

Cloud Scale Technology

*This guide is designed to highlight
Veritas Cloud Scale Technology.*

*For more information about Veritas
products and solutions, please visit
[veritas.com](https://www.veritas.com).*

Contents

- Introduction 4
- Executive Summary 4
- Target Audience 4
- About Veritas Cloud Scale Technology. 4
- Cloud Scale Technology Components 5
- Aks/Eks Prerequisites. 6
- Conclusion 8

Revision History

Version	Date	Changes	Author
1.00	10/7/2022	Initial Version	Neil Glick

Introduction

Executive Summary

With the majority of enterprises accelerating their multi-cloud strategies in pursuit of greater business agility and operational efficiencies, they need a data management solution that can help them achieve these goals. Veritas Cloud Scale Technology redefines data management for the next decade. Cloud Scale Technology's service elasticity and modern web-scale technologies enable NetBackup™ to operate cloud-natively within a cloud, yet deliver a consistent experience across multiple clouds to improve cloud return on investment (ROI), service resiliency, and security while reducing operational complexity and costs.

Target Audience

This document is targeted toward the prospect or customer interested in learning more about how Veritas Cloud Scale Technology can help ensure that their data is protected in the cloud at all times.

About Veritas Cloud Scale Technology

Cloud Scale Technology is a new generation of the proven NetBackup architecture. It is modernized to be cloud-native and ensure scalable, cost-effective, and elastic operation, and it will serve as the foundation for autonomous data management. Cloud Scale Technology is designed to operate cloud-natively and use technologies such as artificial intelligence (AI), machine learning (ML), containers, and microservices along with web-scale IT techniques such as service elasticity and hyper-automation.

Cloud-Native Characteristics
A containerized, Kubernetes-based deployment model that can be used to create a new cloud-native NetBackup environment, or complement an existing one that spans the data center and the cloud
A microservices-based architecture that provides the portability to work within multiple clouds, and resiliency for service availability
Elastic services that autonomously grow and shrink as needed to optimize cloud resource usage and costs
API-driven microservices that enable cross-domain workflow automation

Table 1. Cloud Scale Technology cloud-native characteristics

Here are Additional Examples of Benefits that Only NetBackup Powered by Cloud Scale Technology Can Provide:

- Leverage global, end-to-end deduplication on-premises and throughout the cloud, with multiple options of where and how deduplication is performed
- Reduce storage costs as much as 90 percent by combining cloud-native snapshots with NetBackup's automated snapshot lifecycle management; NetBackup combines cloud-native services such as Azure snapshots with deduplication and intelligent automation to seamlessly reduce the size of snapshots and store them on lower-cost storage
- Scale and protect cloud workloads with efficient and secure object storage support; access data directly from backup object storage, providing a virtual file system view and malware scanning of data that is compressed, encrypted, and deduplicated
- Simplified deployment directly from public cloud marketplaces and native tools

Cloud Scale Technology Components

The Cloud Scale Technology solution is deployed into your Azure or AWS Kubernetes cluster environment. AKS must be created with appropriate network and configuration settings. You can have multiple Cloud Scale Technology deployments in the same AKS or EKS cluster. Each Cloud Scale Technology deployment runs in a dedicated namespace on a dedicated node pool.

The following are the components needed for Cloud Scale Technology:

- **MDS (MetaData service)**

MDS is an independent and stackable service that provides a single system view of Cloud Scale Technology; it is an etcd cluster running inside the MDS pods, and these pods run on different AKS/EKS nodes; the pod name has a format of <cr-name>-uss-mds-<1,2...>the number of pods that get created depends on the number of Cloud Scale Technology engines in AKS/EKS cluster these pods are controlled by the MSDP operator

- 1 or 2 Cloud Scale Technology engines: 1 pod
- 3 or 4 Cloud Scale Technology engines: 3 pods
- 5 or more Cloud Scale Technology engines: 5 pods

- **Cloud Scale Technology Controller**

The Controller is a singleton service and the entry point of Cloud Scale Technology that monitors and repairs the MSDP Engines; it controls and manages the application-level business of the Cloud Scale Technology; the Deployment object name has a format of <cr-name>-uss-controller; it is controlled by the MSDP operator

- **Cloud Scale Technology/MSDP Engine**

MSDP Engines provide the ability to write deduplicated data to the cloud storage; the name of a MSDP engine pod is the corresponding fully qualified domain name (FQDN) of the static IP that is specified in the CR; each MSDP engine pod has MSDP services such as spad, spool, and ocsd running; they are controlled by the MSDP operator.




Figure 1. Cloud Scale Technology architecture

AKS/EKS Prerequisites

Azure Kubernetes Cluster

Your Azure Kubernetes cluster must be created with appropriate network and configuration settings.

Supported Azure Kubernetes cluster version is 1.21.x and later.

- Availability zone for AKS cluster must be disabled
- At least one storage class is backed with Azure disk CSI storage driver `disk.csi.azure.com`, and allows volume expansion; it must be in LRS category with Premium SSD. For example, the built-in storage class `managed-csi-premium`; it is recommended that the storage class has Retain reclaim
- Cert-Manager must be installed
- A Kubernetes Secret that contains the MSDP credentials is required
- Enable AKS Uptime SLA; AKS Uptime SLA is recommended for better resiliency for information about AKS Uptime SLA and to enable it, see [Azure Kubernetes Service \(AKS\) Uptime SLA](#)
- Azure container registry (ACR); use existing ACR or create a new one; your Kubernetes cluster must be able to access this registry to pull the images
- You must have a dedicated node pool for Cloud Scale Technology created; Azure availability zone must be disabled; Azure autoscaling allows your node pool to scale dynamically as required; if Azure autoscaling is not enabled, ensure the node number is not less than Cloud Scale Technology size; it is recommended that you set the minimum node number to one or more to bypass some limitations in AKS
- Client machine to access AKS cluster
 - A separate computer that can access and manage your AKS cluster and ACR
 - It must have the Linux operating system
 - It must have Docker daemon, the Kubernetes command-line tool (`kubectl`), and Azure CLI installed; the Docker storage size must be more than 6 GB; the version of `kubectl` must be `v1.19.x` or later; the version of Azure CLI must meet the AKS cluster requirements
 - If AKS is a private cluster, see [Create a private Azure Kubernetes Service cluster](#)
- If the internal IPs are used, reserve the internal IPs (avoid the IPs that are reserved by other systems) for MSDP scaleout and add DNS records for all of them in your DNS configuration; the Azure static public IPs can be used but this is not recommended; if Azure static public IPs are used, create them in the node resource group for the AKS cluster; a DNS name must be assigned to each static public IP; the IPs must be in the same location of the AKS cluster

AWS Kubernetes Cluster

Ensure that the following prerequisites are met before proceeding with the deployment:

EKS-specific requirements.

Create a Kubernetes cluster with the following guidelines:

- Use Kubernetes version 1.21 onwards
- AWS default CNI is used during cluster creation
- Create a nodegroup with only one availability zone; instance type should be at least m5.4xlarge configuration

Note: Using separate nodegroups is required for the NetBackup Operator, NetBackup, and MSDP deployments

The nodepool uses AWS' manual or autoscaling group feature which allows your nodepool to scale by provisioning and de-provisioning the nodes as required automatically.

Note: All nodes in the node group must be running on the Linux operating system.

- Minimum required policies in IAM role:
 - AmazonEKSClusterPolicy
 - AmazonEKSWorkerNodePolicy
 - AmazonEC2ContainerRegistryReadOnly
 - AmazonEKS_CNI_Policy
 - AmazonEKSServicePolicy
- Use an existing AWS Elastic Container Registry or create a new one and ensure that the EKS has full access to pull images from the elastic container registry
- Deploy AWS load balancer controller add-on in the cluster
- Install cert-manager by using the following command:
 - `$ kubectl apply -f`
 - github.com/cert-manager/cert-manager/releases/download/v1.8.0/cert-manager.yaml
- The FQDN that will be provided in primary server CR and media server CR specifications in networkLoadBalancer section must be DNS resolvable to the provided IP address
- Create a storage class with EBS storage type and retain reclaim policy; provide the name of the created EBS storage type in the storage class name field during the deployment of primary server/media server CR
- If the NetBackup client is outside VPC, or if you want to access the WEB UI from outside VPC, then NetBackup client CIDR must be added with all NetBackup ports in security group inbound rule of cluster
- To obtain the cluster security group, run the following command:
 - `aws eks describe-cluster --name <my-cluster> --query cluster.resourcesVpcConfig.clusterSecurityGroupId`

Host-Specific Requirements

- Install AWS CLI
- Install Kubectl CLI
- Configure docker to enable the push of the container images to the container registry
- Create the OIDC provider for the AWS EKS cluster
- Create an IAM service account for the AWS EKS cluster
- If an IAM role needs access to the EKS cluster, run the following command from the system that already has access to the EKS cluster:
 - `kubectl edit -n kube-system configmap/aws-auth`
- Log in to the AWS environment to access the Kubernetes cluster by running the following command on AWS CLI:
 - `aws eks --region <region_name> update-kubeconfig --name <cluster_name>`
- Free space of approximately 8.5 GB on the location where you copy and extract the product installation TAR package file; if using docker locally, there should be approximately 8 GB available on the `/var/lib/docker` location so that the images can be loaded to the docker cache, before being pushed to the container registry

Conclusion

As more workloads are moved to the cloud, cloud administrators need an enterprise backup solution built with the cloud in mind. Cloud Scale Technology gives Cloud Administrators a scalable, highly available, fault tolerant enterprise solution that is resilient and cost effective.

About Veritas

Veritas Technologies is a leader in multi-cloud data management. Over 80,000 customers—including 95 percent of the Fortune 100—rely on Veritas to help ensure the protection, recoverability, and compliance of their data. Veritas has a reputation for reliability at scale, which delivers the resilience its customers need against the disruptions threatened by cyberattacks, like ransomware. No other vendor is able to match the ability of Veritas to execute, with support for 800+ data sources, 100+ operating systems, 1,400+ storage targets, and 60+ clouds through a single, unified approach. Powered by Cloud Scale Technology, Veritas is delivering today on its strategy for Autonomous Data Management that reduces operational overhead while delivering greater value. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

VERITAS™

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact
information visit:
veritas.com/company/contact