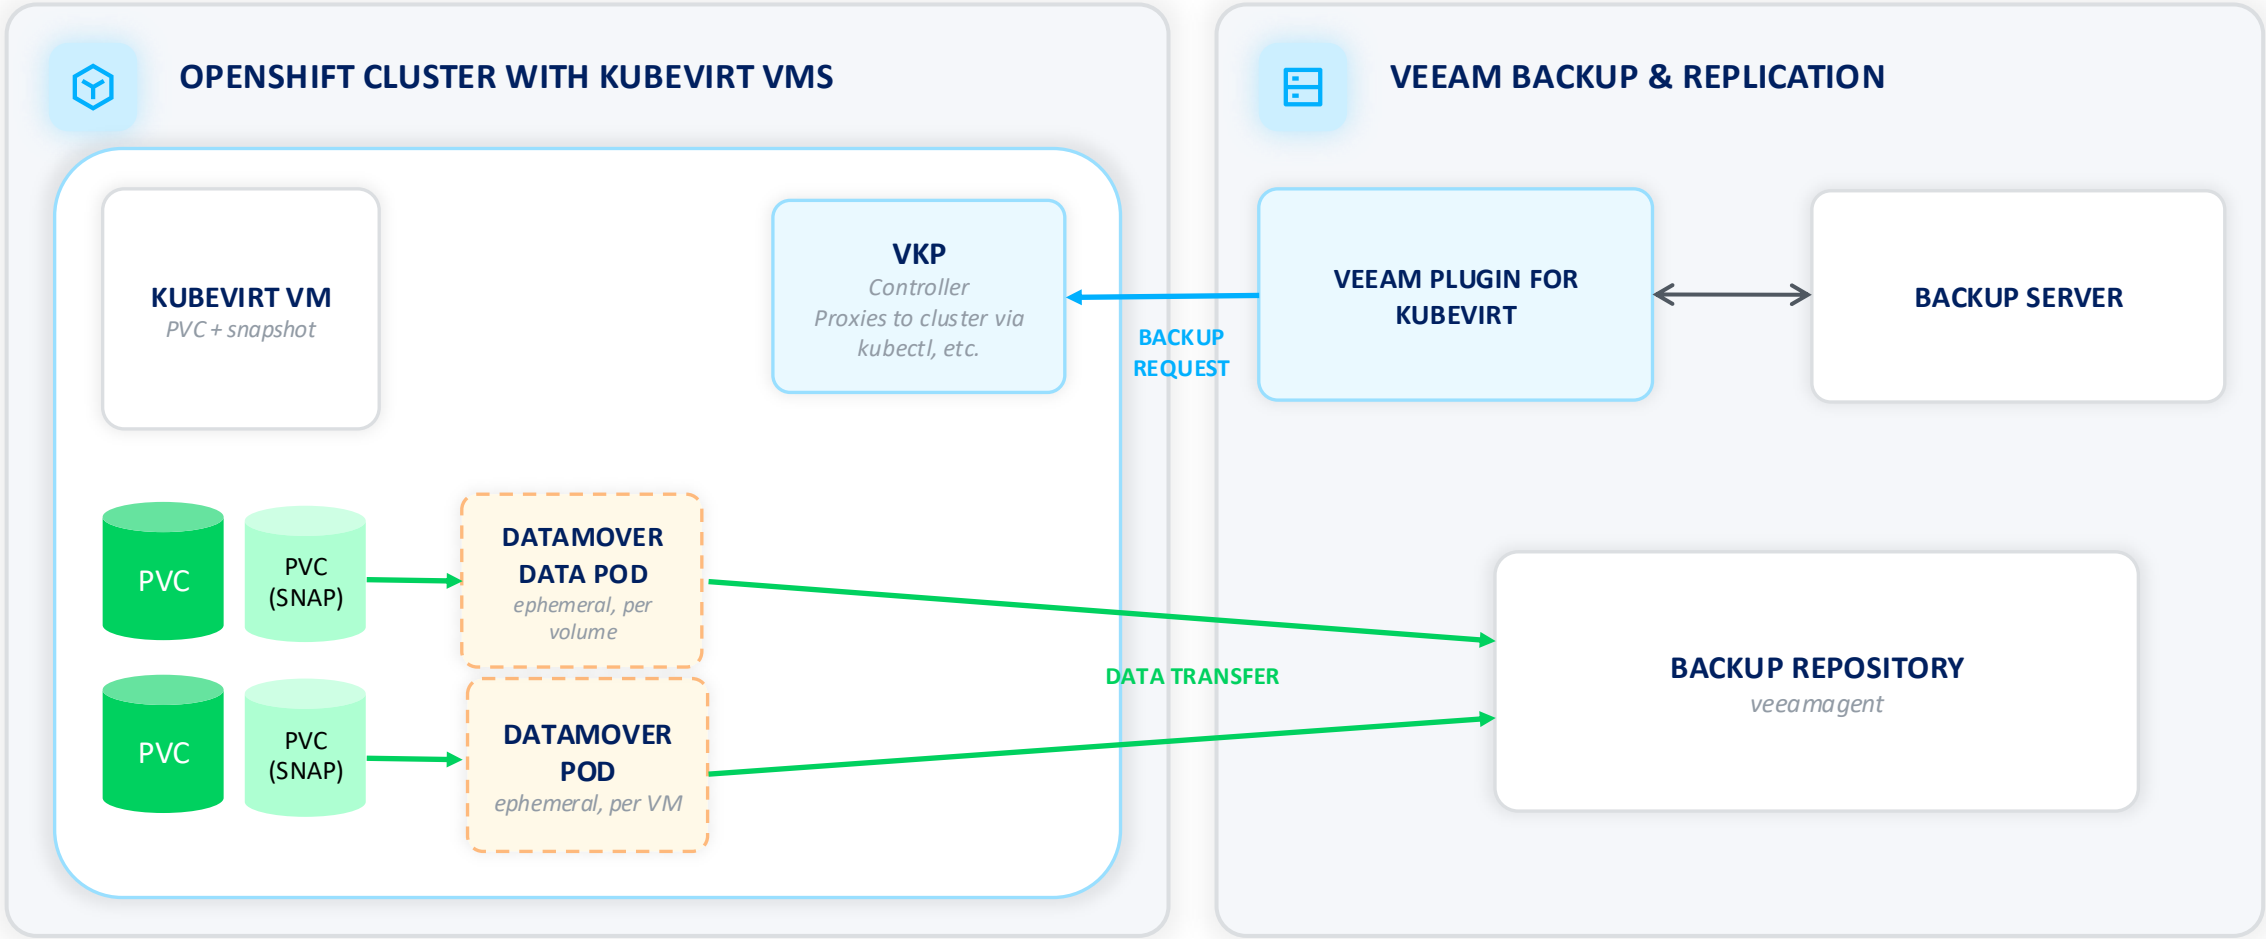
Isolated private networks within or across namespaces.

© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

VEEAM KUBEVIRT PROXY (VKP)

How VBR integrates via the new plug-in + in-cluster-proxy architecture

BACKUP DATA-FLOW WITH VKP



© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

PREREQUISITES

VEEAM BACKUP & REPLICAION HOST

VBR 13.1+ with the Plugin for KubeVirt installed.

RED HAT OPENSHIFT CLUSTER

Healthy, running cluster, version 4.20+.

MANAGEMENT HOST

Any machine (can be the VBR host itself) with oc (OpenShift CLI, 4.20+) and helm (4.1+) able to reach the cluster.

```
oc version
oc login --server="URL:PORT" -u USERNAME -p PASSWORD
oc get nodes
helm version
```

© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

DEPLOYING VkP VIA HELM

1 **ADD & UPDATE THE HELM REPO**

```
helm repo add vkp https://charts.veeam.com/
helm repo update
```

2 **INSTALL KROXY (ACCEPTING THE EULA)**

```
helm install vkp vkp/vkp --namespace=vkp \
  --set eula.accept=true --create-namespace
```

3 **ANNOTATE THE VOLUMESNAPSHOTCLASS**

```
oc get volumesnapshotclass
oc annotate volumesnapshotclass SNAPSHOT_CLASS_NAME \
  vkp.veeam.com/is-snapshot-class=true
```

4 **RETRIEVE VBR CONNECTION DETAILS**

```
oc get route -n vkp
oc get service -n vkp
oc get secret vkp-token -n vkp -o jsonpath='{.data.token}' | base64 -d
```

© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

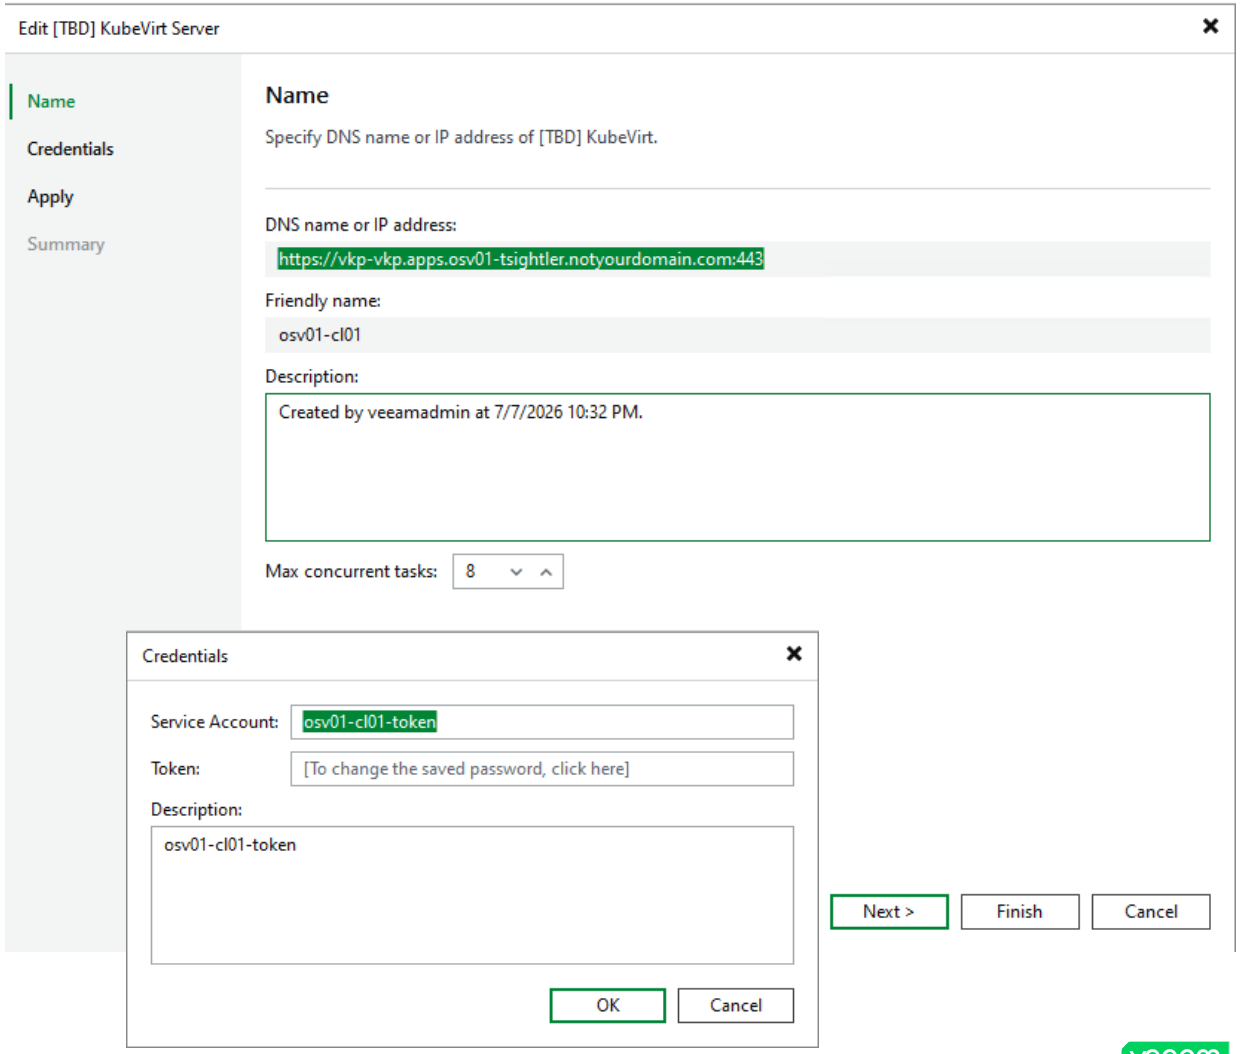
DEPLOYING VKP VIA HELM (cont)

RETRIEVE VBR CONNECTION DETAILS

```
oc get route -n vkp
NAME   HOST/PORT                                          PATH
SERVICES   PORT   TERMINATION     WILDCARD
vkp    vkp-vkp.apps.osv01-tsightler.notyourdomain.com          vkp
8080   edge/Redirect   None

oc get service -n vkp
NAME   TYPE       CLUSTER-IP       EXTERNAL-IP   PORT(S)
AGE
vkp    NodePort   172.30.241.217   <none>        8080:31944/TCP
30d

oc get secret vkp-token -n vkp -o jsonpath='{.data.token}' |
base64 -d
eyJhbGciOiJSUzI1NiIsImtpZCI6IndjTUJpZFdFbTZwdDduMk54THdkSTF6cWlMeW
ZibkpQc2tla0trbmZja00ifQ.eyJpc3MiOiJrdWJlcm5ldGVzL3NlcnZpY2VhY2Nvd
W50Iiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9uYW1lc3BhY2UiOiJ2a3A
iLCJrdWJlcm5ldGVzLmlvL3NlcnZpY2VhY2NvdW50L3NlY3JldC5uYW1lIjoidmtwL
XRva2VuIiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZXJ2aWNlLWFjY29
1bnQubmFtZSI6InZrcCIsImt1YmVybmV0ZXMuaW8vc2VydmljZWFjY291bnQvc2Vyd
mljZS1hY2NvdW50LnVpZCI6ImJjOTkwZDU5LTUyYzAtNGViZS1hZjIxLTJiZjZjZTU
4YmZhOCIsInN1YiI6InN5c3RlbTpzZXJ2aWNlYWNjb3VudDp2a3A6dmtwIn0...
```



© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

SIZING & THROUGHPUT

2.5 vCPU	0.5 GB RAM	≈0.5 GB/s	≈0.45 GB/s
Per Backup Worker (1 core Reader + 1.5 cores Data Mover)	Per Backup Worker (Reader + Data Mover combined)	Incremental throughput (≈1.8 TB/h)	Full throughput (≈1.6 TB/h)

MAX CONCURRENT TASKS	EXPECTED THROUGHPUT	CLUSTER CPU CORES	CLUSTER RAM	NETWORK
4	≈2.0 GB/s (≈7.2 TB/h)	10	2 GB	10 Gbps
8	≈4.0 GB/s (≈14.4 TB/h)	20	4 GB	25 Gbps
16	≈8.0 GB/s (≈28.8 TB/h)	40	8 GB	50 Gbps
32	≈16.0 GB/s (≈57.6 TB/h)	80	16 GB	100 Gbps

Sizing assumes all concurrent tasks operating at peak incremental throughput (500 MB/s per task) and ~2× source-side data reduction before transmission to the repository. Full backups run at slightly lower throughput and should be expected during initial seeding or scheduled active fulls.

ADDITIONAL CONSIDERATIONS

- Unlike other Backup Worker platforms, this uses a fixed 2.5 vCPU allocation per concurrent task rather than a scalable single worker size -- aggregate throughput scales by increasing Max Concurrent Tasks, not by resizing pods.
- Ensure the cluster has sufficient allocatable CPU/memory/worker nodes for the configured Max Concurrent Tasks, on top of capacity for other cluster workloads.
- Full and incremental backups place similar load on the cluster (both read the entire source disk), so plan backup windows accordingly.

© 2026 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Thank You!