

NetBackup™ Deployment Guide for Amazon Elastic Kubernetes Services (EKS) Cluster

Release 10.1.1

VERITAS™

NetBackup™ Deployment Guide for Amazon Elastic Kubernetes Services (EKS) Cluster

Last updated: 2022-12-09

Legal Notice

Copyright © 2022 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Veritas Technologies LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introduction to NetBackup on EKS	9
	About NetBackup deployment on Amazon Elastic Kubernetes (EKS) cluster	9
	Required terminology	11
	User roles and permissions	12
	About MSDP Scaleout	17
	About MSDP Scaleout components	17
	Limitations in MSDP Scaleout	18
Chapter 2	Deployment with environment operators	19
	About deployment with the environment operator	19
	Prerequisites	19
	Contents of the TAR file	21
	Known limitations	21
	Deploying the operators manually	21
	Deploying NetBackup and MSDP Scaleout manually	27
	Deploying NetBackup and Snapshot Manager manually	32
	Configuring the <code>environment.yaml</code> file	39
	Uninstalling NetBackup environment and the operators	52
	Applying security patches	54
Chapter 3	Assessing cluster configuration before deployment	59
	How does the webhook validation works	59
	Webhooks validation execution details	60
	How does the Config-Checker utility work	61
	Config-Checker execution and status details	62
Chapter 4	Deploying NetBackup	64
	Preparing the environment for NetBackup installation on EKS	64
	Recommendations of NetBackup deployment on EKS	73
	Limitations of NetBackup deployment on EKS	74
	About primary server CR and media server CR	74
	After installing primary server CR	76

	After Installing the media server CR	76
	Monitoring the status of the CRs	77
	Updating the CRs	79
	Deleting the CRs	80
	Configuring NetBackup IT Analytics for NetBackup deployment	81
	Managing NetBackup deployment using VxUpdate	82
	Migrating the node group for primary or media servers	83
Chapter 5	Upgrading NetBackup	84
	Preparing for NetBackup upgrade	84
	Upgrading NetBackup operator	85
	Upgrading NetBackup application	86
	Upgrade NetBackup from previous versions	88
	Procedure to rollback when upgrade fails	97
Chapter 6	Deploying Snapshot Manager	102
	Overview	102
	Prerequisites	102
	Installing the docker images	104
Chapter 7	Migration and upgrade of Snapshot Manager	108
	Migration and updating of Snapshot Manager	108
Chapter 8	Deploying MSDP Scaleout	114
	Deploying MSDP Scaleout	114
	Prerequisites	115
	Installing the docker images and binaries	117
	Initializing the MSDP operator	118
	Configuring MSDP Scaleout	119
	Using MSDP Scaleout as a single storage pool in NetBackup	121
	Configuring the MSDP cloud in MSDP Scaleout	122
Chapter 9	Upgrading MSDP Scaleout	123
	Upgrading MSDP Scaleout	123
Chapter 10	Monitoring NetBackup	125
	Monitoring the application health	125
	Telemetry reporting	127

	About NetBackup operator logs	128
	Expanding storage volumes	129
	Allocating static PV for Primary and Media pods	130
Chapter 11	Monitoring MSDP Scaleout	136
	About MSDP Scaleout status and events	136
	Monitoring with Amazon CloudWatch	138
	The Kubernetes resources for MSDP Scaleout and MSDP operator	142
Chapter 12	Monitoring Snapshot Manager deployment	144
	Overview	144
	Logs of Snapshot Manager	144
	Configuration parameters	145
Chapter 13	Managing the Load Balancer service	146
	About the Load Balancer service	146
	Notes for Load Balancer service	149
	Opening the ports from the Load Balancer service	149
Chapter 14	Performing catalog backup and recovery	151
	Backing up a catalog	151
	Restoring a catalog	153
Chapter 15	Managing MSDP Scaleout	156
	Adding MSDP engines	156
	Adding data volumes	157
	Expanding existing data or catalog volumes	158
	Manual storage expansion	158
	MSDP Scaleout scaling recommendations	159
	MSDP Cloud backup and disaster recovery	160
	About the reserved storage space	160
	Cloud LSU disaster recovery	161
	MSDP multi-domain support	164
	Configuring Auto Image Replication	164
	About MSDP Scaleout logging and troubleshooting	165
	Collecting the logs and the inspection information	166

Chapter 16	About MSDP Scaleout maintenance	167
	Pausing the MSDP Scaleout operator for maintenance	167
	Logging in to the pods	167
	Reinstalling MSDP Scaleout operator	168
	Migrating the MSDP Scaleout to another node group	168
Chapter 17	Uninstalling MSDP Scaleout from EKS	170
	Cleaning up MSDP Scaleout	170
	Cleaning up the MSDP Scaleout operator	171
Chapter 18	Uninstalling Snapshot Manager	173
	Uninstalling Snapshot Manager from EKS	173
Chapter 19	Troubleshooting	175
	View the list of operator resources	176
	View the list of product resources	177
	View operator logs	180
	View primary logs	180
	Pod restart failure due to liveness probe time-out	180
	Socket connection failure	181
	Resolving an invalid license key issue	182
	Resolving an issue where external IP address is not assigned to a NetBackup server's load balancer services	183
	Resolving the issue where the NetBackup server pod is not scheduled for long time	183
	Resolving an issue where the Storage class does not exist	184
	Resolving an issue where the primary server or media server deployment does not proceed	185
	Resolving an issue of failed probes	186
	Resolving token issues	187
	Resolving an issue related to insufficient storage	188
	Resolving an issue related to invalid nodepool	188
	Resolving a token expiry issue	189
	Resolve an issue related to KMS database	190
	Resolve an issue related to pulling an image from the container registry	190
	Resolving an issue related to recovery of data	191
	Check primary server status	192
	Pod status field shows as pending	192
	Ensure that the container is running the patched image	193

Getting EEB information from an image, a running container, or persistent data	198
Resolving the certificate error issue in NetBackup operator pod logs	200
Resolving the primary server connection issue	201
Primary pod is in pending state for a long duration	201
Host mapping conflict in NetBackup	202
NetBackup messaging queue broker take more time to start	202
Local connection is getting treated as insecure connection	203
Issue with capacity licensing reporting which takes longer time	204
Backing up data from Primary server's /mnt/nbdata/ directory fails with primary server as a client	204
Wrong EFS ID is provided in <code>environment.yaml</code> file	204
Primary pod is in ContainerCreating state	205
Webhook displays an error for PV not found	206
Appendix A	
CR template	208
Secret	208
MSDP Scaleout CR	209

Introduction to NetBackup on EKS

This chapter includes the following topics:

- [About NetBackup deployment on Amazon Elastic Kubernetes \(EKS\) cluster](#)
- [Required terminology](#)
- [User roles and permissions](#)
- [About MSDP Scaleout](#)
- [About MSDP Scaleout components](#)
- [Limitations in MSDP Scaleout](#)

About NetBackup deployment on Amazon Elastic Kubernetes (EKS) cluster

NetBackup provides the product deployment solution on Amazon Elastic Kubernetes (EKS) cluster, in the Amazon Web Services (AWS) Cloud. The solution facilitates an orchestrated deployment of the NetBackup components on EKS.

You can deploy NetBackup on EKS for scaling the capacity of the NetBackup host to serve a large number of requests concurrently running on the NetBackup primary server at its peak performance capacity.

This guide provides you two distinct methods of deployment. The first and the recommended one is by using the environment operators. In this method, you can deploy the entire NetBackup environment with ease. You can deploy, one primary, and optionally, one media with one or more replicas, and one MSDP Scaleout with

four to 16 replicas. The guide describes a very comprehensive method to deploy, configure, and remove the NetBackup components using the environment operators.

You can also go for a discrete deployment of the NetBackup components without using the environment operator. This method is not the recommended method of deployment.

Supported platforms

Currently we support Amazon Elastic Kubernetes Service.

About the guide

This guide contains the following sections:

- Introduction to NetBackup and MSDP Scaleout —preparatory steps to ensure that your EKS cluster and hardware environment meet the deployment requirements.
- Deploying with environment operators—deploy the entire NetBackup environment primary, media, and MSDP Scaleout servers together in a comprehensive way. This is the most recommended method of deployment.
- Assessing cluster configuration before deployment—check the deployment environment to verify that the environment meets the requirements, before starting the primary server and media server deployments.
- Deploying NetBackup—deploying NetBackup without the environment operator.
- Deploying MSDP Scaleout—deploying MSDP Scaleout without the environment operator.
- Monitoring NetBackup—monitor application health, view logs, expand storage volume and so on.
- Monitoring MSDP Scaleout—monitor status, alerts, events, AWS container insights, and so on.
- Configuring the Load Balancer service—configure the load balancer to access NetBackup from private IPs.
- Performing catalog backup and recovery—how to backup the catalog and recover.
- Configuring MSDP Scaleout—adding MSDP Scaleout engines and data volumes, disaster recovery, scaling and so on.
- Maintaining MSDP Scaleout—running maintenance, logging, reinstalling the operator, and so on.
- Uninstalling MSDP Scaleout from EKS—uninstall and cleanup the cluster and the operator.

- Scenarios for troubleshooting.

The intended audience for this document includes backup, cloud, system administrators, and architects.

Note: NetBackup deployment for EKS offers only English language support and it does not support OpsCenter.

Required terminology

The table describes the important terms for NetBackup deployment on EKS cluster. For more information visit the link to Kubernetes documentation.

Table 1-1 Important terms

Term	Description
Pod	A Pod is a group of one or more containers, with shared storage and network resources, and a specification for how to run the containers. For more information on Pods, see Kubernetes Documentation .
StatefulSet	StatefulSet is the workload API object used to manage stateful applications and it represents a set of Pods with unique, persistent identities, and stable hostnames. For more information on StatefulSets, see Kubernetes Documentation .
Job	Kubernetes jobs ensure that one or more pods execute their commands and exit successfully. For more information on Jobs, see Kubernetes Documentation .
ConfigMap	A ConfigMap is an API object used to store non-confidential data in key-value pairs. For more information on ConfigMaps, see Kubernetes Documentation .
Service	A Service enables network access to a set of Pods in Kubernetes. For more information on Service, see Kubernetes Documentation .
Persistent Volume Claim	A PersistentVolumeClaim (PVC) is a request for storage by a user. For more information on Persistent Volumes, see Kubernetes Documentation .
Persistent Volume	A PersistentVolume (PV) is a piece of storage in the cluster that has been provisioned by an administrator or dynamically provisioned using storage classes. For more information on Persistent Volumes, see Kubernetes Documentation .

Table 1-1 Important terms (*continued*)

Term	Description
Custom Resource	A Custom Resource (CR) is an extension of the Kubernetes API that is not necessarily available in a default Kubernetes installation. For more information on Custom Resources, see Kubernetes Documentation .
Custom Resource Definition	The CustomResourceDefinition (CRD) API resource lets you define custom resources. For more information on CustomResourceDefinitions, see Kubernetes Documentation .
Secret	A Secret is an object that contains a small amount of sensitive data such as a password, a token, or a key. Such information might otherwise be put in a Pod specification or in a container image. For more information on Secrets, see Kubernetes Documentation .
ServiceAccount	A service account provides an identity for processes that run in a Pod. For more information on configuring the service accounts for Pods, see Kubernetes Documentation .
ClusterRole	An RBAC Role or ClusterRole contains rules that represent a set of permissions. Permissions are purely additive (there are no "deny" rules). For more information on ClusterRole, see Kubernetes Documentation .
ClusterRoleBinding	A role binding grants the cluster-wide permissions defined in a role to a user or set of users. For more information on ClusterRoleBinding, see Kubernetes Documentation .
Namespace	Kubernetes supports multiple virtual clusters backed by the same physical cluster. These virtual clusters are called namespaces. For more information on Namespaces, see Kubernetes Documentation .

User roles and permissions

Note the following for user authentication:

- An Administrator must define the custom user credentials by creating a secret; and then provide the secret name at the time of primary server deployment.
- A custom user is assigned the role of a NetBackup Security Administrator and can access the NetBackup Web UI after deployment.
- A custom user will be persisted during the pods restart or upgrade.

- For the custom user, you can change only the password after the deployment. The changed password will be persisted. If the username is changed after the deployment, an error message will be logged in the Operator pod.
- You can delete the secret after the primary server deployment. In that case, if you want to deploy or scale the media servers, you must create a new secret with the same username which was used in the primary server CR. The password can be the same or different. If you change the password, it is also changed in the primary server pod, and gets persisted.
- Do not create a local user in the pods (using the `kubectl exec` or `useradd` commands) as this user may or may not be persisted.
- The Amazon Web Service user is supported through Single Sign-on (SSO). For the detailed user integration information, refer to the *NetBackup Administrator's Guide Volume 1*.
- An **nbitanalyticsadmin** user is available in primary server container. This user is used as **Master Server User ID** while creating data collector policy for data collection on NetBackup IT Analytics portal.
- Service account that is used for this deployment is **netbackup-account** and it is defined in the `operator_deployment.yaml`.
- NetBackup runs most of the primary server services and daemons as non-root user (**nbsvcusr**) and only **root** and **nbsvcusr** are supported as a service account user.
- ClusterRole named **netbackup-role** is set in the NetBackup Operator to define the cluster wide permissions to the resources. This is defined in the `operator_deployment.yaml`.
- Appropriate roles and EKS specific permissions are set to the cluster at the time of cluster creation.
- After successful deployment of the primary and media servers, the operator creates a custom Kubernetes role with name `<resourceNamePrefix>-admin` whereas `resourceNamePrefix` is given in primary server or media server CR specification.

The following permissions are provided in the respective namespaces:

Resource name	API group	Allowed operations
ConfigMaps	default	<ul style="list-style-type: none"> ■ Create ■ Delete ■ Get ■ List ■ Patch ■ Update ■ Watch
Nodes	default	<ul style="list-style-type: none"> ■ Get ■ List

This role can be assigned to the NetBackup Administrator to view the pods that were created, and to execute into them. For more information on the access control, see [Kubernetes Access Control Documentation](#).

Note: One role would be created, only if primary and media servers are in same namespace with the same resource name prefix.

Role-based authentication (RBAC)

NetBackup Operator deployment uses a `serviceAccount` and it must have the following permissions:

Table 1-2

Resource Name	API Group	Allowed Operations
ConfigMaps	default	<ul style="list-style-type: none"> ■ Create ■ Delete ■ Get ■ List ■ Patch ■ Update ■ Watch
Nodes	default	<ul style="list-style-type: none"> ■ Get ■ List

Table 1-2 (continued)

Resource Name	API Group	Allowed Operations
PersistentVolumeClaims	default	<ul style="list-style-type: none">■ Create■ Delete■ Get■ List■ Patch■ Update■ Watch
Pods	default	<ul style="list-style-type: none">■ Create■ Delete■ Get■ List■ Patch■ Update■ Watch
Pods/exec	default	<ul style="list-style-type: none">■ Create■ Get
Secret	default	<ul style="list-style-type: none">■ Get■ List■ Watch
Services	default	<ul style="list-style-type: none">■ Create■ Delete■ Get■ List■ Patch■ Update■ Watch
StatefulSet	app	<ul style="list-style-type: none">■ Create■ Delete■ Get■ List■ Patch■ Update■ Watch

Table 1-2 (continued)

Resource Name	API Group	Allowed Operations
Jobs	batch	<ul style="list-style-type: none"> ■ Create ■ Delete ■ Get ■ List
Primary servers	netbackup.veritas.com	<ul style="list-style-type: none"> ■ Create ■ Delete ■ Get ■ List ■ Patch ■ Update ■ Watch
PrimaryServers/status	netbackup.veritas.com	<ul style="list-style-type: none"> ■ Get ■ Patch ■ Update
Media servers	netbackup.veritas.com	<ul style="list-style-type: none"> ■ Create ■ Delete ■ Get ■ List ■ Patch ■ Update ■ Watch
MediaServers/status	netbackup.veritas.com	<ul style="list-style-type: none"> ■ Get ■ Patch ■ Update
Secrets	netbackup.veritas.com	Watch
Secrets/status	netbackup.veritas.com	<ul style="list-style-type: none"> ■ Get ■ Patch ■ Update
Roles	rbac.authorization.k8s.io	<ul style="list-style-type: none"> ■ Create ■ Get ■ List ■ Watch

Table 1-2 (continued)

Resource Name	API Group	Allowed Operations
Storageclasses	storage.k8s.io	<ul style="list-style-type: none"> ■ Get ■ List
Deployment	app	<ul style="list-style-type: none"> ■ Get ■ List ■ Update ■ Watch

About MSDP Scaleout

MSDP Scaleout is based on MSDP. It empowers MSDP with high resilience and scalability capabilities to simplify management and reduce total cost of ownership.

It runs on multiple nodes to represent a single storage pool for NetBackup and other Veritas products to use. You can seamlessly scale out and scale up a MSDP Scaleout on demand. MSDP Scaleout automatically does failure detection and repair in the background.

It is deployed separately in NetBackup environment. The deployment process is with minimal user intervention. The core MSDP services run on each node to expose the storage optimized services, and manage a part of the cluster level data and metadata. Each MSDP Scaleout node is called MSDP engine.

See [“Deploying MSDP Scaleout”](#) on page 114.

About MSDP Scaleout components

Following are the MSDP Scaleout components:

- MDS (MetaData service)

MDS is an independent and stackable service that provides a single system view of MSDP Scaleout. It's an etcd cluster running inside the MDS pods. These pods run on different 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 MSDP Scaleout engines in EKS cluster. These pods are controlled by the MSDP operator.

 - 1 or 2 MSDP Scaleout engines: 1 pod
 - 3 or 4 MSDP Scaleout engines: 3 pods
 - 5 or more MSDP Scaleout engines: 5 pods

- **MSDP Scaleout Controller**
Controller is a singleton service and the entry point of MSDP Scaleout that monitors and repairs MSDP Engines. It controls and manages the application-level business of the MSDP Scaleout. The Deployment object name has a format of **<cr-name>-uss-controller**. It is controlled by the MSDP operator.
- **MSDP Scaleout Engine**
MSDP Engines provide the ability to write deduplicated data to the storage. The name of a MSDP engine pod is the corresponding FQDN of the static IP that is specified in the CR. Each MSDP engine pod has MSDP services such as spad, spoold, and ocsd running. They are controlled by the MSDP operator.

Limitations in MSDP Scaleout

MSDP Scaleout has the following limitations:

- It is not fully compliant with Federal Information Processing Standards (FIPS). The internal services MSDP operator, MSDP Controller, and MDS of a MSDP Scaleout are not compliant with FIPS.
MSDP is FIPS compliant. For more information, see the *NetBackup Deduplication Guide*.
- Does not support SELinux.
- Supports only NBCA. Does not support ECA.
- Node group cross availability zone is not supported.
- Limited EKS node failure tolerance.
Backup and restore can fail if EKS node fails. If MSDP operator detects the MSDP Scaleout pod failure, it attempts to restart it and perform a repair operation automatically. The repair operation can be delayed if AWS infrastructure or Kubernetes do not allow the pod to be restarted.
An AWS EBS volume cannot be attached to two different nodes at the same time. When the node to which AWS EBS volume is attached fails, MSDP operator cannot run the same pod with the same AWS EBS volume on another node until the failed node is repaired or deleted by EKS.
EKS node auto-repair may take more than 20 minutes to finish. In some cases, it may be necessary to bring the node backup manually.
See [Amazon Elastic Kubernetes Service \(EKS\) Documentation](#)
- IPv6 is not supported.

Deployment with environment operators

This chapter includes the following topics:

- [About deployment with the environment operator](#)
- [Deploying the operators manually](#)
- [Deploying NetBackup and MSDP Scaleout manually](#)
- [Deploying NetBackup and Snapshot Manager manually](#)
- [Configuring the environment.yaml file](#)
- [Uninstalling NetBackup environment and the operators](#)
- [Applying security patches](#)

About deployment with the environment operator

This section describes the deployment of the Veritas NetBackup and MSDP Scaleout on Amazon Elastic Kubernetes Service in AWS cloud. You can start by deploying the two environment operators that together manage the NetBackup environment, the primary server, the media servers, and the MSDP Scaleout storage servers.

Prerequisites

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

- A Kubernetes cluster in Amazon Elastic Kubernetes Service in AWS with multiple nodes. Using separate node group is recommended for the NetBackup servers,

MSSDP Scaleout deployments and for different media server objects. It is required to have separate node pool for Snapshot Manager data plane.

- Taints and tolerations allows you to mark (taint) a node so that no pods can schedule onto it unless a pod explicitly *tolerates* the taint. Marking nodes instead of pods (as in node affinity/anti-affinity) is particularly useful for situations where most pods in the cluster must avoid scheduling onto the node.

Taints are set on the node group while creating the node group in the cluster. Tolerations are set on the pods.

To use this functionality, user must create the node group with the following detail:

- Add a label with certain key value. For example `key = nbpool, value = nbnodes`
- Add a taint with the same key and value which is used for label in above step with effect as *NoSchedule*.
For example, `key = nbpool, value = nbnodes, effect = NoSchedule`
- Access to a container registry that the Kubernetes cluster can access, like an Amazon Elastic Kubernetes Service Container Registry.
- Install Cert-Manager. You can use the following command to install the Cert-Manager:

```
$ kubectl apply -f  
https://github.com/jetstack/cert-manager/releases/download/v1.6.0/cert-manager.yaml
```

For details, see <https://cert-manager.io/docs/installation/>
- A workstation or VM running Linux with the following:
 - Configure `kubectl` to access the cluster.
 - Install AWS CLI to access AWS resources.
 - Configure `docker` to be able to push images to the container registry.
 - Free space of approximately 8.5GB on the location where you copy and extract the product installation TAR package file. If using `docker` locally, there should be approximately 8GB 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.
 - AWS network load balancer controller add-on must be installed for using network load balancer capabilities.
 - AWS EFS-CSI driver must be installed for statically provisioning the PV or PVC in EFS for primary server.

Contents of the TAR file

Download the TAR file from the Veritas download center.

The TAR file contains the following:

Table 2-1 TAR contents

Item	Description
OCI images in the <code>/images</code> directory	These docker image files that are loaded and then copied to the container registry to run in Kubernetes. They include NetBackup and MSDP Scaleout application images and the operator images.
MSDP kubectl plug-in at <code>/bin/kubectl-msdp</code>	Used to deploy and manage the MSDP Scaleout operator tasks.
Configuration(.yaml) files at <code>/operator</code> directory	You can edit these to suit your configuration requirements before installation.
Sample product (.yaml) files at <code>/samples</code> directory	You can use these as templates to define your NetBackup environment.
<code>README.md</code>	Readme file.

Known limitations

Here are some known limitations.

- Changes to the CorePattern which specifies the path used for storing core dump files in case of a crash are not supported. CorePattern can only be set during initial deployment.
- Changes to MSDP Scaleout credential autoDelete, which allows automatic deletion of credential after use, is not supported. The autoDelete value can only be set during initial deployment.

Deploying the operators manually

To perform these steps, log on to the Linux workstation or VM where you have extracted the TAR file.

To deploy the operators

- 1 Install the MSDP kubect1 plug-in at some location which is set in the path environment variable of your shell. For example, copy the file `kubect1-msdp` to `/usr/local/bin/`.
- 2 Run the following commands to load each of the product images to the local docker instance.

```
$ docker load -i netbackup-main-10.1.1.tar.gz
```

```
$ docker load -i netbackup-operator-10.1.1.tar.gz
```

```
$ docker load -i pdcluster-17.0.tar.gz
```

```
$ docker load -i pdde-17.0.tar.gz
```

```
$ docker load -i pdk8soptr-17.0.tar.gz
```

```
$ docker load -i  
netbackup-flexsnap-$(SNAPSHOT_MANAGER_VERSION).tar.gz
```

Run the command `docker image ls` to confirm that the product images are loaded properly to the docker cache.

3 Run the following commands to re-tag the images to associate them with your container registry, keep the image name and version same as original:

```
$ REGISTRY=<<AccountID>.dkr.ecr.<region>.amazonaws.com

$ docker tag netbackup/main:10.1.1
${REGISTRY}/netbackup/main:10.1.1

$ docker tag netbackup/operator:10.1.1
${REGISTRY}/netbackup/operator:10.1.1

$ docker tag uss-engine:17.0 ${REGISTRY}/uss-engine:17.0

$ docker tag uss-controller:17.0 ${REGISTRY}/uss-controller:17.0

$ docker tag uss-mds:17.0 ${REGISTRY}/uss-mds:17.0

$ docker tag msdp-operator:17.0 ${REGISTRY}/msdp-operator:17.0

$ docker tag veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}

$ docker tag
veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}
```

4 Login using the following command:

```
docker login -u AWS -p $(aws ecr get-login-password --region  
<region-name>) <account-id>.dkr.ecr.<region-name>.amazonaws.com
```

If the repository is not created, then create the repository using the following command:

```
aws ecr create-repository --repository-name <image-name> --region  
<region-name>
```

For example, `aws ecr create-repository --repository-name veritas/flexsnap-datamover --region us-east-2`

5 Run the following commands to push the images to the container registry.

```
$ docker push ${REGISTRY}/netbackup/main:10.1.1  
$ docker push ${REGISTRY}/netbackup/operator:10.1.1  
$ docker push ${REGISTRY}/uss-engine:17.0  
$ docker push ${REGISTRY}/uss-controller:17.0  
$ docker push ${REGISTRY}/uss-mds:17.0  
$ docker push ${REGISTRY}/msdp-operator:17.0  
$ docker push  
${REGISTRY}/veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}  
$ docker push  
${REGISTRY}/veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}
```

- 6 Create a namespace for deploying the NetBackup and MSDP Scaleout operators. These instructions use the default ***netbackup-operator-system*** namespace but a custom namespace is also supported, run:

```
$ kubectl create namespace netbackup-operator-system
```

- 7 Install the MSDP Scaleout operator in the created namespace, using this command. To run this command you must define a full image name in step 3, define a storage class for storing logs from the MSDP operator, and define node selector labels (optional) for scheduling the MSDP operator pod on specific nodes. See “Prerequisites” on page 19.

```
$ kubectl msdp init --image ${REGISTRY}/msdp-operator:17.0  
--storageclass x --namespace netbackup-operator-system -l  
key1=value1
```

- 8 To verify that the MSDP Scaleout operator is running, run:

```
$ kubectl get all --namespace netbackup-operator-system
```

Here, we are using the namespace created in step 5.

The **msdp-operator** pod should show status as *Running*.

- 9 In this step, configure the namespace, image name, and node selector to use for the NetBackup operator image by editing the provided configuration yaml files.
 - (Optional) Perform this step only when using a custom namespace. Edit the file `operator/kustomization.yaml` and change `namespace` to your custom namespace. For example: `namespace: my-custom-namespace`
 - Edit the file `operator/kustomization.yaml` and change **newName** and **newTag**. For example:

```
images:  
  - name: netbackupoperator  
    newName: example.com/netbackup/operator  
    newTag: 'SNAPSHOT_MANAGER_VERSION'
```

- Edit the `operator/patches/operator_patch.yaml` file to add or remove node selectors and toleration that control what nodes Kubernetes may schedule the operator to run on. Use the key value pair same as given during node group creation. For example:

```
nodeSelector:  
  nbpool: nbnodes  
# Support node taints by adding pod tolerations equal to the
```

```

specified nodeSelectors
  # For Toleration NODE_SELECTOR_KEY used as a key and
  NODE_SELECTOR_VALUE as a value.
tolerations:
  - key: nbpool
    operator: "Equal"
    value: nbnodes

```

- 10** Configure the namespace, image name, and node selector to use for NetBackup Snapshot Manager operator image by editing the provided configuration yaml files. Edit the `operator/kustomization.yaml` file and change **newName** and **newTag**. Also change Snapshot Manager operator's node selector and toleration (`CONTROL_NODE_KEY` and `CONTROL_NODE_VALUE`).

The value of `CONTROL_NODE_KEY` and `CONTROL_NODE_VALUE` should match with the value of the fields listed in `operator/patches/operator_patch.yaml` > `nodeSelector (labelKey, labelValue)` and `tolerations (key, value)`, so that the Snapshot Manager operator will also run on the same node as NetBackup operator. For example:

```

images:
- name: cloudpointoperator
  newName: example.com/veritas/flexsnap-deploy
  newTag: 'SNAPSHOT_MANAGER_VERSION'

patches:
- target:
  kind: Deployment
  name: flexsnap-operator
  patch: |
  - op: replace
    path: /spec/template/spec/tolerations/0/key
    value: nbu-control-pool
  - op: replace
    path: /spec/template/spec/tolerations/0/value
    value: nbupool
  - op: replace
    path: /spec/template/spec/affinity/nodeAffinity/requiredDuringSched
    value: nbu-control-pool
  - op: replace
    path: /spec/template/spec/affinity/nodeAffinity/requiredDuringSched
    value: nbupool

```

- 11 To install the NetBackup and Snapshot Manager operator, run the following command from the installer's root directory:

```
$ kubectl apply -k operator
```

- 12 To verify if the operators are running, run:

```
$ kubectl get all --namespace netbackup-operator-system
```

Verify that `pod/netbackup-operator` and `pod/flexsnap-operator` STATUS is showing as *Running*.

Deploying NetBackup and MSDP Scaleout manually

After the operators are deployed, you can deploy the NetBackup and MSDP Scaleout environment.

To deploy NetBackup primary, media, and MSDP Scaleout components:

- 1 Create a Kubernetes namespace where your new NetBackup environment will run. Run the command:

```
kubectl create namespace nb-example
```

Where, *nb-example* is the name of the namespace. The Primary, Media, and MSDP Scaleout application namespace must be different from the one used by the operators. It is recommended to use two namespaces. One for the operators, and a second one for the applications.

- 2 Create a secret to hold the primary server credentials. Those credentials are configured in the NetBackup primary server, and other resources in the NetBackup environment use them to communicate with and configure the primary server. The secret must include fields for `username` and `password`. If you are creating the secret by YAML, the type should be `opaque` or `basic-auth`. For example:

```
apiVersion: v1
  kind: Secret
  metadata:
    name: primary-credentials
    namespace: nb-example
  type: kubernetes.io/basic-auth
  stringData:
    username: nbuser
    password: p@ssw0rd
```

You can also use this command to create a secret.

```
$ kubectl create secret generic primary-credentials --namespace
nb-example --from-literal=username='nbuser'
--from-literal=password='p@ssw0rd'
```

- 3 Create a KMS DB secret to hold Host Master Key ID (`HMKID`), Host Master Key passphrase (`HMKpassphrase`), Key Protection Key ID (`KPKID`), and Key Protection Key passphrase (`KPKpassphrase`) for NetBackup Key Management Service. If creating the secret by YAML, the type should be `_opaque_`. For example:

```
apiVersion: v1
  kind: Secret
  metadata:
    name: example-key-secret
    namespace: nb-example
  type: Opaque
  stringData:
    HMKID: HMKID
    HMKpassphrase: HMKpassphrase
    KPKID: KPKID
    KPKpassphrase: KPKpassphrase
```

You can also create a secret using `kubectl` from the command line:

```
$ kubectl create secret generic example-key-secret --namespace
nb-namespace --from-literal=HMKID="HMKID"
--from-literal=HMKpassphrase="HMKpassphrase"
--from-literal=KPKID="KPKID"
--from-literal=KPKpassphrase="KPKpassphrase"
```

For more details on NetBackup deduplication engine credential rules, see: https://www.veritas.com/content/support/en_US/article.100048511

- 4 Create a secret to hold the MSDP Scaleout credentials for the storage server. The secret must include fields for `username` and `password` and must be located in the same namespace as the Environment resource. If creating the secret by YAML, the type should be `_opaque_` or `_basic-auth_`. For example:

```
apiVersion: v1
  kind: Secret
  metadata:
    name: msdp-secret1
    namespace: nb-example
  type: kubernetes.io/basic-auth
  stringData:
    username: nbuser
    password: p@ssw0rd
```

You can also create a secret using `kubectl` from the command line:

```
$ kubectl create secret generic msdp-secret1 --namespace
nb-example --from-literal=username='nbuser'
--from-literal=password='p@ssw0rd'
```

Note: You can use the same secret for the primary server credentials (from step 2) and the MSDP Scaleout credentials, so the following step is optional. However, to use the primary server secret in an MSDP Scaleout, you must set the `credential.autoDelete` property to *false*. The sample file includes an example of setting the property. The default value is *true*, in which case the secret may be deleted before all parts of the environment have finished using it.

- 5 (Optional) Create a secret to hold the KMS key details. Specify KMS Key only if the KMS Key Group does not already exist and you need to create.

Note: When reusing storage from previous deployment, the KMS Key Group and KMS Key may already exist. In this case, provide KMS Key Group only.

If creating the secret by YAML, the type should be `_opaque_`. For example:

```
apiVersion: v1
  kind: Secret
  metadata:
    name: example-key-secret
    namespace: nb-example
  type: Opaque
  stringData:
    username: nbuser
    passphrase: 'test passphrase'
```

You can also create a secret using `kubectl` from the command line:

```
$ kubectl create secret generic example-key-secret --namespace
nb-example --from-literal=username="nbuser"
--from-literal=passphrase="test passphrase"
```

You may need this key for future data recovery. After you have successfully deployed and saved the key details. It is recommended that you delete this secret and the corresponding key info secret.

- 6 Configure the `samples/environment.yaml` file according to your requirements. This file defines a primary server, media servers, and scale out MSDP Scaleout storage servers. See [“Configuring the `environment.yaml` file”](#) on page 39. for details.
- 7 Apply the environment yaml file, using the same application namespace created in step 1.

```
$ kubectl apply --namespace nb-example --filename environment.yaml
```

Use this command to verify the new environment resource in your cluster:

```
$ kubectl get --namespace nb-example environments
```

The output should look like:

```
NAME                AGE
environment-sample  2m
```

After a few minutes, NetBackup finishes starting up on the primary server, and then the media servers and MSDP Scaleout storage servers you configured in the environment resource start appearing. Run:

```
$ kubectl get --namespace nb-example  
all,environments,primaryservers,mediaservers,msdpscaleouts
```

The output should show:

- All pod status as Ready and Running

NAME	READY	STATUS
pod/dedupel-uss-controller-	1/1	Running
pod/dedupel-uss-mds-1	1/1	Running

- For `msdpscaleout` `SIZE = READY`, for example: 4=4.

NAME	SIZE	READY
<code>msdpscaleout.msdp.veritas.com/dedupel</code>	4	4

- `environment.netbackup` should show `STATUS` as `Success`

NAME	STATUS
<code>environment.netbackup.veritas.com/environment-sample</code>	Success

- 8 To start using your newly deployed environment sign-in to NetBackup web UI. Open a web browser and navigate to `https://<primaryserver>/webui/login` URL.

The primary server is the host name or IP address of the NetBackup primary server.

You can retrieve the primary server's hostname by using the command:

```
$ kubectl describe primaryserver.netbackup.veritas.com/<primary  
server CR name>--namespace <namespace_name>
```

Refer to **Deploying MSDP Scaleout** from the guide [NetBackup™ Deployment Guide for Amazon Elastic Kubernetes Services \(EKS\) Cluster](#)

Deploying NetBackup and Snapshot Manager manually

After the operators are deployed as mentioned in the following section, you can deploy the NetBackup and Snapshot Manager environment:

See [“Deploying the operators manually”](#) on page 21.

To deploy NetBackup primary, media and Snapshot Manager components

- 1 Create a Kubernetes namespace where your new NetBackup environment will run. Run the following command:

```
kubectl create namespace nb-example
```

Where, *nb-example* is the name of the namespace. The Primary, Media, and Snapshot Manager application namespace must be different from the one used by the operators. It is recommended to use two namespaces. One for the operators, and a second one for the applications.

- 2 Create a secret to hold the Snapshot Manager credentials. The secret must include fields for **username** and **password**. If you are creating the secret by YAML, the type should be opaque or basic-auth.

For example:

```
apiVersion: v1
stringData:
  password: p@ssw0rd
  username: cpuser
kind: Secret
metadata:
  name: cp-creds
  namespace: ns-155
type: Opaque
```

You can also use this command to create a secret.

```
kubectl create secret generic cp-credentials --namespace
nb-example --from-literal=username='cpuser'
--from-literal=password='p@ssw0rd'
```

- 3** Configure the `samples/environment.yaml` file according to your requirements. This file defines a primary server, media servers, and Snapshot Manager servers. See [“Configuring the `environment.yaml` file”](#) on page 39. for details.

- 4 Apply the `environment.yaml` file, using the same application namespace created in the above step.

```
kubectl apply --namespace nb-example --filename environment.yaml
```

Use this command to verify the new environment resource in your cluster:

```
kubectl get --namespace nb-example environments
```

After a few minutes, NetBackup finishes starting up the primary server, media servers and Snapshot Manager servers in the sequence that you configured in the environment resource. Snapshot Manager is registered with NetBackup and cloud provider is configured automatically.

Run the following command:

```
kubectl get --namespace netbackup-environment  
all,environments,primaryservers,cpservers,mediaservers
```

The output would be displayed as follows:

```
$ kubectl get all,environments,primaryservers,mediaservers,  
cpservers,msdpscaleouts -n netbackup-environment
```

NAME	READY	STATUS	REST.
pod/flexsnap-agent			
-33e649abd383410ea618751f7b2eb8ae-598b8b747-957xd	1/1	Running	0
pod/flexsnap-agent			
-688d478bc8-hxgrk	1/1	Running	0
pod/flexsnap-api-gateway			
-69cbbfc844-f6whb	1/1	Running	0
pod/flexsnap-certauth			
-6f65894b69-njfhf	1/1	Running	0
pod/flexsnap-coordinator			
-749649c7-fgs4t	1/1	Running	0
pod/flexsnap-fluentd-collector			
-7445f6fb9f-bqfqq	1/1	Running	0
pod/flexsnap-fluentd-csvfz	1/1	Running	0
pod/flexsnap-fluentd-ht7w8	1/1	Running	0
pod/flexsnap-fluentd-lkdgt	1/1	Running	0
pod/flexsnap-fluentd-rzvrv	1/1	Running	0
pod/flexsnap-fluentd-spgkc	1/1	Running	0
pod/flexsnap-listener-664674-phjnd	1/1	Running	0
pod/flexsnap-mongodb-f6b744df5-p4hfv	1/1	Running	0
pod/flexsnap-nginx-8647f57db8-rzkt5	1/1	Running	0
pod/flexsnap-notification-7db95868f5-dpx7z	1/1	Running	0
pod/flexsnap-rabbitmq-0	1/1	Running	0
pod/flexsnap-scheduler-68d8b75d75-5q4fk	1/1	Running	0

pod/nbux-marketplace-10-239-207-44. vxindia.veritas.com	1/1	Running	0
pod/nbux-marketplace-10-239-207-45. vxindia.veritas.com	2/2	Running	0
pod/nbux-marketplace-10-239-207-46. vxindia.veritas.com	2/2	Running	0
pod/nbux-marketplace-10-239-207-47. vxindia.veritas.com	2/2	Running	0
pod/nbux-eks-dedupel-uss-agent-56r7h	1/1	Running	0
pod/nbux-eks-dedupel-uss-agent-hvfw7	1/1	Running	0
pod/nbux-eks-dedupel-uss-agent-jx46x	1/1	Running	0
pod/nbux-eks-dedupel-uss-agent-pz7w8	1/1	Running	0
pod/nbux-eks-dedupel-uss-agent-r2kkt	1/1	Running	0
pod/nbux-eks-dedupel-uss-agent-vx8gc	1/1	Running	0
pod/nbux-eks-dedupel-uss-controller-0	1/1	Running	0
pod/nbux-eks-dedupel-uss-controller-1	1/1	Running	0
pod/nbux-eks-dedupel-uss-mds-1	1/1	Running	0
pod/nbux-eks-dedupel-uss-mds-2	1/1	Running	0
pod/nbux-eks-dedupel-uss-mds-3	1/1	Running	0
pod/nbux-eks-medial-media-0	1/1	Running	0
pod/nbux-eks-primary-0	1/1	Running	0

NAME	TYPE	CLUSTER-IP	EXTERNAL
service/flexsnap-api-gateway	ClusterIP	172.20.92.70	<none>
service/flexsnap-certauth	ClusterIP	172.20.222.22	<none>
service/flexsnap-fluentd-service	ClusterIP	172.20.141.61	<none>
service/flexsnap-mongodb	ClusterIP	172.20.157.102	<none>
service/flexsnap-nginx	LoadBalancer	172.20.187.1	
		nbux-eks-cpserver-1-2b677a3f6b6ffe48.elb.us-west-2	443:31318/TCP
service/flexsnap-rabbitmq	ClusterIP	172.20.33.99	<none>
service/ip-10-239-207-44-host-nbux-marketplace-10-239-207-44-vxindia-ve	LoadBalancer	172.20.186.216	k8s-ns155-ip102392-8d0152d6a0-8d77c9a84d1
elb.us-west-2.amazonaws.com			10082:30397/TCP,10102:31873/TCP,10086:32374
			443:30732/TCP,111:30721/TCP,662:32206/TCP,875:32361/TCP,892:31540/TCP,204
			45209:31944/TCP,58329:30149/TCP,139:31587/TCP,445:31252/TCP 73m
service/ip-10-239-207-45-host-nbux-marketplace-10-239-207-45-vxindia-ve	LoadBalancer	172.20.23.36	k8s-ns155-ip102392-410db5113c-791da4601d5
elb.us-west-2.amazonaws.com			10082:31116/TCP,10102:31904/TCP,10086:32468
			443:32693/TCP,111:32658/TCP,662:31151/TCP,875:30175/TCP,892:
			31126/TCP,2049:31632/TCP,45209:32602/TCP,58329:31082/TCP,139:31800/TCP,44
			73m
service/ip-10-239-207-46-host-nbux-marketplace-10-239-207-46-vxindia-ve			

```

LoadBalancer 172.20.6.179 k8s-ns155-ip102392-1c160eed54-c975ee450ed
elb.us-west-2.amazonaws.com 10082:31927/TCP,10102:31309/TCP,10086:30285
443:32648/TCP,111:32348/TCP,662:32170/TCP,875:31854/TCP,892:30842/TCP,204
45209:30002/TCP,58329:32408/TCP,139:30882/TCP,445:32017/TCP 73m
service/ip-10-239-207-47-host-nbux-marketplace-10-239-207-47-vxindia-ve
LoadBalancer 172.20.221.124 k8s-ns155-ip102392-bb1f5b1cbe-22e4275c1af
elb.us-west-2.amazonaws.com 10082:30137/TCP,10102:30727/TCP,10086:32649
443:30474/TCP,111:32690/TCP,662:31740/TCP,875:30437/TCP,892:32532/TCP,204
45209:31259/TCP,58329:31070/TCP,139:32393/TCP,445:30296/TCP 73m
service/nbux-eks
-dedupel-uss-controller ClusterIP 172.20.197.231 <none> 10
service/nbux-eks
-dedupel-uss-mds ClusterIP None <none> 23
service/nbux-eks
-dedupel-uss-mds-client ClusterIP 172.20.226.45 <none> 23
service/nbux-eks
-medial-media-0 LoadBalancer 172.20.146.140
nbux-eks-medial-media-0-1a525c3f2587c8cf.elb.us-west-2.amazonaws.com
13782:31414/TCP
service/nbux-eks-primary LoadBalancer 172.20.108.104
nbux-eks-primary-c0f9de7ce7e231cd.elb.us-west-2.amazonaws.com
13781:30252/TCP,13782:31446/TCP,1556:31033/TCP,443:31517/TCP,8443:32237/
TCP,22:30

```

NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	NODE
daemonset.apps/ flexsnap-fluentd	5	5	5	5	5	<n
daemonset.apps/ nbux-eks- dedupe1-uss-agent	6	6	6	6	6	ag =m

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
deployment.apps/ flexsnap-agent	1/1	1	1	43m
deployment.apps/ flexsnap-agent- 33e649abd383410ea618751f7b2eb8ae	1/1	1	1	28m
deployment.apps/ flexsnap-api-gateway	1/1	1	1	43m
deployment.apps/ flexsnap-certauth	1/1	1	1	44m
deployment.apps/ flexsnap-coordinator	1/1	1	1	43m
deployment.apps/ flexsnap-coordinator	1/1	1	1	43m

flexsnap-fluentd-collector	1/1	1	1	43m
deployment.apps/				
flexsnap-listener	1/1	1	1	43m
deployment.apps/				
flexsnap-mongodb	1/1	1	1	43m
deployment.apps/				
flexsnap-nginx	1/1	1	1	43m
deployment.apps/				
flexsnap-notification	1/1	1	1	43m
deployment.apps/				
flexsnap-scheduler	1/1	1	1	43m
NAME		DESIRED	CURRENT	READY
replicaset.apps/				
flexsnap-agent-				
33e649abd383410ea618751f7b2eb8ae-598b8b747		1	1	1
replicaset.apps/				
flexsnap-agent-688d478bc8		1	1	1
replicaset.apps/				
flexsnap-api-gateway-69cbbfc844		1	1	1
replicaset.apps/				
flexsnap-certauth-6f65894b69		1	1	1
replicaset.apps/				
flexsnap-coordinator-749649c7		1	1	1
replicaset.apps/				
flexsnap-fluentd-collector-7445f6fb9f		1	1	1
replicaset.apps/				
flexsnap-listener-664674		1	1	1
replicaset.apps/				
flexsnap-mongodb-f6b744df5		1	1	1
replicaset.apps/				
flexsnap-nginx-8647f57db8		1	1	1
replicaset.apps/				
flexsnap-notification-7db95868f5		1	1	1
replicaset.apps/				
flexsnap-scheduler-68d8b75d75		1	1	1
NAME	READY	AGE		
statefulset.apps/				
flexsnap-rabbitmq	1/1	43m		
statefulset.apps/				
nbux-eks-dedupe1-uss-controller	2/2	73m		
statefulset.apps/				
nbux-eks-medial-media	1/1	55m		
statefulset.apps/				

```

nbux-eks-primary          1/1      101m
NAME                      READY    AGE      STATUS
environment.netbackup.
veritas.com/nbux-eks      4/4      102m     Success
NAME                      TAG      AGE      STATUS
primaryserver.netbackup.
veritas.com/nbux-eks      10.1.1.0085  102m     Success
NAME                      TAG      AGE      PRIMARY SERVER
mediaserver.netbackup.
veritas.com/nbux-eks-medial  10.1.1.0085  56m     nbux-marketplace
                                -10-239-207-42.vxindia.veritas.
NAME                      TAG      AGE      STATUS
cpserver.netbackup.
veritas.com/nbux-eks-cpserver-1  10.1.1.0.1073  44m     Success
NAME                      AGE     TAG      SIZE     READY
msdp*scaleout.msdp.
veritas.com/nbux-eks-dedupe1  75m    17.1.0085  4        4
    
```

Configuring the `environment.yaml` file

The following configurations apply to all the components:

Table 2-2 Common environment parameters

Parameter	Description
name: environment-sample	Specify the name of the environment in your cluster.
namespace: example-ns	Specify the namespace where all the NetBackup resources are managed. If not specified here, then it will be the current namespace when you run the command <code>kubectl apply -f</code> on this file.
containerRegistry: example-registry.com	Specify a container registry that the cluster has access. NetBackup images are pushed to this registry.

Table 2-2 Common environment parameters (*continued*)

Parameter	Description
tag: 10.1.1	This tag is used for all images in the environment. Specifying a `tag` value on a sub-resource affects the images for that sub-resource only. For example, if you apply an EEB that affects only primary servers, you might set the `primary.tag` to the custom tag of that EEB. The primary server runs with that image, but the media servers and MSDP scaleouts continue to run images tagged `10.1.1`. Beware that the values that look like numbers are treated as numbers in YAML even though this field needs to be a string; quote this to avoid misinterpretation.
licenseKeys:	List the license keys that are shared among all the sub-resources. Licenses specified in a sub-resource are appended to this list and applied only to the sub-resource.
paused: false	Specify whether the NetBackup operator attempts to reconcile the differences between this YAML specification and the current Kubernetes cluster state. Only set it to true during maintenance.
configCheckMode: default	This controls whether certain configuration restrictions are checked or enforced during setup. Other allowed values are skip and dryrun.
corePattern: /corefiles/core.%e.%p.%t	Specify the path to use for storing core files in case of a crash.
loadBalancerAnnotations: service.beta.kubernetes.io/aws-load-balancer-target-group-attributes example-subnet1 name	Specify the annotations to be added for the network load balancer

Note: If NetBackup is upgraded from 10.0.0.1, then delete the following configuration from the `environment.yaml` file from **spec.loadBalancerAnnotations** section:

```
service.beta.kubernetes.io/aws-load-balancer-target-group-attributes:
preserve_client_ip.enabled=true
```

The following section describes Snapshot Manager related parameters. You may also deploy without any Snapshot Manager. In that case, remove the `cpServer` section entirely from the configuration file.

Table 2-3 Snapshot Manager parameters

Parameter	Description
cpServer:	This specifies Snapshot Manager configurations.
-name	Currently only single instance of Snapshot Manager deployment is supported. It is also possible to have no Snapshot Managers configured; in this case, delete the cpServer section itself.
containerRegistry	(Optional) Specify a container registry that the cluster has access. Snapshot Manager images are pushed to this registry which overrides the one defined in Common environment parameters table above.
tag:	This tag overrides the one defined in Common environment parameters table above. The Snapshot Manager images are shipped with tags different from the NetBackup primary, media, and MSDP images.
credential:secretName	This defines the credentials for Snapshot Manager. It refers to a secret in the same namespace as this environment resource with values for username and password.
networkLoadBalancer: annotations	Annotations to be provided to the network load balancer. All networkLoadBalancer annotations are supported. These values are merged with the values provided in the loadBalancerAnnotations . The duplicate values provided here, override the corresponding values in the loadBalancerAnnotations .
networkLoadBalancer: ipaddr	IP address to be assigned to the network load balancer.
networkLoadBalancer: fqdn	FQDN to be assigned to the network load balancer.
log.capacity	Size for log volume.
log.storageClassName	Storage class for log volume. It must be EFS based storage class.
data.capacity	Size for data volume.
data.storageClassName	EBS based storage class for data volume.
controlPlane.nodePool	Name of the control plane node pool.
controlPlane.labelKey	Label and taint key of the control plane.
controlPlane.labeValue	Label and taint value of the control plane.

Table 2-3 Snapshot Manager parameters (*continued*)

Parameter	Description
<code>dataPlane.nodePool</code>	Name of the data plane node pool.
<code>dataPlane.labelKey</code>	Label and taint key of the data plane.
<code>dataPlane.labelValue</code>	Label and taint value of the data plane.
<code>proxySettings.vx_http_proxy</code> :	Address to be used as the proxy for all HTTP connections. For example, "http://proxy.example.com:8080/"
<code>proxySettings.vx_https_proxy</code> :	Address to be used as the proxy for all HTTPS connections. For example, "http://proxy.example.com:8080/"
<code>proxySettings.vx_no_proxy</code> :	Address that are allowed to bypass the proxy server. You can specify host name, IP addresses and domain names in this parameter. For example, "localhost,mycompany.com,169.254.169.254"

The following configurations apply to the primary server. The values specified in the following table can override the values specified in the table above.

Table 2-4 Environment parameters for the primary server

Paragraph	Description
<code>paused: false</code>	Specifies whether the NetBackup operator attempts to reconcile the differences between this YAML specification and the current Kubernetes cluster state. Set it to <i>true</i> only during maintenance. This applies only to the environment object. To pause reconciliation of the managed primary server, for example, you must set <code>spec.primary.paused</code> . Setting <code>spec.paused:true</code> ceases updates to the managed resources, including updates to their <code>paused</code> status. Entries in the media servers and MSDP scaleouts lists also support the <code>paused</code> field. The default value is <i>false</i> .
<code>primary</code>	Specifies attributes specific to the primary server resources. Every environment has exactly one primary server, so this section cannot be left blank.

Table 2-4 Environment parameters for the primary server (*continued*)

Paragraph	Description
name: primary-name	Set resourceNamePrefix to control the name of the primary server. The default value is the same as the environment's name.
tag: 10.1.1-special	To use a different image tag specifically for the primary server, uncomment this value and provide the desired tag. This overrides the tag specified in the common section.
nodeSelector: labelKey: kubernetes.io/os labelValue: linux	Specify a key and value that identifies nodes where the primary server pod runs. Note: This labelKey and labelValue must be the same label key:value pair used during node group creation which would be used as a toleration for primary server.
networkLoadBalancer: annotations: service.beta.kubernetes.io/aws-load-balancer-subnets: example-subnet1 name ipList: - ipAddr: 4.3.2.1 fqdn: primary.example.com	Uncomment the annotations to specify additional primary server-specific annotations. These values are merged with the values given in the loadBalancerAnnotations above. Any duplicate values given here override the corresponding values above. Next, specify the hostname and IP address of the primary server.
credSecretName: primary-credential-secret	This determines the credentials for the primary server. Media servers use these credentials to register themselves with the primary server.
itAnalyticsPublicKey: ssh-rsaxxx	If using NetBackup IT Analytics, uncomment this and provide the SSH public key. IT Analytics uses this to access the primary server.

Table 2-4 Environment parameters for the primary server (*continued*)

Paragraph	Description
kmsDBSecret: kms-secret	Secret name which contains the Host Master Key ID (HMKID), Host Master Key passphrase (HMKpassphrase), Key Protection Key ID (KPKID) and Key Protection Key passphrase (KPKpassphrase) for NetBackup Key Management Service. The secret should be 'Opaque', and can be created either using a YAML or the following example command: <pre>kubectl create secret generic kms-secret --namespace nb-namespace --from-literal=HMKID="HMK@ID" --from-literal=HMKpassphrase="HMK@passphrase" --from-literal=KPKID="KPK@ID" --from-literal=KPKpassphrase="KPK@passphrase"</pre>
licenseKeys:	To specify additional license keys that are applied only to the primary server, uncomment this and provide the license key(s). In this example, the primary server would have the "X" license key defined in the previous section, followed by this "Y" key.
catalog: capacity: 100Gi storageClassName: <EFS_ID>	This storage applies to the primary server for the NetBackup catalog, log and data volumes. The primary server catalog volume must be at least 100 Gi.
log: capacity: 30Gi storageClassName: <EBS based storage class>	Log volume must be at least 30Gi.
data: capacity: 30Gi storageClassName: <EBS based storage class>	The primary server data volume must be at least 30Gi.

The following section describes the media server configurations. If you do not have a media server either remove this section from the configuration file entirely, or define it as an empty list.

Note: The environment name or media server name in `environment.yaml` file must always be less than 22 characters.

Table 2-5 Media server related parameters

parameters	Description
mediaServers: - name: media1	This specifies media server configurations. This is given as a list of media servers, but most environments will have just one, with multiple replicas. It's also possible to have zero media servers; in that case, either remove the media servers section entirely, or define it as an empty list: <code>mediaServers: []</code>
replicas: 1	Specifies the number of replicas of this media server. Minimum number of supported replicas is 1.
tag: <i>10.1.1-special</i>	To use a different image tag specifically for the media servers, uncomment this value and provide the desired tag. This overrides the tag specified above in the common table.
nodeSelector: labelKey: <i>kubernetes.io/os</i> labelValue: <i>linux</i>	Specify a key and value that identifies nodes where media-server pods will run. Note: This labelKey and labelValue must be the same label key:value pair used during node group creation which would be used as a toleration for media server.
data: capacity: <i>50Gi</i> storageClassName: <i><EBS based storage class></i>	This storage applies to the media server data volumes. The minimum data size for a media server is 50 Gi.
log capacity: <i>30Gi</i> storageClassName: <i><EBS based storage class></i>	This storage applies to the media server log volumes. Log volumes must be at least 30Gi.

Table 2-5 Media server related parameters (*continued*)

parameters	Description
networkLoadBalancer: <i>annotations:</i> -service.beta.kubernetes.io/aws-load-balancer-subnets: <i>example-subnet1 name</i> ipList: ipAddr: 4.3.2.2 fqdn: <i>media1-1.example.com</i> ipAddr: 4.3.2.3 fqdn: <i>media1-2.example.com</i>	Uncomment annotations to specify additional media-server specific annotations. These values are merged with the values given in the loadBalancerAnnotations. The duplicate values given here, override the corresponding values in the loadBalancerAnnotations. The number of entries in the IP list should match the replica count specified above.

Note the following:

To use gp3 (EBS based storage class), user must specify provisioner for storage class as `ebs.csi.aws.com` and must install EBS CSI driver. For more information on installing the EBS CSI driver, see [Amazon EBS CSI driver](#). Example, for gp3 storage class:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: gp3
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
allowVolumeExpansion: true
provisioner: ebs.csi.aws.com
volumeBindingMode: WaitForFirstConsumer
parameters:
  type: gp3
```

The following section describes MSDP-related parameters. You may also deploy without any MSDP scaleouts. In that case, remove the `msdpScaleouts` section entirely from the configuration file.

Table 2-6 MSDP Scaleout related parameters

Parameter	Description
msdpScaleouts: - name: dedupe1	This specifies MSDP Scaleout configurations. This is given as a list, but it would be rare to need more than one scaleout deployment in a single environment. Use the `replicas` property below to scale out. It's also possible to have zero MSDP scaleouts; in that case, either remove the msdpScaleouts section entirely, or define it to an empty list: msdpScaleouts: []
tag: '17.0'	This tag overrides the one defined in the table 1-3. It is necessary because the MSDP Scaleout images are shipped with tags different from the NetBackup primary and media images.
replicas: 4	This is the scaleout size of this MSDP Scaleout component. It is a required value, and it must be between 4 and 16 inclusive. Note: Scale-down of the MSDP Scaleout replicas after deployment is not supported.
serviceIPFQDNs: ipAddr: 1.2.3.4 fqdn: dedupe1-1.example.com ipAddr: 1.2.3.5 fqdn: dedupe1-2.example.com ipAddr: 1.2.3.6 fqdn: dedupe1-3.example.com ipAddr: 1.2.3.7 fqdn: dedupe1-4.example.com	These are the IP addresses and host names of the MSDP Scaleout servers. The number of the entries should match the number of the replicas specified above.
kms: keyGroup: <i>example-key-group</i>	Specifies the initial key group and key secret to be used for KMS encryption. When reusing storage from a previous deployment, the key group and key secret may already exist. In this case, provide the keyGroup only.

Table 2-6 MSDP Scaleout related parameters (*continued*)

Parameter	Description
keySecret: <i>example-key-secret</i>	Specify keySecret only if the key group does not already exist and needs to be created. The secret type should be Opaque, and you can create the secret either using a YAML or the following command: <pre>kubectl create secret generic example-key-secret --namespace nb-namespace --from-literal=username="devuser" --from-literal=password="test password"</pre>
loadBalancerAnnotations: service.beta.kubernetes.io/aws-load-balancer-internal: <i>true</i>	For MSDP scaleouts, the default value for the AWS-load-balancer-internal annotation is <code>false</code> , which may cause the MSDP Scaleout services in this Environment to be accessible publicly. To make sure that they use private IP addresses, specify <code>true</code> here or in the loadBalancerAnnotations above in Table 1-3.
credential: secretName: <i>msdp-secret1</i>	This defines the credentials for the MSDP Scaleout server. It refers to a secret in the same namespace as this environment resource. Secret can be either of type 'Basic-auth' or 'Opaque'. You can create secrets using a YAML or by using the following command: <pre>kubectl create secret generic <msdp-secret1> --namespace <nb-namespace> --from-literal=username=<"devuser"> --from-literal=password=<"Y@123abCdEf"></pre>
autoDelete: <i>false</i>	Optional parameter. Default value is true. When set to true, the MSDP Scaleout operator deletes the MSDP secret after using it. In such case, the MSDP and primary secrets must be distinct. To use the same secret for both MSDP scaleouts and the primary server, set autoDelete to false.

Table 2-6 MSDP Scaleout related parameters (*continued*)

Parameter	Description
catalog: capacity: <i>1Gi</i> storageClassName: <i>gp2</i>	This storage applies to MSDP Scaleout to store the catalog and metadata. The catalog size may only be increased for capacity expansion. Expanding the existing catalog volumes cause short downtime of the engines. Recommended size is 1/100 of backend data capacity.
dataVolumes: capacity: <i>5Gi</i> storageClassName: <i>gp2</i>	This specifies the data storage for this MSDP Scaleout resource. You may increase the size of a volume or add more volumes to the end of the list, but do not remove or re-order volumes. Maximum 16 volumes are allowed. Appending new data volumes or expanding existing ones will cause short downtime of the Engines. Recommended volume size is 5Gi-32Ti.
log: capacity: <i>20Gi</i> storageClassName: <i>gp2</i>	Specifies log volume size used to provision Persistent Volume Claim for Controller and MDS Pods. In most cases, 5-10 Gi capacity should be big enough for one MDS or Controller Pod to use.
nodeSelector: labelKey: <i>kubernetes.io/os</i> labelValue: <i>linux</i>	Specify a key and value that identifies nodes where MSDP Scaleout pods will run.

For more information on Snapshot Manager related parameters, refer to the following:

See [“Installing the docker images”](#) on page 104.

Edit restricted parameters post deployment

Do not change these parameters post initial deployment. Changing these parameters may result in an inconsistent deployment.

Table 2-7 Edit restricted parameters post deployment

Parameter	Description
name	Specifies the prefix name for the primary, media, and MSDP Scaleout server resources.

Table 2-7 Edit restricted parameters post deployment (*continued*)

Parameter	Description
ipAddr, fqdn and loadBalancerAnnotations	<p>The values against ipAddr, fqdn and loadBalancerAnnotations against following fields should not be changed post initial deployment. This is applicable for primary, media, and MSDP Scaleout servers. For example:</p> <ul style="list-style-type: none"> - The loadBalancerAnnotations for loadBalancerAnnotation service.beta.kubernetes.io/aws-load-balancer -internal: <pre>example-subnet service.beta.kubernetes.io/ aws-load-balancer -internal: ##~#apos;true##~#apos;</pre> - The IP and FQDNs values defined for Primary, Media and MSDPScaleout ipList: <pre>- ipAddr: 4.3.2.1 fqdn: primary.example.com ipList: - ipAddr: 4.3.2.2 fqdn: medial-1.example.com - ipAddr: 4.3.2.3 fqdn: medial-2.example.com serviceIPFQDNs: - ipAddr: 1.2.3.4 fqdn: dedupe1-1.example.com - ipAddr: 1.2.3.5 fqdn: dedupe1-2.example.com - ipAddr: 1.2.3.6 fqdn: dedupe1-3.example.com fqdn: dedupe1-4.example.com</pre>

Table 2-8 Snapshot Manager server related parameters

parameters	Description
cpServer: - name: cpServer-name	This specifies Snapshot Manager server configurations. This is given as a list of Snapshot Manager servers, but most environments will have just one, with multiple replicas.

Table 2-8 Snapshot Manager server related parameters (*continued*)

parameters	Description
tag: <i>10.1.1-special</i>	To use a different containerRegistry specifically for the Snapshot Manager server, uncomment this value and provide the desired containerRegistry. This overrides the containerRegistry specified above.
nodeSelector: controlPlane: <i>cpcontrol</i> dataPlane: <i>cpdata</i>	Details of the label to be used for identification of Kubernetes nodes reserved for the Snapshot Manager Servers. If controlPlane is not specified here it will read it from <i>primary.nodeSelector</i> . In that case the nodepool should have appropriate taint and label added to it.[the nodepool name mentioned will have value of <i>primary.nodeSelector.labelValue</i>] Note: The nodepool name mentioned will have value of <i>primary.nodeSelector.labelValue</i> .
data: capacity: <i>100Gi</i> storageClassName: <i><EBS based storage class></i>	This storage applies to the Snapshot Manager server data volumes. The minimum data size for a Snapshot Manager server is 100 Gi.
log capacity: <i>5Gi</i> storageClassName: <i><EBS based storage class></i>	This storage applies to the Snapshot Manager server log volumes. Log volumes must be at least 5 Gi.
networkLoadBalancer: annotations: -service.beta.kubernetes.io/aws-load-balancer-subnets: example-subnet1 name ipList: ipAddr: <i>4.3.2.2</i> fqdn: <i>media1-1.example.com</i> ipAddr: <i>4.3.2.3</i> fqdn: <i>media1-2.example.com</i>	Snapshot ManagerUncomment annotations to specify additional -server specific annotations. These values are merged with the values given in the spec.loadBalancerAnnotations. The duplicate values given here, override the corresponding values in the spec.loadBalancerAnnotations. The number of entries in the IP list should match the replica count specified above.

Table 2-8 Snapshot Manager server related parameters (*continued*)

parameters	Description
credential: secretName: <i>cp-creds</i>	This defines the credentials for the Snapshot Manager server. It refers to a secret in the same namespace as this environment resource. The secret name can be created using a YAML or the following example command: <pre># kubectl create secret generic cp-creds --namespace nb-namespace --from-literal=username="admin" --from-literal=password="CloudPoint@123"</pre>

Note the following:

To use `efs-sc` (EFS based storage class), user must specify provisioner for storage class as `efs.csi.aws.com` and must install EFS CSI driver. For example:

```
Storage class:

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: efs-sc
provisioner: efs.csi.aws.com
parameters:
  provisioningMode: efs-ap
  fileSystemId: <EFS ID>
  directoryPerms: "700"
reclaimPolicy: Retain
volumeBindingMode: Immediate
```

Uninstalling NetBackup environment and the operators

You can uninstall the NetBackup primary, media, and MSDP Scaleout environment and the operators as required. You need to uninstall the NetBackup environment before you uninstall the operators.

Note: Replace the environment custom resource names as per your configuration in the steps below.

To uninstall the NetBackup environment

- 1 To remove the environment components from the application namespace, run:

```
$ kubectl delete
environment.netbackup.veritas.com/environment-sample --namespace
<namespce_name>
```

- 2 Wait for all the pods, services and resources to be terminated. To confirm, run

```
$ kubectl get --namespace <namespce_name>
all,environments,primaryservers,mediaservers,msdpscaleouts
```

You should get a message that no resources were found in the *nb-example* namespace.

- 3 To identify and delete any outstanding persistent volume claims, run the following:

```
$ kubectl get pvc --namespace <namespce_name>
$ kubectl delete pvc <pvc-name>
```

- 4 To locate and delete any persistent volumes created by the deployment, run:

```
$ kubectl get pv
$ kubectl delete pv <pv-name> --grace-period=0 --force
```

Note: Certain storage drivers may cause physical volumes to get stuck in the terminating state. To resolve this issue, remove the finalizer, using the command: `$ kubectl patch pv <pv-name> -p '{"metadata":{"finalizers":null}}'`

Note: Navigate to mounted EFS directory and delete the content from `primary_catalog` folder by running the `rm -rf /efs/` command.

- 5 To delete the application namespace, run:

```
$ kubectl delete ns <namespace name>
```

To uninstall the operators

- 1 To uninstall the NetBackup operator run the following command from the installation directory.

```
$ kubectl delete -k operator
```

- 2 To uninstall the MSDP Scaleout operator and remove the operator's namespace, run.

```
$ kubectl msdp delete --namespace <namespace name>
```

Note: Do not remove the MSDP Scaleout operator first as it may corrupt the NetBackup operator.

To uninstall NetBackup operator and Snapshot Manager

- ◆ To uninstall the NetBackup operator and Snapshot Manager operator and remove the operator's namespace, run the following command:

```
$ kubectl delete -k operator
```

For more information on uninstalling the Snapshot Manager, refer to the following section:

See [“Uninstalling Snapshot Manager from EKS”](#) on page 173.

Applying security patches

This section describes how to apply security patches for operator and application images.

In the instructions below, we assume that the operators were deployed to the *netbackup-operator-system* namespace (the default namespace suggested by the deployment script), and that an environment resource named *nb-env* was deployed to a namespace named *nb-example*.

Although it is not necessary to manually shut down NetBackup primary server or media servers, it's still a good idea to quiesce scheduling so that no jobs get interrupted while pods are taken down and restarted.

Prepare the images

To prepare the images to apply patches

- 1 Unpack the tar file on a system where docker is able to push to the container registry, and `kubectl` can access the cluster.
- 2 Decide on a unique tag value to use for MSDP Scaleout images. The unique tag should be in `version-postfix` format, For example, `17.0-update1`. Set the **DD_TAG** environment variable accordingly and run `deploy.sh`:

```
DD_TAG=17.0-update1 ./deploy.sh
```

- 3 In the menu that appears, select option **1** to install the operators.
- 4 Enter the fully qualified domain name of the container registry.
For example: `example.dkr.ecr.us-east-2.amazonaws.com/`.
When the script prompts to load images, answer `yes`.
- 5 When the script prompts to tag and push images, wait. Open another terminal window and re-tag the MSDP Scaleout images as:

```
docker tag msdp-operator:17.0 msdp-operator:17.0-update1  
  
docker tag uss-controller:17.0 uss-controller:17.0-update1  
  
docker tag uss-engine:17.0 uss-engine:17.0-update1  
  
docker tag uss-mds:17.0 uss-mds:17.0-update1
```

- 6 Return to the deploy script and when prompted, enter `yes` to tag and push the images. Wait for the images to be pushed, and then the script will pause to ask another question. The remaining questions are not required, so press `Ctrl+c` to exit the deploy script.

Update the NetBackup operator

- 1 Get the image ID of the existing NetBackup operator container and record it for later. Run:

```
kubectl get pod -n netbackup-operator-system -l  
nb-control-plane=nb-controller-manager -o jsonpath --template  
"{.items[*].status.containerStatuses[?(@.name=='netbackup-operator')].imageID}{'\n'}"
```

The command prints the name of the image and includes the SHA-256 hash identifying the image. For example:

```
example.dkr.ecr.us-east-2.amazonaws.com/
```

- 2 To restart the NetBackup operator, run:

```
pod=$(kubectl get pod -n netbackup-operator-system -l  
nb-control-plane=nb-controller-manager -o jsonpath --template  
'{.items[*].metadata.name}')  
  
kubectl delete pod -n netbackup-operator-system $pod
```

- 3 Re-run the `kubectl` command from earlier to get the image ID of the NetBackup operator. Confirm that it's different from what it was before the update.

Update the MSDP Scaleout operator

- 1 Get the image ID of the existing MSDP Scaleout operator container and save it for later use. Run:

```
kubectl get pods -n netbackup-operator-system -l  
control-plane=controller-manager -o jsonpath --template  
"{.items[*].status.containerStatuses[?(@.name=='manager')].imageID}{'\n'}"
```

- 2 Re-initialize the MSDP Scaleout operator using the new image.

```
kubectl msdp init -n netbackup-operator-system --image <account  
id>.dkr.ecr.<region>.amazonaws.com/<registry>:<tag>/msdp-operator:17.0-update1
```

- 3 Re-run the `kubectl` command from earlier to get the image ID of the MSDP Scaleout operator. Confirm that it's different from what it was before the update.

Update the primary server or media servers

- 1 Look at the list of pods in the application namespace and identify the pod or pods to update. The primary-server pod's name typically end with "primary-0" and media-server pods end with "media-0", "media-1", etc. Hereafter, pod will be referred to as \$pod. Run:

```
kubectl get pods -n nb-example
```

- 2 Get the image ID of the existing NetBackup container and record it for later. Run:

```
kubectl get pods -n nb-example $pod -o jsonpath --template  
"{.status.containerStatuses[*].imageID}{'\n'}"
```

- 3 Look at the list of StatefulSets in the application namespace and identify the one that corresponds to the pod or pods to be updated. The name is typically the same as the pod, but without the number at the end. For example, a pod named nb-primary-0 is associated with statefulset nb-primary. Hereafter the statefulset will be referred to as \$set. Run:

```
kubectl get statefulsets -n nb-example
```

- 4 Restart the statefulset. Run:

```
kubectl rollout restart -n nb-example statefulset $set
```

The pod or pods associated with the statefulset are terminated and be re-created. It may take several minutes to reach the "Running" state.

- 5 Once the pods are running, re-run the kubectl command from step 2 to get the image ID of the new NetBackup container. Confirm that it's different from what it was before the update.

Update the MSDP Scaleout containers

- 1 Look at the list of pods in the application namespace and identify the pods to update. The controller pod have "uss-controller" in its name, the MDS pods have "uss-mds" in their names, and the engine pods are be named like their fully qualified domain names. Run:

```
kubectl get pods -n nb-example
```

- 2 Get the image IDs of the existing MSDP Scaleout containers and record them for later. All the MDS pods use the same image, and all the engine pods use the same image, so it's only necessary to get three image IDs, one for each type of pod.

```
kubectl get pods -n nb-example $engine $controller $mds -o  
jsonpath --template "{range  
.items[*]}{.status.containerStatuses[*].imageID}{'\n'}{end}"
```

- 3 Edit the Environment resource and change the `spec.msdpScaleouts[*].tag` values to the new tag used earlier in these instructions.

```
kubectl edit environment -n nb-example nb-env

...
spec:
  ...
  msdpScaleouts:
  - ...
    tag: "17.0-update1"
```

- 4 Save the file and close the editor. The MSDP Scaleout pods are terminated and re-created. It may take several minutes for all the pods to reach the "Running" state.
- 5 Run `kubectl get pods`, to check the list of pods and note the new name of the `uss-controller` pod. Then, once the pods are all ready, re-run the `kubectl` command above to get the image IDs of the new MSDP Scaleout containers. Confirm that they're different from what they were before the update.

Assessing cluster configuration before deployment

This chapter includes the following topics:

- [How does the webhook validation works](#)
- [Webhooks validation execution details](#)
- [How does the Config-Checker utility work](#)
- [Config-Checker execution and status details](#)

How does the webhook validation works

- Webhooks are implemented to validate the CR input provided in the `sample/environment.yaml` file which is the interface of NetBackup installation on the EKS cluster.
- For each user input in the `sample/environment.yaml` file a validation webhook is implemented.
- If any of the input value is not in the required form, then webhooks displays an error and prevents the creation of an environment.
- For primary server deployment, following webhook validations have been implemented:
 - Validate RetainReclaimPolicy: This check verifies that the storage classes used for PVC creation in the CR have reclaim policy as **Retain**. The check fails if any of the webhook do not have the **Retain** reclaim policy.

- Validate MinimumVolumeSize: This check verifies that the PVC storage capacity meets the minimum required volume size for each volume in the CR. The check fails if any of the volume capacity sizes does not meet the following requirements for Primary server.
 - Catalog volume size: 100Gi
 - Log volume size: 30Gi
 - Data volume size: 30Gi
- Validate CSI driver: This will verify that the PV created is provisioned using the `efs.csi.aws.com` driver, that is, AWS Elastic file system (EFS) for volumes catalog. If any other driver type is used, the webhook fails.
- Validate AWS Elastic file system (EFS) controller add-on: Verifies if the AWS Elastic file system (EFS) controller add-on is installed on the cluster. This AWS Elastic file system (EFS) controller is required to use EFS as persistence storage for pods which will be running on cluster. Webhooks will check the EFS controller add-on is installed and it is running properly. If no, then validation error is displayed.
- AWS Load Balancer Controller add-on check: Verifies if the AWS load balancer controller add-on is installed on the cluster. This load balancer controller is required to use load balancer in the cluster. Webhooks will check the load balancer controller add-on is installed and it is running properly. If no, then a validation error is displayed.

Webhooks validation execution details

Note the following points.

- A Webhook is an HTTP call back: An HTTP POST that occurs when an event-notification is sent through HTTP POST. A web application implementing Webhooks will POST a message to a URL when certain tasks happen.
- Webhooks are called when the following command is applied to create/update the environment to validate the CR input provided into the yaml file:

```
kubectl apply -f sample/environment.yaml
```
- Webhook validates each check in sequence. Even if one of the validation fails then a validation error is displayed and the execution is stopped.
- The error must be fixed and the `environment.yaml` file must be applied so that the next validation check is performed.
- The environment is created only after webhook validations are passed.

How does the Config-Checker utility work

The Config-Checker utility performs checks on the deployment environment to verify that the environment meets the requirements, before starting the primary server and media server deployments.

How does the Config-Checker works:

- **RetainReclaimPolicy check:**

This check verifies that the storage classes used for PVC creation in the CR have reclaim policy as **Retain**. The check fails if any of the storage classes do not have the **Retain** reclaim policy.

For more information, see [Persistent Volumes Reclaiming](#)

- **Provisioner check:**

- Primary server: This will verify that the PV created using the storage class is provisioned using the `efs.csi.aws.com` driver, that is, AWS Elastic file system (EFS) for catalog volume. This will verify that the storage type provided is Amazon Elastic Block Store(Amazon EBS) for data and log volume. If any other driver type is used, the Config-Checker fails.

- Media server: This will verify that the storage type provided is Amazon Elastic Block Store (Amazon EBS) for data and log volume. Config-Checker fails if this requirement is not met for media server.

- **MinimumVolumeSize check:**

This check verifies that the PVC storage capacity meets the minimum required volume size for each volume in the CR. The check fails if any of the volume capacity sizes does not meet the requirements.

Following are the minimum volume size requirements:

- Primary server:

- Data volume size: 30Gi
- Catalog volume size: 100Gi
- Log volume size: 30Gi

- Media server:

- Data volume size: 50Gi
- Log volume size: 30Gi

- **AWS Load Balancer Controller add-on check:**

This check verifies if the AWS Load Balancer Controller add-on is installed in the cluster. This load balancer controller is required for load balancer in the

cluster. If this check fails, user must deploy the AWS Load Balancer Controller add-on

- **Volume expansion check:**

This check verifies the storage class name given for Primary server data and log volume and for Media server data and log volumes has `AllowVolumeExpansion = true`. If Config-Checker fails with this check then it gives a warning message and continues with deployment of NetBackup media servers.

Config-Checker execution and status details

Note the following points.

- Config-Checker is executed as a separate job in Kubernetes cluster for both the primary server and media server CRs respectively. Each job creates a pod in the cluster. Config-checker creates the pod in the operator namespace.

Note: Config-checker pod gets deleted after 4 hours.

- Execution summary of the Config-Checker can be retrieved from the Config-Checker pod logs using the `kubectl logs <configchecker-pod-name> -n <operator-namespace>` command.
This summary can also be retrieved from the operator pod logs using the `kubectl logs <operator-pod-name> -n <operator-namespace>` command.
- Following are the Config-Checker modes that can be specified in the Primary and Media CR:
 - **Default:** This mode executes the Config-Checker. If the execution is successful, the Primary and Media CRs deployment is started.
 - **Dryrun:** This mode only executes the Config-Checker to verify the configuration requirements but does not start the CR deployment.
 - **Skip:** This mode skips the Config-Checker execution of Config-Checker and directly start the deployment of the respective CR.
- Status of the Config-Checker can be retrieved from the primary server and media server CRs by using the `kubectl describe <PrimaryServer/MediaServer> <CR name> -n <namespace>` command.
For example, `kubectl describe primaryservers environment-sample -n test`
- Following are the Config-Checker statuses:

- **Success:** Indicates that all the mandatory config checks have successfully passed.
- **Failed:** Indicates that some of the config checks have failed.
- **Running:** Indicates that the Config-Checker execution is in progress.
- **Skip:** Indicates that the Config-Checker is not executed because the `configcheckmode` specified in the CR is skipped.
- If the Config-Checker execution status is Failed, you can check the Config-Checker job logs using `kubectl logs <configchecker-pod-name> -n <operator-namespace>`. Review the error codes and error messages pertaining to the failure and update the respective CR with the correct configuration details to resolve the errors.
For more information about the error codes, refer to [NetBackup™ Status Codes Reference Guide](#).
- If Config-Checker ran in **dryrun** mode and if user wants to run Config-Checker again with same values in Primary or Media server YAML as provided earlier, then user needs to delete respective CR of Primary or Media server. And then apply it again.
 - If it is primary server CR, delete primary server CR using the `kubectl delete -f <environment.yaml>` command.
Or
If it is media server CR, edit the Environment CR by removing the media server section in the `environment.yaml` file. Before removing the **mediaServer** section, you must save the content and note the location of the content. After removing section apply environment CR using `kubectl apply -f <environment.yaml>` command.
 - Apply the CR again. Add the required data which was deleted earlier at correct location, save it and apply the yaml using `kubectl apply -f <environment.yaml>` command.

Deploying NetBackup

This chapter includes the following topics:

- [Preparing the environment for NetBackup installation on EKS](#)
- [Recommendations of NetBackup deployment on EKS](#)
- [Limitations of NetBackup deployment on EKS](#)
- [About primary server CR and media server CR](#)
- [Monitoring the status of the CRs](#)
- [Updating the CRs](#)
- [Deleting the CRs](#)
- [Configuring NetBackup IT Analytics for NetBackup deployment](#)
- [Managing NetBackup deployment using VxUpdate](#)
- [Migrating the node group for primary or media servers](#)

Preparing the environment for NetBackup installation on EKS

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

EKS-specific requirements

- 1 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 and instance type should be of at least **m5.4xlarge** configuration and select the size of attached EBS volume for each node more than 100 GB.

Note: Using separate nodegroups is required for NetBackup servers and MSDP deployments. If more than one mediaServer objects are created then they should use separate nodegroups.

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

- Minimum required policies in IAM role:
 - AmazonEKSClusterPolicy
 - AmazonEKSServicePolicy
 - AmazonEKSWorkerNodePolicy
 - AmazonEC2ContainerRegistryReadOnly
 - AmazonEKS_CNI_Policy
 - AmazonEKSServicePolicy
- 2 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.

- 3** A dedicated node pool for NetBackup must be created with manual scaling or Autoscaling enabled in Amazon Elastic Kubernetes Services cluster. The autoscaling feature allows your node pool to scale dynamically by provisioning and de-provisioning the nodes as required automatically.

The following table lists the node configuration for the primary and media servers.

Node type		D16ds v4
Disk type		P30
vCPU		16
RAM		64 GiB
Total disk size per node (TiB)		1 TB
Number of disks/node		1
Cluster storage size	Small (4 nodes)	4 TB
	Medium (8 nodes)	8 TB
	Large (16 nodes)	16 TB

- 4** Another dedicated node pool must be created for Snapshot Manager (if it has to be deployed) with auto scaling enabled.

Following is the minimum configuration required for Snapshot Manager data plane node pool:

Node type		t3.large
RAM		8 GB
Number of nodes		Minimum 1 with auto scaling enabled.

Maximum pods per node

Number of IPs required for Snapshot Manager data pool, must be greater than:

the number of nodes (for node's own IP) + (RAM size per node * 2 * number of nodes) + (number of all kube-system pods running on all nodes) + static listener pod + number of nodes(for fluent daemonset)

Number of IPs required for Snapshot Manager control pool, must be greater than:

number of nodes (for node's own IP) + number of flexsnap pods(15) + number of flexsnap services (6) + nginx load balancer IP + no. of additional off host agents + operator + (number of all kube-system pods running on all nodes)

Following are the different scenario's on how the NetBackup Snapshot Manager calculates the number of job which can run at a given point in time, based on the above mentioned formula:

- For 2 CPU's and 8 GB RAM node configuration:

CPU	More than 2 CPU's
RAM	8 GB

Maximum pods per node	<p>Number of IPs required for Snapshot Manager data pool, must be greater than:</p> <p>number of nodes (for node's own IP) + (RAM size per node * 2 * number of nodes) + (number of all kube-system pods running on all nodes) + static listener pod + number of nodes(for fluent daemonset)</p> <p>Number of IPs required for Snapshot Manager control pool, must be greater than:</p> <p>number of nodes (for node's own IP) + number of flexsnap pods(15) + number of flexsnap services (6) + nginx load balancer IP + no. of additional off host agents + operator + (number of all kube-system pods running on all nodes)</p>
Autoscaling enabled	<p>Minimum number =1 and Maximum = 3</p>

Note: Above configuration will run 8 jobs per node at once.

- For 2/4/6 CPU's and 16 GB RAM node configuration:

CPU	More than 2/4/6 CPU's
RAM	16 GB
Maximum pods per node	6 (system) + 4 (Static pods) + 16*2=32 (Dynamic pods) = 42 or more
Autoscaling enabled	Minimum number =1 and Maximum = 3

Note: Above configuration will run 16 jobs per node at once.

- 5** Taints and tolerations allows you to mark (taint) a node so that no pods can schedule onto it unless a pod explicitly *tolerates* the taint. Marking nodes instead of pods (as in node affinity/anti-affinity) is particularly useful for situations where most pods in the cluster must avoid scheduling onto the node.

Taints are set on the node group while creating the node group in the cluster. Tolerations are set on the pods.

To use this functionality, user must create the node group with the following detail:

- Add a label with certain key value. For example `key = nbpool, value = nbnodes`
- Add a taint with the same key and value which is used for label in above step with effect as *NoSchedule*.

For example, `key = nbpool, value = nbnodes, effect = NoSchedule`

Provide these details in the operator yaml as follows. To update the toleration and node selector for operator pod,

Edit the `operator/patch/operator_patch.yaml` file. Provide the same label `key:value` in node selector section and in toleration sections. For example,

```
nodeSelector:
  nbpool: nbnodes
  # Support node taints by adding pod tolerations equal to the specif
  # For Toleartion NODE_SELECTOR_KEY used as a key and NODE_SELECTOR_V
tolerations:
  - key: nbpool
    operator: "Equal"
    value: nbnodes
```

Update the same label `key:value` as `labelKey` and `labelValue` in `nodeselector` section in `environment.yaml` file.

- 6** Deploy aws load balancer controller add-on in the cluster.

For more information on installing the add-on, see [Installing the AWS Load Balancer Controller add-on](#).

- 7** Install cert-manager by using the following command:

```
$ kubectl apply -f
https://github.com/cert-manager/cert-manager/releases/download/v1.8.0/cert-manager.yaml
```

For more information, see [Documentation for cert-manager installation](#).

8 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.

9 Amazon Elastic File System (Amazon EFS) for shared persistence storage. To create EFS for primary server, see [Create your Amazon EFS file system](#).

EFS configuration can be as follow and user can update Throughput mode as required:

Performance mode: General Purpose

Throughput mode: Provisioned (256 MiB/s)

Availability zone: Regional

Note: Throughput mode can be increased at runtime depending on the size of workloads and also if you are seeing performance issue you can increase the Throughput mode till 1024 MiB/s.

Note: To install the add-on in the cluster, ensure that you install the Amazon EFS CSI driver. For more information on installing the Amazon EFS CSI driver, see [Amazon EFS CSI driver](#).

- 10 Create a storage class with Managed disc storage type with `efs.csi.aws.com` and allows volume expansion. It must be in LRS category with Premium SSD. It is recommended that the storage class has, `Retain` reclaim. Such storage class can be used for primary server as it supports Amazon Elastic File System storage only for catalog volume.

For more information on Amazon Elastic File System, see [Create your Amazon EFS file system](#).

For example,

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: efs-sc
provisioner: efs.csi.aws.com
parameters:
  provisioningMode: efs-ap
  fileSystemId: fs-92107410
  directoryPerms: "700"
  gidRangeStart: "1000" # optional
  gidRangeEnd: "2000" # optional
  basePath: "/dynamic_provisioning" # optional
```

- 11 If 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. See [“About the Load Balancer service”](#) on page 146. for more information on NetBackup ports.
 - To obtain the cluster security group, run the following command:

```
aws eks describe-cluster --name <my-cluster> --query cluster.resourcesVpcConfig.clusterSecurityGroupId
```
 - The following link helps to add inbound rule to the security group:
[Add rules to a security group](#)

- 12** Create a storage class with EBS storage type with `allowVolumeExpansion = true` and `ReclaimPolicy=Retain`. This storage class is to be used for data and log for both primary and media servers.

For example,

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: efs-sc
provisioner: efs.csi.aws.com
parameters:
  provisioningMode: efs-ap
  fileSystemId: <EFS ID>
  directoryPerms: "700"
reclaimPolicy: Retain
volumeBindingMode: Immediate
allowVolumeExpansion: true
```

Note: Ensure that you install the Amazon EBS CSI driver to install the add-on in the cluster. For more information on installing the Amazon EBS CSI driver, see [Managing the Amazon EBS CSI driver as an Amazon EKS add-on](#) and [Amazon EBS CSI driver](#).

- 13** The EFS based PV must be specified for Primary server catalog volume with `ReclaimPolicy=Retain`.

Host-specific requirements

- 1** Install AWS CLI.

For more information on installing the AWS CLI, see [Installing or updating the latest version of the AWS CLI](#).

- 2** Install Kubectl CLI.

For more information on installing the Kubectl CLI, see [Installing kubectl](#).

- 3** Configure docker to enable the push of the container images to the container registry.

- 4** Create the OIDC provider for the AWS EKS cluster.

For more information on creating the OIDC provider, see [Create an IAM OIDC provider for your cluster](#).

- 5 Create an IAM service account for the AWS EKS cluster.
For more information on creating an IAM service account, see [Amazon EFS CSI driver](#).
- 6 If an IAM role needs an 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
```


For more information on creating an IAM role, see [Enabling IAM user and role access to your cluster](#).
- 7 Login 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>
```
- 8 Free space of approximately 8.5GB on the location where you copy and extract the product installation TAR package file. If using docker locally, there should be approximately 8GB 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.
- 9 AWS EFS-CSI driver should be installed for static PV/PVC creation of primary catalog volume.

Recommendations of NetBackup deployment on EKS

Note the following recommendations:

- Use AWS Premium storage for data volume in media server CR.
- Use AWS Standard storage for log volume in media server CR.
- For primary server volume (catalog), use `Amazon EFS` as storage type. For media server, primary server volumes, log and data volumes use `Amazon EBS` as storage type.
- Do not delete the disk linked to PV used in primary server and media server CR deployment. This may lead to data loss.
- Ensure that in one cluster, only one NetBackup operator instance is running.
- Do not edit any Kubernetes resource created as part of primary server and media server custom resource. Update is supported through custom resource update only.

- Detailed primary server custom resource deployment and media server custom resource deployment logs are retrieved from NetBackup operator pod logs using the `kubectl logs <netbackup-operator-pod-name> -c netbackup-operator -n <netbackup operator-namespace>` command .
- Deploy primary server custom resource and media server custom resource in same namespace.
- Ensure that you follow the symbolic link and edit the actual persisted version of the file, if you want to edit a file having a symbolic link in the primary server or media server.
- In case of upgrade and during migration, do not delete the `Amazon elastic files` linked to the old PV which is used in primary server CR deployment until the migration is completed successfully. Else this leads to data loss.
- Specify different block storage based volume to obtain good performance when the `nbdeployutil` utility does not perform well on Amazon elastic files based volumes.

Limitations of NetBackup deployment on EKS

Note the following limitations:

- *(Applicable only for media servers)* A storage class that has the storage type as `EFS` is not supported. When the Config-Checker runs the validation for checking the storage type, the Config-Checker job fails if it detects the storage type as `EFS`. But if the Config-Checker is skipped then this validation is not run, and there can be issues in the deployment. There is no workaround available for this limitation. You must clean up the PVCs and CRs and reapply the CRs.
- Media server scale down is not supported. Certain workloads that require media server affinity for the clients would not work.
- External Certificate Authority (ECA) is not supported.
- In case of load balancer service updating the CR with dynamic IP address to static IP address and vice versa is not allowed.
- Media server pods as NetBackup storage targets are not supported. For example, NetBackup storage targets like `AdvancedDisk` and so on are not supported on the media server pods.

About primary server CR and media server CR

Primary server custom resource is used to deploy the NetBackup primary server and media server custom resource is used to deploy the NetBackup media server.

- After the operator is installed, update the custom resource YAMLs to deploy the primary server and media server CRs located in the `samples` folder.
- The primary server CRD and media server CRD are located in `operator_deployment.yaml` in the operator folder where the package is extracted.
- **Name** used in the primary server and media server CRs must not be same. In the primary server CR the **Name** should not contain the word **media** and in the media server CR the **Name** should not contain the word **primary**.

Note: After deployment, you cannot change the **Name** in primary server and media server CR.

- Before the CRs can be deployed, the utility called `Config-Checker` is executed that performs checks on the environment to ensure that it meets the basic deployment requirements. The config-check is done according to the **configCheckMode** and **paused** values provided in the custom resource YAML. See [“How does the Config-Checker utility work”](#) on page 61.
- You can deploy the primary server and media server CRs in same namespace.
- Use the storage class that has the storage type as `Amazon elastic files` for the catalog volume in the primary server CR. For data and log volumes in the media server use the storage type as `EBS`.
- During fresh installation of the NetBackup servers, the value for **keep logs up to** under **log retention configuration** is set based on the log storage capacity provided in the primary server CR inputs. You may change this value if required. To update logs retention configuration, refer the steps mentioned in [NetBackup™ Logging Reference Guide](#).
 - The NetBackup deployment sets the value as per the formula.
Size of logs PVC/PV * 0.8 = Keep logs up value By default, the default value is set to 24GB.
For example: If the user configures the storage size in the CR as 40GB (instead of the default 30GB) then the default value for that option become 32GB automatically based on the formula.

Note: This value will get automatically updated to the value of `bp.conf` file on volume expansion.

- Deployment details of primary server and media server can be observed from the operator pod logs using the following command:

```
kubectl logs <operator-pod-name> -c netbackup-operator -n  
<operator-namespace>
```

After installing primary server CR

Note the following points:

- The primary server CR will create a pod, a statefulset, a load balancer service and a configmap.
- Initially pod will be in not ready state (0/1) when installation is going on in the background. Check the pod logs for installation progress using the following command:

```
kubectl logs <primary-pod-name> -n <namespace>
```

Primary server can be considered as successfully installed and running when the primary server pod's state is **ready (1/1)** and the Statefulset is **ready (1/1)**.

- You can access the NetBackup webUI using the **primary server hostname** that was specified in the primary server CR status in **Primary server details** section.

For example, if the primary server hostname is **nbu-primary**, then you can access the webUI at <https://nbu-primary/webui/login>.

After Installing the media server CR

Note the following points:

- The media server CR will create a pod, a statefulset, a configmap, a loadbalancer service, a persistent volume claim for data volume, and a persistent volume claim for log volume.
- Initially pod will be in not ready state (0/1) when installation is going in the background. Check the pod logs for installation progress using the following command:

```
kubectl logs <media-pod-name> -n <namespace>
```

Media server can be considered as successfully installed and running when the media server pod's state is **ready (1/1)**, and the Statefulset is **ready (1/1)**, for each replica count.

- Details of media server name for each replica can be obtained from media server CR status by running the following command:

```
kubectl describe <MediaServer_cr_name> -n <namespace>
```

Monitoring the status of the CRs

You can view the status and other details of the primary server and media server CRs using the following commands:

- `kubectl get <PrimaryServer/MediaServer> -n <namespace> OR`
- `kubectl describe <PrimaryServer/MediaServer> <CR name> -n <namespace>`

Following table describes the primary server CR and media server CR status fields:

Table 4-1

Section	Field / Value	Description
Primary Server Details Only one hostname and IP address for the respective primary server.	Host Name	Name of the primary server that should be used to access the web UI.
	IP	IP address to access the primary server.
	Version	This indicates that the NetBackup primary server version is installed.
Media Server Details Number of hostname and IP address is equal to the replica count mentioned CR spec.	Host Name	Name of the media server.
	IP	IP address to access the media server.
	Version	This indicates that the NetBackup media server version is installed.

Table 4-1 (continued)

Section	Field / Value	Description
Attributes	Resource Name	Statefulset name of the respective server.
	Primary/Media server name	Name of the primary server or media server deployed.
	Config checker status	Indicates the status of the config checker as passed, failed, or skipped. For more information on the Config-Checker, See “How does the Config-Checker utility work” on page 61.
	SHA Fingerprint	Represents the SHA key fingerprint of the NetBackup primary server. Note: SHAFingerprint represents the SHA256 CA certificate fingerprint of the primary server.
Error Details	Code	A code assigned to an error encountered during the CR deployment or during the config-check operation. For more information on the error, refer to the <i>NetBackup Status Code Reference Guide</i> .
	Message	Message that describes the respective error code.

Table 4-1 (continued)

Section	Field / Value	Description
State	Success /Paused /Failed /Running	<p>Current state of the custom resource, from one of the following:</p> <ul style="list-style-type: none"> ■ Success: Indicates that the deployment of Primary/Media Custom Resource (CR) is successful. However, this does not mean that the installation of the NetBackup primary/media servers is successful. The primary or media server StatefulSets and/or the pods might or might not be in a ready state, irrespective of that the Primary/Media CR state will show as Success. ■ Paused: Indicates that the reconciler is in paused state and deployment of a CR is paused. ■ Failed: Indicates that the deployment of a CR failed with errors. However this does not mean failed installation of the NetBackup Server. Errors can be analyzed from the Operator logs or CR describe. ■ Running: Indicates that the CR deployment is in progress and the resources are getting created.
Events	INIT/FAILOVER/UPGRADE	<p>The events like INIT, FAILOVER and UPGRADE are logged in here.</p> <p>The details of these events can also be added.</p>

Updating the CRs

After the successful deployment of the primary server and media server CRs, you can update the values of only selected specs by editing the respective environment custom resource.

Note: Updating the Kubernetes resources (pod, configmap, services, statefulset etc) created for the CRs is not recommended.

Following tables describe the specs that can be edited for each CR.

Table 4-2 Primary server CR

Spec	Description
paused	Specify <i>True</i> or <i>False</i> as a value, to temporarily stop the respective CR controller. <i>True</i> : Stop the controller. <i>False</i> : Resume the controller.
configCheckMode	Specify <i>default</i> , <i>dryrun</i> or <i>skip</i> as a value. See “Config-Checker execution and status details” on page 62.

Table 4-3 Media server CR

Spec	Description
paused	Specify <i>True</i> or <i>False</i> as a value, to temporarily stop the respective CR controller. <i>True</i> : Stop the controller. <i>False</i> : Resume the controller.
replicas	Represents the replica count of the media server. Media server count can be scaled up by incrementing the replica count. Reducing the replica count is not supported.
configCheckMode	Specify <i>default</i> , <i>dryrun</i> or <i>skip</i> as a value. See “Config-Checker execution and status details” on page 62.

If you edit any other fields, the deployment can go into an inconsistent state.

Deleting the CRs

If you must delete any of the CRs for a valid reason such as for the troubleshooting purpose, or because any of the specs provided were incorrect; you can reinstall the deleted CR after resolving the issues.

Note: Once installed, deleting a CR is not recommended as it will stop the deployment and NetBackup will not work.

Notes:

- Deleting a CR will delete all its child resources like pod, statefulset, services, configmaps, config checker job, config checker pod.
- Deleting operator with `kubectl delete -k <operator_folder_path>` will delete the CRs and its resources except the PVC.
- Persistent volume claim (PVC) will not be deleted upon deleting a CR so that the data is retained in the volumes. Then if you create a new CR with the same name as the deleted one, the existing PVC with that same name will be automatically linked to the newly created pods.
- Do not delete `/mnt/nbdata`, `/mnt/nblogs` and `/mnt/nbdb` folders manually from primary server and media pods. The NetBackup deployment will go into an inconsistent state and will also result in data loss.

Configuring NetBackup IT Analytics for NetBackup deployment

NetBackup IT Analytics can be configured to use with NetBackup primary server in this Kubernetes environment. NetBackup IT Analytics can be configured at the time of primary server deployment or user can update the primary server CR to configure NetBackup IT Analytics.

To configure NetBackup IT Analytics for NetBackup deployment

- 1 Using the `ssh-keygen` command, generates public key and private key on NetBackup IT Analytics data collector.

NetBackup IT Analytics data collector uses passwordless ssh login.
- 2 Update the primary server CR, copy public key generated in previous steps to “itAnalyticsPublicKey” section in spec.
 - Apply the primary server CR changes using `kubectl apply -f environment.yaml -n <namespace>`.
On successful deployment of primary server CR, describe the primary server CR using `kubectl describe PrimaryServer <primary-server-name> -n <namespace>`
 - In status section, verify **It Analytics Configured** is set to **true**.
For more information, refer to the [NetBackup™ Web UI Administrator's Guide](#).

- 3 Create and copy NetBackup API key from NetBackup web UI.
- 4 On NetBackup IT Analytics portal:
 - Navigate to **Admin > Collector Administration > Select respective data collector > Add policy > Veritas NetBackup > Add**.
 - Add required options, specify the NetBackup API in the **API Key** field, and then click **OK**.
 - Select newly added primary server from **NetBackup Master Servers** and provide **nbitanalyticsadmin** as **Master Server User ID**.
 - Provide **privateKey=<path-of-private-key>|password=<passphrase>** as **Master Server Password** and **Repeat Password** whereas **<path-of-private-key>** is the private key created using ssh-keygen in earlier steps and **<passphrase>** is the passphrase used while creating private key via ssh-keygen.
 - Provide appropriate data to data collector policy fields and select collection method as **SSH or WMI protocol to NetBackup Master Server**.

Configuring the primary server with NetBackup IT Analytics tools is supported only once from primary server CR.

For more information about IT Analytics data collector policy, see [Add a Veritas NetBackup Data Collector policy](#) and for more information about adding NetBackup Primary Servers within the Data Collector policy, see [Add/Edit NetBackup Master Servers within the Data Collector policy](#).

To change the already configured public key

- 1 Execute the following command in the primary server pod:

```
kubectl exec -it -n <namespace> <primaryServer-pod-name> --  
/bin/bash
```
- 2 Copy the new public keys in the `/home/nbitanalyticsadmin/.ssh/authorized_keys` and `/mnt/nbdata/.ssh/nbitanalyticsadmin_keys` files.
- 3 Restart the **sshd** service using the `systemctl restart sshd` command.

Managing NetBackup deployment using VxUpdate

VxUpdate package is not shipped with the NetBackup deployment package. You must download and add it in NetBackup primary server.

To manage NetBackup deployment using VxUpdate

- 1 Download the required VxUpdate package on the docker-host used to interact with EKS cluster.
- 2 Copy the VxUpdate package to primary server pod using the `kubectl cp` `<path-vxupdate.sja>` `<primaryServerNamespace>/<primaryServer-pod-name>:<path-on-primary-pod>` command.
- 3 After VxUpdate package is available on primary server Pod, add it to NetBackup repository using any one of the following:
 - Select the VxUpdate package from the NetBackup web UI.
 - Execute the following command in the primary server pod:

```
kubectl exec -it -n <primaryserver-namespace>  
<primaryServer-pod-name> -- bash
```

And run the following command:

```
nbrepo -a <vxupdate-package-path-on-PrimaryPod>
```

After adding the VxUpdate package to `nbrepo`, this package is persisted even after pod restarts.

Migrating the node group for primary or media servers

You can migrate the node group for primary or media servers.

To migrate the node group for primary or media servers

- 1 Edit environment CR object using the the following command:

```
kubectl edit environment <environmentCR_name> -n <namespace>
```
- 2 Change the node selector **labelKey** and **lableValue** to new values for primary/media server.
- 3 Save the environment CR.

This will change the statefulset for respective NetBackup server replica to 0 for respective server. This will terminate the pods. After successful migration, statefulset replicas will be set to original value.

Upgrading NetBackup

This chapter includes the following topics:

- [Preparing for NetBackup upgrade](#)
- [Upgrading NetBackup operator](#)
- [Upgrading NetBackup application](#)
- [Upgrade NetBackup from previous versions](#)
- [Procedure to rollback when upgrade fails](#)

Preparing for NetBackup upgrade

Note: Ensure that you go through this section carefully before starting with the upgrade procedure.

During upgrade, ensure that the following sequence of upgrade is followed:

- Upgrade MSDP operator
- Upgrade NetBackup operator
- Upgrade NetBackup application

In case the above sequence is not followed, user may face data loss.

Preparing for NetBackup upgrade

- 1 Take a backup of all the NetBackup jobs and ensure that all the jobs are suspended.
- 2 Take a catalog backup.
See [“Backing up a catalog”](#) on page 151.

- 3 Copy **DRPackages** files (packages) located at `/mnt/nblogs/DRPackages/` from the pod to the host machine from where EKS cluster is accessed.

Run the following command:

```
kubectl cp  
<primary-pod-namespace>/<primary-pod-name>:/mnt/nblogs/DRPackages  
<Path_where_to_copy_on_host_machine>
```

- 4 Preserve the data of `/mnt/nbdata` and `/mnt/nblogs` on host machine by creating tar and copying it using the following command:

```
kubectl cp  
<primary-pod-namespace>/<primary-pod-name>:<tar_file_name>  
<path_on_host_machine_where_to_preserve_the_data>
```

- 5 Preserve the environment CR object using the following command and operator directory that is used to deploy the NetBackup operator:

```
kubectl -n <namespace> get environment.netbackup.veritas.com  
<environment name> -o yaml > environment.yaml
```

Note: Ensure that you upgrade MSDP operator first.

- 6 In case of existing NetBackup deployments with EBS, change the storage class for catalog data of primary server to trigger data migration. For more information on changing the storage class, refer to the following section:

See [“Upgrade NetBackup from previous versions”](#) on page 88.

Upgrading NetBackup operator

Ensure that all the steps mentioned in the following section are performed before performing the upgrade of NetBackup operator:

See [“Preparing for NetBackup upgrade”](#) on page 84.

Upgrading the NetBackup operator

- 1 Push the new operator images, NetBackup main image to container registry with different tags.
- 2 Update the new image name and tag in images section in `kustomization.yaml` file in operator folder available in the new package folder.

- 3 Update the node selector and tolerations in `operator_patch.yaml` file in `operator/patches` folder in the new package folder.
- 4 To upgrade the operator, apply the new image changes using the following command:

```
kubectl apply -k <operator folder name>
```

After applying the changes, new NetBackup operator pod will start in operator namespace and run successfully.

Upgrading NetBackup application

Ensure that all the steps mentioned in the following section are performed before performing the upgrade of NetBackup application:

See [“Preparing for NetBackup upgrade”](#) on page 84.

Ensure that the following server upgrade sequence is followed:

- Primary server: Upgrade and verify it is successfully upgraded
- MSDP server: Upgrade and verify it is successfully upgraded
- Media server: Upgrade and verify it is successfully upgraded

Upgrading the NetBackup application

- 1 To upgrade the primary server and media server, edit the `environment.yaml` from the new package. Copy all the fields from the preserved environment CR `environment.yaml` which can be obtained from the the following section:

See “[Preparing for NetBackup upgrade](#)” on page 84.

To update the primary server, update **tag** with new image tag in **primary** section in new `environment.yaml` file.

To update the media server, update **tag** with new image tag in **mediaServers** section in new `environment.yaml` file.

Update the **storageClassName** for catalog volume for primary server in Storage subsection of primary section in `environment.yaml` file.

For example,

```
Kind: Environment
...
Spec:
  primary:
    tag: "newtag"
  mediaServers:
    tag: "newtag"
```

Apply the changes using the following command:

```
kubectl apply -f <environment.yaml>
```

Primary server and media server pods would start with new container images respectively.

Note: Upgrade the PrimaryServer first and then change the tag for MediaServer to upgrade. If this sequence is not followed then deployment may go into inconsistent state

Note: MediaServer version should be same or lower than the PrimaryServer version after upgrade. Otherwise the deployment may go into inconsistent state.

- 2 At the time of upgrade, primary server and media server status would be changed to *Running*. Once upgrade is completed, the status would be changed to *Success* again.

Perform the following if upgrade fails in between for primary server or media server

- 1 Check the installation logs using the following command:

```
kubectl logs <PrimaryServer-pod-name/MediaServer-pod-name> -n  
<PrimaryServer/MediaServer-CR-namespace>
```

- 2 If required, check the NetBackup logs by performing exec into the pod using the following command:

```
kubectl exec -it -n <PrimaryServer/MediaServer-CR-namespace>  
<PrimaryServer/MediaServer-pod-name> -- bash
```

- 3 Fix the issue and restart the pod by deleting the respective pod with the following command:

```
kubectl delete < PrimaryServer/MediaServer-pod-name > -n  
<PrimaryServer/MediaServer-CR-namespace>
```

- 4 New pod would be created and upgrade process will be restarted for the respective NetBackup server.

- 5 Data migration jobs create the pods that run before deployment of primary server. Data migration pod exist after migration for one hour only if data migration job failed. The logs for data migration execution can be checked using the following command:

```
kubectl logs <migration-pod-name> -n  
<netbackup-environment-namespace>
```

User can copy the logs to retain them even after job pod deletion using the following command:

```
kubectl logs <migration-pod-name> -n  
<netbackup-environment-namespace> > jobpod.log
```

Note: Downgrade of NetBackup servers is not supported. If this is done, there are chances of inconsistent state of NetBackup deployment.

Upgrade NetBackup from previous versions

Ensure that all the steps mentioned for data migration in the following section are performed before upgrading to the latest NetBackup or installing the latest :

See [“Preparing the environment for NetBackup installation on EKS”](#) on page 64.

- User must have deployed NetBackup on AWS with `EBS` as its storage class.

While upgrading to latest NetBackup, the existing catalog data of primary server will be migrated (copied) from EBS to Amazon elastic files.

- Fresh NetBackup deployment: If user is deploying NetBackup for the first time, then Amazon elastic files will be used for primary server's catalog volume for any backup and restore operations.

Perform the following steps to create EFS when upgrading NetBackup from version 10.0.0.1

- 1 To create EFS for primary server, see [Create your Amazon EFS file system](#).
EFS configuration can be as follow and user can update Throughput mode as required:
Performance mode: General Purpose
Throughput mode: Provisioned (256 MiB/s)
Availability zone: Regional

Note: Throughput mode can be increased at runtime depending on the size of workloads and also if you are seeing performance issue you can increase the Throughput mode till 1024 MiB/s.

- 2 Install `efs-csi-controller` driver on EKS cluster.
For more information on installing the driver, see [Amazon EFS CSI driver](#).
- 3 Note down the EFS ID for further use.
- 4 Mount EFS on any EC-2 instance and create and create two directories on EFS to store NetBackup data.

For more information, see [Mount on EC-2 instance](#).

For example,

```
[root@sych09b03v30 ~]# mkdir /efs
[root@sych09b03v30 ~]# mount -t nfs4 -o nfsvers=4.1,rsize=1048576,
wsiz=1048576,hard,timeo=600,retrans=2,noresvport
<fs-0bde325bc5b8d6969>.efs.us-east-2.amazonaws.com:
/ /efs # change EFS ID
```

After changing the existing storage class from EBS to EFS for data migration, manually create PVC and PV with EFS volume handle and update the yaml file as described in the following procedure:

1. Create new PVC and PV with EFS volume handle.

- **PVC**

CatlogPVC.yaml

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: catalog
  namespace: ns-155
spec:
  accessModes:
    - ReadWriteMany
  storageClassName: ""
  resources:
    requests:
      storage: 100Gi
  volumeName: environment-pv-primary -catalog

```

- **PV**

catalogPV

```

apiVersion: v1
kind: PersistentVolume
metadata:
  name: environment-pv-primary-catalog
  labels:
    topology.kubernetes.io/region: us-east-2
# Give the region as your configuration in your cluster
    topology.kubernetes.io/zone: us-east-2c
# Give the zone of your node instance,

```

can also check with subnet zone in which your node instance is there.

```

spec:
  capacity:
    storage: 100Gi
  volumeMode: Filesystem
  accessModes:
    - ReadWriteMany
  storageClassName: ""
  persistentVolumeReclaimPolicy: Retain
  mountOptions:
    - iam
  csi:
    driver: efs.csi.aws.com
    volumeHandle: fs-07a82a46b4a7d87f8:/nbdata

```

```
#EFS id need to be changed as per your created EFS id
claimRef:
  apiVersion: v1
  kind: PersistentVolumeClaim
  name: catalog # catalog pvc name to which data to be copied namesp
```

PVC for data (EBS)

PVC

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: data-< Primary name >-primary-0
  namespace: ns-155
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: <Storageclass name>
  resources:
    requests:
      storage: 30Gi
```

2. Edit the `environment.yaml` file and change the value of `paused` to `true` in `primary` section and apply the `yaml`.
3. Scale down the primary server using the following commands:
 - To get statefulset name: `kubectl get sts -n < namespace in environment cr (ns-155)>`
 - To scale down the STS: `kubectl scale sts --replicas=0 < STS name > -n < Namespace >`
4. Copy the data using the migration `yaml` file as follows:

catalogMigration.yaml

```
apiVersion: batch/v1
kind: Job
metadata:
  name: rsync-data
  namespace: ns-155
spec:
  template:
    spec:
      volumes:
      - name: source-pvc
        persistentVolumeClaim:
          # SOURCE PVC
          claimName: <EBS PVC name of catalog> # catalog-environment-migr

      # old PVC (EBS) from which data to be copied
      - name: destination-pvc
        persistentVolumeClaim:
          # DESTINATION PVC
          claimName: catalog # new PVC (EFS) to which data will be copi

      securityContext:
        runAsUser: 0
        runAsGroup: 0
      containers:
      - name: netbackup-migration
        image: OPERATOR_IMAGE:TAG

      #image name with tag
        command: ["/migration", '{"VolumesList":[{"Src":"srcPvc","Dest":"
"Verify":true,"StorageType":"catalog","OnlyCatalog":true}]}']

      volumeMounts:
      - name: source-pvc
        mountPath: /srcPvc
      - name: destination-pvc
        mountPath: /destPvc
      restartPolicy: Never
```

dataMigration.yaml

```

apiVersion: batch/v1
kind: Job
metadata:
  name: rsync-data2
  namespace: ns-155
spec:
  template:
    spec:
      volumes:
      - name: source-pvc
        persistentVolumeClaim:
          # SOURCE PVC
          claimName: <EBS PVC name of catalog> # catalog-environment-migr

      # old PVC (EBS) from where data to be copied
      - name: destination-pvc
        persistentVolumeClaim:
          # DESTINATION PVC
          claimName: data (EBS) pvc name # new PVC (EFS) to where data w

      securityContext:
        runAsUser: 0
        runAsGroup: 0
      containers:
      - name: netbackup-migration
        image: OPERATOR_IMAGE:TAG # image name with tag
        command: ["/migration", '{"VolumesList":[{"Src":"srcPvc","Dest":"
"Verify":true,"StorageType":"data","OnlyCatalog":false}]}']

      volumeMounts:
      - name: source-pvc
        mountPath: /srcPvc
      - name: destination-pvc
        mountPath: /destPvc
      restartPolicy: Never

```

5. Delete the migration job once the pods are in complete state.
6. For primary server, delete old PVC (EBS) of catalog volume.

For example, catalog-`<Name_of_primary>-primary-0` and create new PVC with same name (as deleted PVC) which were attached to primary server.

- Follow the naming conventions of static PV and PVC to consume for Primary Server Deployment.

```

catalog-<Name_of_primary>-primary-0
data-<Name_of_primary>-primary-0
Example:
catalog-test-env-primary-0
data-test-env-primary-0
environment.yaml
apiVersion: netbackup.veritas.com/v2
kind: Environment
metadata:
  name: test-env
  namespace: ns-155
spec:
  ...
  primary:
    # Set name to control the name of the primary server.
    The default value is the same as the Environment's metadata.name.
    name: test-env

```

■ **Yaml to create new catalog PVC:**

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: catalog-test-env-primary-0
  namespace: ns-155
spec:
  accessModes:
    - ReadWriteMany
  storageClassName: ""
  resources:
    requests:
      storage: 100Gi
  volumeName: environment-pv-primary-catalog

```

7. Edit the PV (mounted on EFS) and replace the name, resource version, uid with new created PVC to meet the naming convention.

Get the PV's and PVC's using the following commands:

- **To get PVC details:** `kubectl get pvc -n < Namespace >`
- **Use edit command to get PVC details:** `kubectl edit pvc < New PVC(old name) name > -n < Namespace >`

- To pause the reconciler of the particular custom resource, change the *paused: false* value to *paused: true* in the `primaryServer` or `mediaServer` section and save the changes.
 - Scale down the primary server using the following commands:
 - To get statefulset name: `kubectl get sts -n <namespace>`
 - To scale down the STS: `kubectl scale sts --replicas=0 < STS name of primary server> -n <Namespace>`
- 2** Upgrade the MSDP with new build and image tag. Apply the following command to MSDP:
- ```
./kubectl-msdp init --image <Image name:Tag> --storageclass <Storage Class Name> --namespace <Namespace>
```
- 3** Edit the `sample/environment.yaml` file from new build and perform the following changes:
- Add the tag: `<new_tag_of_upgrade_image>` tag separately under primary sections.
  - Provide the EFS ID for **storageClassName** of catalog volume in primary section.

---

**Note:** The provided EFS ID for **storageClassName** of catalog volume must be same as previously used EFS ID to create PV and PVC.

---

- Use the following command to retrieve the previously used EFS ID from PV and PVC:
 

```
kubectl get pvc -n <namespace>
```

From the output, copy the name of catalog PVC which is of the following format:

```
catalog-<resource name prefix>-primary-0
```
- Describe catalog PVC using the following command:
 

```
kubectl describe pvc <pvc name> -n <namespace>
```

Note down the value of **Volume** field from the output.
- Describe PV using the following command:
 

```
kubectl describe pv <value of Volume obtained from above step>
```

Note down the value of **VolumeHandle** field from the output which is the previously used EFS ID.

- For data and logs volume, provide the **storageClassName** and then apply `environment.yaml` file using the following command and ensure that the primary server is upgraded successfully:  

```
kubectl apply -f environment.yaml
```
- Upgrade the MSDP Scaleout by updating the new image tag in `msdp-scaleout` section in `environment.yaml` file.
- Apply `environment.yaml` file using the following command and ensure that MSDP is deployed successfully:  

```
kubectl apply -f environment.yaml
```
- Edit the `environment.yaml` file and update the image tag for Media Server in `mediaServer` section.
- Apply `environment.yaml` file using the following command and ensure that the Media Server is deployed successfully:  

```
kubectl apply -f environment.yaml
```

## Procedure to rollback when upgrade fails

---

**Note:** The rollback procedure in this section can be performed only after assuming that the customer has taken catalog backup before performing the upgrade.

---

### Perform the following steps to rollback from upgrade failure and install the NetBackup version prior to upgrade

- 1 Delete the environment CR object using the following command and wait until all the underlying resources are cleaned up:

```
kubectl delete environment.netbackup.veritas.com <environment name> -n <namespace>
```

For example, primary server CR, media server CR, MSDP CR and their underlined resources.

- 2 Delete the new operator which is deployed during upgrade using the following command:

```
kubectl delete -k <new-operator-directory>
```

This will delete the new operator and new CRDs.

- 3 Apply the NetBackup operator directory which was preserved (the directory which was used to install operator before upgrade) using the following command:  

```
kubectl apply -k <operator_directory>
```
- 4 Get names of PV attached to primary server PVC (data, catalog and log) using the following command:  

```
kubectl get pvc -n <namespace> -o wide
```
- 5 Delete the primary server PVC (data, catalog and log) using the following command:  

```
kubectl delete pvc <pvc-name> -n <namespace>
```
- 6 Delete the PV linked to primary server PVC using the following command:  

```
kubectl delete pv <pv-name> command
```
- 7 Edit the preserved `environment.yaml` file (from older version of NetBackup package directory) and remove **keySecret** section from MSDP Scaleout section. Also change the CR spec *paused: false* to *paused: true* for every object in MSDP Scaleout and media servers section.
- 8 Apply the edited `environment.yaml` file using the following command:  

```
kubectl apply -f <environment.yaml>
```
- 9 After the primary server pod is in ready state (1/1), change the CR spec from *paused: false* to *paused: true* in `environment.yaml` file of the primary server section and reapply the `environment.yaml` using the following command:  

```
kubectl apply -f environment.yaml -n <namespace>
```
- 10 Exec into the primary server pod using the following command:  

```
kubectl exec -it -n <PrimaryServer/MediaServer-CR-namespace>
<primary-pod-name> -- /bin/bash
```

  - Increase the debug logs level on primary server.
  - Create a DRPackages directory at the persisted location using `mkdir /mnt/nblogs/DRPackages` folder.
  - Change ownership of the DRPackages folder to service user using the following command:  

```
chown nbsvcusr:nbsvcusr /mnt/nblogs/DRPackages
```

- 11** Copy the earlier copied DR files to primary pod at /mnt/nblogs/DRPackages using the following command:

```
kubectl cp <Path_of_DRPackages_on_host_machine>
<primary-pod-namespace>/<primary-pod-name>:/mnt/nblogs/DRPackages
```

- 12** Execute the following steps in the primary server pod:

- Change ownership of the files in /mnt/nblogs/DRPackages using the following command:

```
chown nbsvcusr:nbsvcusr <filename>
```

- Deactivate NetBackup health probes using the following command:

```
/opt/veritas/vxapp-manage/nbu-health deactivate
```

- Stop the NetBackup services using the following command:

```
/usr/openv/netbackup/bin/bp.kill_all
```

- Execute the following command:

```
nblastidentity -import -infile
/mnt/nblogs/DRPackages/<filename>.drpkg
```

- Restart all the NetBackup services using the following command:

```
/usr/openv/netbackup/bin/bp.start
```

- 13** Verify if the security settings are enabled.

- 14** Add respective media server entry in host properties using NetBackupAdministration Console as follows:

Navigate to NetBackup Management > Host properties > Master Server > Add Additional server and add media server.

- 15** Restart the NetBackup services in primary server pod and external media server as follows:

- Exec into the primary server pod using command:

```
kubectl exec -it -n <PrimaryServer/MediaServer-CR-namespace>
<primary-pod-name> -- /bin/bash
```

- Run the following command to stop all the services:

```
/usr/openv/netbackup/bin/bp.kill_all
```

After stopping all the services, restart the services using the following command:

```
/usr/openv/netbackup/bin/bp.start_all
```

- Run the following command to stop all the NetBackup services:

```
/usr/openv/netbackup/bin/bp.kill_all
```

After stopping all the services, restart the NetBackup services using the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 16 Configure a storage unit on external media server that is used during catalog backup.

- 17 Perform catalog recovery from NetBackup Administration Console.

For more information, refer to the *Veritas™ NetBackup Troubleshooting Guide*

- 18 Exec into the primary server pod using the following command:

```
kubect1 exec -it -n <PrimaryServer/MediaServer-CR-namespace>
<primary-pod-name> -- /bin/bash
```

- Stop the NetBackup services using the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- Start the NetBackup services using the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- Activate NetBackup health probes using the following command:

```
/opt/veritas/vxapp-manage/nbu-health activate
```

- 19 Restart the NetBackup operator pod, where user must delete the pod using the following command:

```
kuebctl delete <operator-pod-name> -n <namespace>
```

Kubernetes will start new pod after deletion.

- 20 Pause the reconciler for primary, media servers, and msdp scaleouts in the following sequence:

- Change CR spec *paused: true* to *paused: false* in `environment.yaml` file of the primary section and re-apply `environment.yaml` file using the following command:

```
kubect1 apply -f environment.yaml -n <namespace>
```

Wait till primary server is in ready state.

- Change CR spec *paused: true* to *paused: false* in `environment.yaml` file of the msdp scaleouts section and re-apply `environment.yaml` file using the following command:

```
kubect1 apply -f environment.yaml -n <namespace>
```

Wait till primary server is in ready state.

- Change CR spec *paused: true* to *paused: false* in `environment.yaml` file of the media servers section and re-apply `environment.yaml` file using the following command:

```
kubectl apply -f environment.yaml -n <namespace>
```

Wait till primary server is in ready state.

- 21** Verify the rollback is successful by performing backups and recovery jobs.

# Deploying Snapshot Manager

This chapter includes the following topics:

- [Overview](#)
- [Prerequisites](#)
- [Installing the docker images](#)

## Overview

You must deploy Snapshot Manager solution in your Amazon Elastic Kubernetes Services Cluster environment. EKS must be created with appropriate network and configuration settings. Before you deploy the solution, ensure that your environment meets the requirements.

See [“Prerequisites”](#) on page 102.

## Prerequisites

### **A working Amazon Elastic Kubernetes Services cluster (EKS cluster)**

- Amazon Elastic Kubernetes Services cluster
  - Your Amazon Elastic Kubernetes Services cluster must be created with appropriate network and configuration settings. Supported Amazon Elastic Kubernetes Services cluster version is 1.21.x and later.
  - Two storage classes with the following configurations is required:

| Storage class field  | Data                  | Log             |
|----------------------|-----------------------|-----------------|
| provisioner          | kubernetes.io/aws-ebs | efs.csi.aws.com |
| reclaimPolicy        | Retain                | Retain          |
| allowVolumeExpansion | True                  | True            |

---

**Note:** It is recommended to use a separate EFS for Snapshot Manager deployment and primary server catalog.

---

- A Kubernetes Secret that contains the credentials is required.
- Amazon Elastic Kubernetes Service containerSnapshot Manager registry (ECR) Use existing ECR or create a new one. Your Kubernetes cluster must be able to access this registry to pull the images from.
- Node group  
You must have a dedicated node group created for Snapshot Managers scalable workflow and data movement services. It is recommended to use the NBU primary server's node group for Snapshot Manager's control services. User must assign same IAM role with required permissions to both the node groups, that is control node group and data node group.
- Client machine to access EKS cluster
  - A separate computer that can access and manage your EKS cluster and ECR.
  - It must have Linux operating system.
  - It must have Docker daemon, the Kubernetes command-line tool (kubectl), and AWS CLI installed.  
The Docker storage size must be more than 6 GB. The version of kubectl must be v1.21.x or later. The version of AWS CLI must meet the EKS cluster requirements.
  - If EKS is a private cluster, see [Create a private Amazon EKS Service cluster](#).
- Static Internal IPs  
If the internal IPs are used, reserve the internal IPs (avoid the IPs that are reserved by other systems) for Snapshot Manager and add forward and reverse DNS records for all of them in your DNS configuration.  
The AWS static public IPs can be used but is not recommended.

# Installing the docker images

The Snapshot Manager package

`netbackup-flexsnap-$(SNAPSHOT_MANAGER_VERSION).tar.gz` for Kubernetes includes the following:

- A docker image for Snapshot Manager operator
- 8 docker images for Snapshot Manager: `flexsnap-certauth`, `flexsnap-rabbitmq`, `flexsnap-fluentd`, `flexsnap-datamover`, `flexsnap-nginx`, `flexsnap-mongodb`, `flexsnap-core`, `flexsnap-deploy`

## To install the docker images

- 1 Download `netbackup-flexsnap-$(SNAPSHOT_MANAGER_VERSION).tar.gz` from the Veritas site.
- 2 Load the docker images to your docker storage.

```
docker load -i
netbackup-flexsnap-$(SNAPSHOT_MANAGER_VERSION).tar.gz
```

### 3 Tag the images.

```
$ docker tag veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}

$ docker tag
veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}

$ docker tag veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}
${REGISTRY}/veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}
```

---

**Note:** Ensure that you use the same tag as that of Snapshot Manager image version. Custom tag cannot be used.

---

#### 4 Push the images.

```
$ docker push
${REGISTRY}/veritas/flexsnap-certauth:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-rabbitmq:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-fluentd:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-datamover:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-nginx:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-mongodb:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-core:${SNAPSHOT_MANAGER_VERSION}

$ docker push
${REGISTRY}/veritas/flexsnap-deploy:${SNAPSHOT_MANAGER_VERSION}
```

#### Configure Snapshot Manager

After you push the docker images to Amazon Elastic container registry, then initialize Snapshot Manager (flexsnap) operator and configure Snapshot Manager. The Snapshot Manager operator starts with NetBackup operator. For more information, refer to the following section:

See [“Deploying the operators manually”](#) on page 21.

#### Configure Snapshot Manager

- 1 Use existing dedicated namespace for Snapshot Manager to run:

```
kubectl create ns <sample-namespace>
```

- 2 Create a Snapshot Manager Secret. The Secret is used in CR.

```
kubectl create secret generic cp-creds
--from-literal=username='admin' --from-literal=password='Cloudpoint@123' -n
$ENVIRONMENT_NAMESPACE
```

See [“Deploying NetBackup and Snapshot Manager manually”](#) on page 32.

- 3 Edit the cpServer CR section of the `environment.yaml` file in the text editor.

See [“Configuring the `environment.yaml` file”](#) on page 39.

- 4 Apply the CR file to the EKS cluster:

```
kubectl apply -f <sample-cr-yaml>
```

- 5 Monitor the configuration process:

```
kubectl get all -n <namespace> -o wide
```

- 6 Verify the status by running the following command:

```
kubectl get cpservers -n <sample-namespace>
```

---

**Note:** If Snapshot Manager is uninstalled and installed again with same NetBackup primary server, then generate reissue token and edit the Snapshot Manager after enabling it.

---

### Reissue token

- 1 Login to NetBackup Web UI.
- 2 Navigate to **Security > Host Mappings**.
- 3 On the Host mappings page, identify the Snapshot Manager IP and select **Action**.
- 4 Click on **Generate reissue token**.
- 5 Provide a name for the reissue token, number of days for validation and click on **Generate**.
- 6 Note down the generated token number.
- 7 Navigate to **Workloads > Cloud**.
- 8 Select Snapshot Manager tab and click on **Action**.
- 9 Select **Enable** option and click on **Edit**.
- 10 Click on **Validate** and **Accept** the certificate.
- 11 Provide Username, Password, and Token (noted down in the above step) credentials and click on **Save**.

# Migration and upgrade of Snapshot Manager

This chapter includes the following topics:

- [Migration and updating of Snapshot Manager](#)

## Migration and updating of Snapshot Manager

### Migrating Snapshot Manager

Users can manually migrate Snapshot Manager registered with NetBackup to Kubernetes Service cluster environment by performing the following steps:

1. Disable Snapshot Manager from NetBackup.
2. Stop services on the Snapshot Manager VM.
3. Create and attach a volume to the Snapshot Manager VM:
  - Create a new volume from AWS cloud console in the same region where nodegroup of NetBackup is there.

---

**Note:** Copy and save the volume ID which would be used later.

---

- Attach the volume to Snapshot Manager VM. Create a file system and mount the volume:.

For example,

```
mkdir /mnt/test/
```

```
mkfs.ext4 /dev/xvdg
```

```
mount /dev/xvdg /mnt/test/
```

```
mount | grep xvdg
```

- To check the mount point, use the following command:  
df -h
- Copy the contents from /cloudpoint/mongodb to the new volume.  
For example, cp -r /cloudpoint/mongodb/\* /mnt/test/
- Copy flexsnap.conf and bp.conf configuration files to the VM from where cluster is accessible.
- Unmount and detach the volume from VM as follows:
  - To unmount: umount /mnt/test/
  - To detach navigate to the console and detach.

#### 4. From Cluster Access VM perform the following steps:

- Create configuration maps using the following command:  
kubect1 create cm --from-file=<path to flexsnap.conf > -n  
<environment namespace>  
kubect1 create cm nbuconf --from-file=<path to bp.conf> -n  
<environment namespace>
- Create Snapshot Manager credential Secret using the following command:  
kubect1 create secret generic cp-creds  
--from-literal=username='<username>'  
--from-literal=password='<password>' -n <environment-namespace>
- Create static PV and PVCs using the volume.
- Create StorageClass volume for EBS:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
 name: gp2-reclaim
parameters:
 fsType: ext4
 type: gp2
provisioner: kubernetes.io/aws-ebs
reclaimPolicy: Retain
allowVolumeExpansion: true
volumeBindingMode: WaitForFirstConsumer
```

- Create StorageClass volume for EFS:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
 name: efs-sc
provisioner: efs.csi.aws.com
parameters:
 provisioningMode: efs-ap
 fileSystemId: <EFS ID>
 directoryPerms: "700"
reclaimPolicy: Retain
volumeBindingMode: Immediate
```

---

**Note:** It is recommended to use a separate EFS for Snapshot Manager deployment and primary server catalog.

---

- Create mongodb Persistent Volume and Persistent Volume Claim as follows:

**mongodb\_pv.yaml :**

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: mongodb-pv <pv name>
spec:
 capacity:
 storage: 30Gi
 accessModes:
 - ReadWriteOnce
 awsElasticBlockStore:
 volumeID: aws://us-east-2a/vol-00de395edff9fa8fb <need to give in
 fsType: ext4
 persistentVolumeReclaimPolicy: Retain
 storageClassName: gp2-reclaim <storageclass name>
 volumeMode: Filesystem
 claimRef:
 apiVersion: v1
 kind: PersistentVolumeClaim
 name: <PVC name>
 namespace: <netbackup-environment>
```

**mongodb-pvc.yaml :**

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
 name: mongodb-pvc
 namespace: <netbackup-environment>
spec:
 accessModes:
 - ReadWriteOnce
 storageClassName: gp2-reclaim <Storage class name>
 resources:
 requests:
 storage: 30Gi
 volumeName: mongodb-pv <PV name>
```

For more information, refer to [Leveraging AWS EBS for Kubernetes Persistent Volumes](#).

- Create mongodb Persistent Volume and Persistent Volume Claim by applying the above yamls file as follows:

```
kubectl apply -f <path_to_mongodb_pv.yaml>
```

```
kubectl apply -f <path_to_mongodb_pvc.yaml> -n <environment-namespace>
```

- Ensure that the newly created PVC is bound to the PV.  
For example, `kubectl get pvc -n <netbackup-environment> | grep mongodb`

5. Edit the `environment.yaml` file to upgrade NetBackup (primary/media/MSDP) and add section for Snapshot Manager as follows in the `environment.yaml` file:

```
cpServer:
- name: cp-cluster-deployment
 tag: <version>
 containerRegistry: <container registry>
 credential:
 secretName: cp-creds
 networkLoadBalancer:
 annotations:
 ipAddr: <IP address>
 fqdn: <FQDN>
 storage:
 log:
 capacity: <log size>
```

```

 storageClassName: <EFS based storage class>
 data:
 capacity: <mongodb PV size>
 storageClassName: <mongodb PV storage class>
nodeSelector:
 #controlPlane:
 #nodepool: <Control node pool>
 #labelKey: <Control node label key>
 #labelValue: <Control node label value>
 dataPlane:
 nodepool: <Data node pool>
 labelKey: <Data node label key>
 labelValue: <Data node label value>

```

6. Apply the `environment.yaml` file using the following command:

```

kubectl apply -f <path to environment.yaml> -n
<environment-namespace>

```

7. Re-register the Snapshot Manager from WebUI using the edit option of existing Snapshot Manager entry and provide reissue token if the Snapshot Manager name (fqdn/ip) is same as VM deployment. For more information on reissuing the token, refer to the following section:

See [“Reissue token”](#) on page 107.

---

**Note:** Automatic cloud provider additions would be skipped, if there was any AWS plugin addition present pre-migration. User can manually add the cloud with specific region if needed. Ensure that cluster region is configured with the plugin, as it is required to obtain the cluster details and correctly calculate the capability of Snapshot Manager.

---

After migration, duplicate plug-in entries might be displayed in the NetBackup Web UI. Manually delete these duplicated plugin information from the `/usr/opensv/var/global/CloudPoint_plugin.conf` file.

## Updating Snapshot Manager

User can update few parameters on the existing deployed Snapshot Manager by making changes in the `cpServer` section of `environment.yaml` file and apply it.

Only **data size** field can be changed in `cpServer` section of CR. For update operation to work, set the value of `allowVolumeExpansion` parameter to **true** in the storage classes used.

The following table lists the parameters that can be modified during update of Snapshot Manager:

| <b>Parameters</b>                   | <b>Edit during update</b> | <b>Edit during upgrade</b> |
|-------------------------------------|---------------------------|----------------------------|
| name                                | No                        | No                         |
| tag                                 | No                        | Yes                        |
| containerRegistry                   | No                        | Yes                        |
| credential: secretName              | No                        | No                         |
| networkLoadBalancer:<br>annotations | No                        | Yes                        |
| networkLoadBalancer: fqdn           | No                        | No                         |
| networkLoadBalancer: ipAddr         | No                        | Yes                        |
| data.capacity                       | Yes                       | Yes                        |
| data.storageClassName               | No                        | No                         |
| cpServer:NodeSelector:ControlPlane  | No                        | Yes                        |
| cpServer:NodeSelector:DataPlane     | No                        | Yes                        |
| proxySettings: vx_http_proxy        | No                        | Yes                        |
| proxySettings: vx_https_proxy       | No                        | Yes                        |
| proxySettings: vx_no_proxy          | No                        | Yes                        |

# Deploying MSDP Scaleout

This chapter includes the following topics:

- [Deploying MSDP Scaleout](#)
- [Prerequisites](#)
- [Installing the docker images and binaries](#)
- [Initializing the MSDP operator](#)
- [Configuring MSDP Scaleout](#)
- [Using MSDP Scaleout as a single storage pool in NetBackup](#)
- [Configuring the MSDP cloud in MSDP Scaleout](#)

## Deploying MSDP Scaleout

You must deploy MSDP Scaleout solution in your Amazon Elastic Kubernetes Service Cluster environment. EKS must be created with appropriate network and configuration settings.

You can have multiple MSDP Scaleout deployments in the same EKS cluster. Ensure that each MSDP Scaleout deployment runs in a dedicated namespace on a dedicated node group.

Before you deploy the solution, ensure that your environment meets the requirements.

See [“Prerequisites”](#) on page 115.

**Table 8-1** MSDP Scaleout deployment steps

| Step   | Task                                                     | Description                                                                                  |
|--------|----------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Step 1 | Install the docker images and binaries.                  | See <a href="#">“Installing the docker images and binaries”</a> on page 117.                 |
| Step 2 | Initialize MSDP operator.                                | See <a href="#">“Initializing the MSDP operator”</a> on page 118.                            |
| Step 3 | Configuring MSDP Scaleout.                               | See <a href="#">“Configuring MSDP Scaleout”</a> on page 119.                                 |
| Step 4 | Use MSDP Scaleout as a single storage pool in NetBackup. | See <a href="#">“Using MSDP Scaleout as a single storage pool in NetBackup”</a> on page 121. |

See [“Cleaning up MSDP Scaleout”](#) on page 170.

See [“Cleaning up the MSDP Scaleout operator”](#) on page 171.

## Prerequisites

### A working Amazon Elastic Kubernetes Service (EKS cluster)

- AWS Kubernetes cluster
  - Your AWS Kubernetes cluster must be created with appropriate network and configuration settings.  
Supported AWS Kubernetes cluster version is 1.21.x and later.
  - The node group in EKS should not cross availability zone.
  - At least one storage class that is backed with Amazon EBS CSI storage driver [ebs.csi.aws.com](#) or with the default provisioner [kubernetes.io/aws-ebs](#), and allows volume expansion. The built-in storage class is **gp2**. It is recommended that the storage class has "Retain" reclaim policy.
  - AWS Load Balancer controller must be installed on EKS.
  - A Kubernetes Secret that contains the MSDP credentials is required.  
See [“Secret”](#) on page 208.
- Amazon Elastic Container Registry (ECR)  
Use existing ECR or create a new one. Your Kubernetes cluster must be able to access this registry to pull the images from.

- **Node Group**

You must have a dedicated node group for MSDP Scaleout created. The node group should not cross availability zone.

The AWS Auto Scaling allows your node group to scale dynamically as required. If AWS Auto Scaling is not enabled, ensure the node number is not less than MSDP Scaleout size.

It is recommended that you set the minimum node number to 1 or more to bypass some limitations in EKS.
- **Client machine to access EKS cluster**
  - A separate computer that can access and manage your EKS cluster and ECR.
  - It must have Linux operating system.
  - It must have Docker daemon, the Kubernetes command-line tool (kubectl), and AWS 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 AWS CLI must meet the EKS cluster requirements.
  - If EKS is a private cluster, see [Creating an private Amazon EKS cluster](#).
- If the internal IPs are used, reserve N internal IPs and make sure they are not used. N matches the MSDP-X cluster size which is to be configured.

These IPs are used for network load balancer services. For the private IPs, please do not use the same subnet with the node group to avoid IP conflict with the secondary private IPs used in the node group.

For the DNS name, you can use the Private IP DNS name amazon provided, or you can create DNS and Reverse DNS entries under Route53.

## Existing NetBackup environment

MSDP Scaleout connects to the existing NetBackup environment with the required network ports 1556 and 443. The NetBackup primary server should be 10.0 or later. The NetBackup environment can be anywhere, locally or remotely. It may or may not be in EKS cluster. It may or may not be in the same EKS cluster.

If the NetBackup servers are on AWS cloud, besides the NetBackup configuration requirements, the following settings are recommended. They are not MSDP-specific requirements, they just help your NetBackup environment run smoothly on AWS cloud.

- Add the following in `/usr/opensv/netbackup/bp.conf`

```
HOST_HAS_NAT_ENDPOINTS = YES
```

- Tune sysctl parameters as follows:

```
net.ipv4.tcp_keepalive_time=120

net.ipv4.ip_local_port_range = 14000 65535

net.core.somaxconn = 1024
```

Tune the max open files to 1048576 if you run concurrent jobs.

## Installing the docker images and binaries

The MSDP package `VRTSpddek.tar.gz` for Kubernetes includes the following:

- A docker image for MSDP operator
- 3 docker images for MSDP Scaleout: `uss-controller`, `uss-mds`, and `uss-engine`
- A kubectl plugin: `kubectl-msdp`

### To install the docker images and binaries

- 1 Download `VRTSpddek.tar.gz` from the Veritas site.
- 2 Load the docker images to your docker storage.

```
tar -zxvf VRTSpddek.tar.gz

ls VRTSpddek-*/images/*.tar.gz|xargs -i docker load -i {}
```

- 3 Copy MSDP kubectl plugin to a directory from where you access EKS host. This directory can be configured in the `PATH` environment variable so that kubectl can load `kubectl-msdp` as a plugin automatically.

For example,

```
cp ./VRTSpddek-*/bin/kubectl-msdp /usr/local/bin/
```

- 4 Push the docker images to the ECR.

- Log in.

```
aws ecr get-login-password \
--region <region> \
| docker login \
--username AWS \
--password-stdin \
<aws_account_id>.dkr.ecr.<region>.amazonaws.com
```

- Create a repository.

See AWS documentation [Creating a private repository](#)

- Push the docker images to ECR. Keep the image name and version same as original.

```
for image in msdp-operator uss-mds uss-controller
uss-engine; do \
docker image tag $image:<version> <your-ecr-url>/$image:<version>; \
docker push <your-ecr-url>/$image:<version>; \
done
```

## Initializing the MSDP operator

Run the following command to initialize MSDP operator.

```
kubectl msdp init -i <ecr-url>/msdp-operator:<version> -s
<storage-class-name> [-l agentpool=<nodegroup-name>]
```

You can use the following `init` command options.

**Table 8-2**      `init` command options

| Option | Description                                                                                                                                                                                                                                                                                       |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -i     | MSDP operator images on your ECR.                                                                                                                                                                                                                                                                 |
| -s     | The storage class name.                                                                                                                                                                                                                                                                           |
| -l     | Node selector of the MSDP operator.<br>By default, each node group has a unique label with key-value pair <code>agentpool=&lt;nodegroup-name&gt;</code> . If you have assigned a different and cluster-wise unique label for the node group, you can use that instead of <code>agentpool</code> . |
| -c     | Core pattern of the operator pod.<br>Default value: <code>"/core/core.%e.%p.%t"</code>                                                                                                                                                                                                            |
| -d     | Enable debug-level logging in MSDP operator.                                                                                                                                                                                                                                                      |
| -a     | The maximum number of days to retain the old log files.<br>Range: 1-365<br>Default value: 28                                                                                                                                                                                                      |

**Table 8-2** init command options (*continued*)

| Option | Description                                                                        |
|--------|------------------------------------------------------------------------------------|
| -u     | The maximum number of old log files to retain.<br>Range: 1-20<br>Default value: 20 |
| -n     | Namespace scope for this request.<br>Default value: msdp-operator-system           |
| -o     | Generate MSDP operator CRD YAML.                                                   |
| -h     | Help for the <code>init</code> command.                                            |

This command installs Custom Resource Definitions (CRD) **msdp-scaleouts.msdp.veritas.com** and deploys MSDP operator in the Kubernetes environment. MSDP operator runs with Deployment Kubernetes workload type with single replica size in the default namespace **msdp-operator-system**.

MSDP operator also exposes the following services:

- **Webhook service**  
The webhook service is consumed by Kubernetes api-server to mutate and validate the user inputs and changes of the MSDP CR for the MSDP Scaleout configuration.
- **Metrics service**  
The metric service is consumed by Kubernetes/EKS for Amazon CloudWatch integration.

You can deploy only one MSDP operator instance in an EKS cluster.

Run the following command to check the MSDP operator status.

```
kubectl -n msdp-operator-system get pods -o wide
```

## Configuring MSDP Scaleout

After you push the docker images to ACR and initialize MSDP operator, configure MSDP Scaleout.

**To configure MSDP Scaleout**

- 1 Create a dedicated namespace for MSDP Scaleout to run.

```
kubectl create ns <sample-namespace>
```

- 2 Create an MSDP Scaleout Secret. The Secret is used in CR.

```
kubectl apply -f <secret-yaml-file>
```

See “[Secret](#)” on page 208.

- 3 Display the custom resource (CR) template.

```
kubectl msdp show -c
```

- 4 Save the CR template.

```
kubectl msdp show -c -f <file path>
```

- 5 Edit the CR file in the text editor.

- 6 Apply the CR file to the EKS cluster.

---

**Caution:** Add `MSDP_SERVER = <first Engine FQDN>` in `/usr/openv/netbackup/bp.conf` file on the NetBackup primary server before applying the CR YAML.

---

```
kubectl apply -f <sample-cr-yaml>
```

## 7 Monitor the configuration progress.

```
kubectl get all -n <namespace> -o wide
```

In the **STATUS** column, if the readiness state for the controller, MDS and engine pods are all **Running**, it means that the configuration has completed successfully.

In the **READINESS GATES** column for engines, 1/1 indicates that the engine configuration has completed successfully.

## 8 If you specified `spec.autoRegisterOST.enabled: true` in the CR, when the MSDP engines are configured, the MSDP operator automatically registers the storage server, a default disk pool, and a default storage unit in the NetBackup primary server.

A field **ostAutoRegisterStatus** in the Status section indicates the registration status. If **ostAutoRegisterStatus.registered** is **True**, it means that the registration has completed successfully.

You can run the following command to check the status:

```
kubectl get msdp-scaleouts.msdp.veritas.com -n <sample-namespace>
```

You can find the storage server, the default disk pool, and storage unit on the Web UI of the NetBackup primary server.

# Using MSDP Scaleout as a single storage pool in NetBackup

If you did not enable automatic registration of the storage server (`autoRegisterOST`) in the CR, you can configure it manually using the NetBackup Web UI.

See “[MSDP Scaleout CR](#)” on page 209.

**To use MSDP Scaleout as a single storage pool in NetBackup**

- 1 Follow the OpenStorage wizard with storage type **PureDisk** to create the storage server using the first Engine FQDN.

MSDP storage server credentials are defined in the Secret resource.

For more information, see *Create a Cloud storage, OpenStorage, or AdvancedDisk storage server* topic of the *NetBackup Web UI Administrator's Guide*.

- 2 Follow the MSDP wizard to create the disk pool.

For more information, see *Create a disk pool* topic of the *NetBackup Web UI Administrator's Guide*.

- 3 Follow the MSDP wizard to the storage unit.

For more information, see *Create a disk pool* topic of the *NetBackup Web UI Administrator's Guide*.

You can use MSDP Scaleout like the legacy single-node MSDP.

## Configuring the MSDP cloud in MSDP Scaleout

After you configure the local LSU, you can also configure MSDP cloud in MSDP Scaleout.

For more information about MSDP cloud support, see the *NetBackup Deduplication Guide*.

# Upgrading MSDP Scaleout

This chapter includes the following topics:

- [Upgrading MSDP Scaleout](#)

## Upgrading MSDP Scaleout

You can upgrade the MSDP Scaleout solution to the latest version in your EKS environment.

### To upgrade MSDP Scaleout

- 1 Install the new **kubect1** plug-in and push the new docker images to your container registry.

See [“Installing the docker images and binaries”](#) on page 117.

- 2 Run the following command to upgrade the MSDP operator.

```
kubect1 msdp init -i <new-operator-image> -s <storage-class-name>
-l agentpool=<nodepool-name> -n <operator-namespace>
```

All the options except `-i` option must be same as earlier when the operator was deployed initially.

- 3 Run the following command to change the `spec.version` in the existing CR resources.

```
kubectl edit msdp-scaleout <cr-name>
```

Wait for a few minutes. MSDP operator upgrades all the pods and other MSDP Scaleout resources automatically.

---

**Note:** If you use the environment operator for the MSDP Scaleout deployment, change the version string for MSDP Scaleout in the environment operator CR only. Do not change the version string in the MSDP Scaleout CR.

---

- 4 Upgrade process restarts the pods. The NetBackup jobs are interrupted during the process.

#### Upgrade NetBackup 10.0, 10.0.0.1 or 10.1 to NetBackup 10.1.1

After you upgrade NetBackup 10.0 (10.0.0.1 or 10.1) to NetBackup 10.1.1 or MSDP Scaleout to 17.0, if you find the storage server not supporting Instant Access capability on Web UI, or if you fail to select MSSQL recovery point to create MSSQL Instant Access on Web UI, then perform the following steps manually to refresh the Instant Access capability in NetBackup.

1. Login to NetBackup primary server.
2. Execute the following commands to refresh the MSDP capabilities on NetBackup primary server:

```
nbdevconfig -getconfig
```

```
nbdevconfig -setconfig
```

For example,

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -getconfig -stype
PureDisk -storage_server [storage server] >
```

```
/tmp/tmp_pd_config_file
```

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server [storage server] -stype PureDisk -configlist
```

```
/tmp/tmp_pd_config_file
```

3. Restart the NetBackup Web Management Console service (nbwmc) on NetBackup primary server.

For example,

```
/usr/opensv/netbackup/bin/nbwmc terminate
```

```
/usr/opensv/netbackup/bin/nbwmc start
```

# Monitoring NetBackup

This chapter includes the following topics:

- [Monitoring the application health](#)
- [Telemetry reporting](#)
- [About NetBackup operator logs](#)
- [Expanding storage volumes](#)
- [Allocating static PV for Primary and Media pods](#)

## Monitoring the application health

Kubernetes Liveness and Readiness probes are used to monitor and control the health of the NetBackup primary server and media server pods. The probes collectively also called as health probes, keep checking the availability and readiness of the pods, and take designated actions in case of any issues. The kubelet uses liveness probes to know when to restart a container, and readiness probes to know when a container is ready. For more information, refer to the Kubernetes documentation.

[Configure Liveness, Readiness and Startup Probes | Kubernetes](#)

The health probes monitor the following for the NetBackup deployment:

- Mount directories are present for the `data/catalog` at `/mnt/nbdata` and the log volume at `/mnt/nblogs`.
- `bp.conf` is present at `/usr/opensv/netbackup`
- NetBackup services are running as expected.

Following table describes the actions and time intervals configured for the probes:

Table 10-1

| Action                        | Description                                                                                                                              | Probe name      | Primary server (seconds) | Media server (seconds) |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|-----------------|--------------------------|------------------------|
| Initial delay                 | This is the delay that tells kubelet to wait for a given number of seconds before performing the first probe.                            | Readiness Probe | 120                      | 60                     |
|                               |                                                                                                                                          | Liveness Probe  | 300                      | 90                     |
| Periodic execution time       | This action specifies that kubelet should perform a probe every given number of seconds.                                                 | Readiness Probe | 30                       | 30                     |
|                               |                                                                                                                                          | Liveness Probe  | 90                       | 90                     |
| Threshold for failure retries | This action specifies that kubelet should retry the probe for given number of times in case a probe fails, and then restart a container. | Readiness Probe | 1                        | 1                      |
|                               |                                                                                                                                          | Liveness Probe  | 5                        | 5                      |

Health probes are run using the `nbu-health` command. If you want to manually run the `nbu-health` command, the following options are available:

- **Disable**  
This option disables the health check that will mark pod as not ready (0/1).
- **Enable**  
This option enables the already disabled health check in the pod. This marks the pod in ready state(1/1) again if all the NetBackup health checks are passed.
- **Deactivate**  
This option deactivates the health probe functionality in pod. Pod remains in ready state(1/1). This will avoid pod restarts due to health probes like liveness, readiness probe failure. This is the temporary step and not recommended to use in usual case.
-

- **Activate**

This option activates the health probe functionality that has been deactivated earlier using the **deactivate** option.

You can manually disable or enable the probes if required. For example, if for any reason you need to exec into the pod and restart the NetBackup services, the health probes should be disabled before restarting the services, and then they should be enabled again after successfully restarting the NetBackup services. If you do not disable the health probes during this process, the pod may restart due to the failed health probes.

---

**Note:** It is recommended to disable the health probes only temporarily for troubleshooting purposes. When the probes are disabled, the web UI is not accessible in case of the primary server pod, and the media server pods cannot be scaled up. Then the health probes must be enabled again to successfully run NetBackup.

---

#### To disable or enable the health probes

- 1 Execute the following command in the Primary or media server pod as required:

```
kubectl exec -it -n <namespace> <primary/media-server-pod-name>
-- /bin/bash
```

- 2 To disable the probes, run the `/opt/veritas/vxapp-manage/nbu-health disable` command. Then the pod goes into the **not ready** (0/1) state.

- 3 To enable the probes, run the `"/opt/veritas/vxapp-manage/nbu-health enable"` command. Then the pod will be back into the **ready** (1/1) state.

You can check pod events in case of probe failures to get more details using the `kubectl describe <primary/media-pod-name> -n <namesapce>` command.

## Telemetry reporting

Telemetry reporting entries for the NetBackup deployment on EKS are indicated with the **EKS based deployments** text.

- By default, the telemetry data is saved at the `/var/veritas/nbtelemetry/` location. The default location will not persisted during the pod restarts.
- If you want to save telemetry data to persisted location, then execute the `kubectl exec -it -n <namespace> <primary/media_server_pod_name> - /bin/bash` command in the pod using the and execute telemetry command using

`/usr/opensv/netbackup/bin/nbtelemetry` with `--dataset-path=DESIRED_PATH` option.

- Exec into the primary server pod using the following command:  

```
kubectl exec -it -n <namespace> <primary/media_server_pod_name> -- /bin/bash
```
- Execute telemetry command using  
`/usr/opensv/netbackup/bin/nbtelemetry` with `--dataset-path=DESIRED_PATH`

---

**Note:** Here `DESIRED_PATH` must be `/mnt/nbdata` or `/mnt/nblogs`.

---

## About NetBackup operator logs

Note the following about the NetBackup operator logs.

- NetBackup operator logs can be checked using the operator pod logs using the `kubectl logs <Netbackup-operator-pod-name> -c netbackup-operator -n <netbackup-opertaor-namespace>` command.
- NetBackup operator provides different log levels that can be changed before deployment of NetBackup operator.

The following log levels are provided:

- -1 - Debug
- 0 - Info
- 1 - Warn
- 2 - Error

By default, the log level is 0.

It is recommended to use 0, 1, or 2 log level depending on your requirement. Before you deploy NetBackup operator, you can change the log levels using `operator_patch.yaml`.

After deployment if user changes operator log level, to reflect it, user has to perform the following steps:

- Apply the operator changes using the `kubectl apply -k <operator-folder>` command.
- Restart the operator pod. Delete the pod using the `kubectl delete pod/<netbackup-opertaor-pod-name> -n <namespace>` command. Kubernetes will recreate the NetBackup operator pod again after deletion.

- Config-Checker jobs that run before deployment of primary server and media server creates the pod. The logs for config checker executions can be checked using the `kubectl logs <configchecker-pod-name> -n <netbackup-operator-namespace>` command.
- Installation logs of NetBackup primary server and media server can be retrieved using any of the following methods:
  - Run the `kubectl logs <PrimaryServer/MediaServer-Pod-Name> -n <PrimaryServer/MediaServer namespace>` command.
  - Execute the following command in the primary server/media server pod and check the `/mnt/nblogs/setup-server.log` file:  

```
kubectl exec -it <PrimaryServer/MediaServer-Pod-Name> -n <PrimaryServer/MediaServer-namespace> -- bash
```
- Data migration jobs create the pods that run before deployment of primary server. The logs for data migration execution can be checked using the following command:  

```
kubectl logs <migration-pod-name> -n <netbackup-environment-namespace>
```
- Execute the following respective commands to check the event logs that shows deployment logs for PrimaryServer and MediaServer:
  - For primary server: `kubectl describe PrimaryServer <PrimaryServer name> -n <PrimaryServer-namespace>`
  - For media server: `kubectl describe MediaServer<MediaServername> -n<MediaServer-namespace>`

## Expanding storage volumes

You can update storage capacity of already created persistent volume claim for primary server and media server. Expanding storage volume for particular replica of respective CR object is not supported. In case of media server user needs to update volumes for all the replicas of particular media server object.

PVC will expand as per the new size and it will be available to volume mounts in primaryServer pod.

### To expand volume of data and log volumes for primary and media server

Amazon EFS is an elastic file system, it does not enforce any file system capacity limits. The actual storage capacity value in persistent volumes and persistent volume claims is not used when creating the file system. However, because storage capacity

is a required field in Kubernetes, you must specify a valid value. This value does not limit the size of your Amazon EFS file system.

- 1 Edit the environment custom resource using the `kubectl edit Environment <environmentCR_name> -n <namespace>` command.
- 2 To pause the reconciler of the particular custom resource, change the *paused: false* value to *paused: true* in the `primaryServer` or `mediaServer` section and save the changes. In case of multiple media server objects change `Paused` value to `true` for respective media server object only.
- 3 Edit `StatefulSet` of primary server or particular media server object using the `kubectl edit <statfulset name> -n <namespace>` command, change replica count to 0 and wait for all pods to terminate for the particular CR object.
- 4 Update all the persistent volume claim which expects capacity resize with the `kubectl edit pvc <pvcName> -n <namespace>` command. In case of particular media server object, resize respective PVC with expected storage capacity for all its replicas.
- 5 Update the respective custom resource section using the `kubectl edit Environment <environmentCR_name> -n <namespace>` command with updated storage capacity for respective volume and change *paused: false*. Save updated custom resource.

To update the storage details for respective volume, add storage section with specific volume and its capacity in respective `primaryServer` or `mediaServer` section in environment CR.

Earlier terminated pod and `StatefulSet` must get recreated and running successfully. Pod should get linked to respective persistent volume claim and data must have been persisted.

- 6 Run the `kubectl get pvc -n <namespace>` command and check for **capacity** column in result to check the persistent volume claim storage capacity is expanded.
- 7 (Optional) Update the log retention configuration for NetBackup depending on the updated storage capacity.

For more information, refer to the *NetBackup™ Administrator's Guide*, Volume I

## Allocating static PV for Primary and Media pods

When you want to use a disk with specific performance parameters, you can statically create the PV and PVC. You must allocate static PV and PVC before deploying the NetBackup server for the first time.

## To allocate static PV for Media and Primary pods

### 1 Create storage class in cluster.

See “[How does the Config-Checker utility work](#)” on page 61.

This newly created storage class name is used while creating PV and PVC's and should be mentioned for Catalog, Log, Data volume in the environment CR mediaServer section at the time of deployment.

For more information on creating the storage class, see [Storage class](#).

For example,

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
 name: gp2-reclaim
provisioner: kubernetes.io/aws-ebs
reclaimPolicy: Retain
allowVolumeExpansion: true
volumeBindingMode: Immediate
parameters:
 fsType: ext4
 type: gp2
```

For more information about the static provisioning for EFS, see [Static Provisioning](#).

### 2 Calculate number of disks required.

The following persistent volumes are required by Media pods:

- Data and Log volume disk per replica of media server.

Use the following format to form PVC names.

For primary server:

For media server

- data-[resourceNamePrefix\\_of\\_media](#)-media-[media server replica number](#). Count starts from 0>
- logs-[resourceNamePrefix\\_of\\_media](#)-media-[media server replica number](#). Count starts from 0>

**Example 1** If user wants to deploy a media server with replica count 3. For this scenario, you must create total 8 disks, 8 PV and 8 PVCs.

Name of the Media PVC assuming resourceNamePrefix\_of\_media is **testmedia**. 6 disks, 6 PV and 6 PVCs for media server.

For data of PrimaryServer:

For logs:

- logs-testprimary-primary-0

Following will be the names for media server volumes

For data:

- data-testmedia-media-0
- data-testmedia-media-1
- data-testmedia-media-2

For log:

- logs-testmedia-media-0
- logs-testmedia-media-1
- logs-testmedia-media-2

|                  |                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Example 2</b> | If user wants to deploy a media server with replica count 5                       | For this scenario, you must create 12 disks, 12 PV and 12 PVCs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                  | Names of the Media PVC assuming resourceNamePrefix_of_media is <b>testmedia</b> . | 10 disks, 10 PV and 10 PVCs for media server.<br>Following will be the names for primary server volumes<br>For data:<br>For logs:<br>Following will be the names for media server volumes<br>For data: <ul style="list-style-type: none"><li>■ data-testmedia-media-0</li><li>■ data-testmedia-media-1</li><li>■ data-testmedia-media-2</li><li>■ data-testmedia-media-3</li><li>■ data-testmedia-media-4</li></ul> For log: <ul style="list-style-type: none"><li>■ logs-testmedia-media-0</li><li>■ logs-testmedia-media-1</li><li>■ logs-testmedia-media-2</li><li>■ logs-testmedia-media-3</li><li>■ logs-testmedia-media-4</li></ul> |

- 3** Create the required number of AWS EBS volumes and save the VolumeId of newly created volumes.

For more information on creating EBS volumes, see [EBS volumes](#).

(For Primary Server volumes) Create the required number of EFS. User can use single EFS to mount catalog of primary. For example, VolumeHandle in **PersistentVolume** spec will be as follows:

```
<file_system_id>:/catalog
```

#### 4 Create PVs for each disks.

To create the PVs, specify the created storage class and VolumeID (ID of the EBS volumes received in step 3). The PV must be created using the **claimRef** field and provide PVC name for its corresponding namespace.

For example, if you are creating PV for catalog volume, storage required is 128GB and namespace is **test**. PVC named **catalog-testprimary-primary-0** is linked to this PV when PVC is created in the namespace test.

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: catalog
spec:
 accessModes:
 - ReadWriteMany
 awsElasticBlockStore:
 fsType: xfs
 volumeID: aws://us-east-2c/vol-xxxxxxxxxxxxxxxxxxxxx
 capacity:
 storage: 128Gi
 persistentVolumeReclaimPolicy: Retain
 storageClassName: gp2-retain
 volumeMode: Filesystem
 claimRef:
 apiVersion: v1
 kind: PersistentVolumeClaim
 name: catalog-testprimary-primary-0
 namespace: test
```

- 5 Create PVC with correct PVC name (step 2), storage class and storage.

For example,

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
 name: catalog-testprimary-primary-0
 namespace: test
spec:
 storageClassName: gp2-retain
accessModes:
 - ReadWriteMany
resources:
 requests:
 storage: 128Gi
```

- 6 Deploy the Operator.
- 7 Use previously created storage class names for the volumes in primary section and mediaServers section in environment CR spec and deploy environment CR.

# Monitoring MSDP Scaleout

This chapter includes the following topics:

- [About MSDP Scaleout status and events](#)
- [Monitoring with Amazon CloudWatch](#)
- [The Kubernetes resources for MSDP Scaleout and MSDP operator](#)

## About MSDP Scaleout status and events

The MSDP Scaleout CR status includes the readiness state, the storage space utilization (via PersistentVolumeClaim) of each Controller, MDS, and Engine pod.

In the initial configuration of MSDP Scaleout, the readiness state of each pod changes from "false" to "true" in the first few minutes. When the state of all the pods changes to "true", it indicates MSDP Scaleout is ready for use.

You can check the storage space utilization routinely to plan MSDP Scaleout autoscaling before the storage space runs out.

## To check the MSDP Scaleout status and events

### 1 Check the status and the events under the namespace for MSDP Scaleout.

```
kubectl -n <sample-namespace> describe msdp-scaleout
<sample-cr-name>
```

### 2 Check the MSDP Scaleout events.

```
kubectl -n <sample-namespace> get events
[--sort-by='{.lastTimestamp}']
```

### 3 Check the storage space utilization.

```
kubectl -n <sample-namespace> get msdp-scaleout <sample-cr-name>
-o json
```

#### Example of the of the status format:

```
kubectl -n sample-cr-namespace get msdp-scaleout sample-cr -o json
```

```
kubectl get -n demo msdp-scaleout msdp-app -o json |jq .status
{
 "controllers": [
 {
 "apiVersions": [
 "1.0"
],
 "name": "msdp-app-uss-controller",
 "nodeName": "ip-x-x-x-x.ec2.internal",
 "productVersion": "16.0.1-0035",
 "pvc": [
 {
 "pvcName": "msdp-app-uss-controller-log",
 "stats": {
 "availableBytes": "9878.00Mi",
 "capacityBytes": "9951.27Mi",
 "percentageUsed": "0.58%",
 "usedBytes": "57.27Mi"
 }
 }
],
 "ready": "True"
 }
],
 "engines": [
 {
 "ip": "x.x.x.x",
```

```
"name": "ip-x-x-x-x.ec2.internal",
"nodeName": "ip-x-x-x-x.ec2.internal",
"pvc": [
 {
 "pvcName": "ip-x-x-x-x.ec2.internal-catalog",
 "stats": {
 "availableBytes": "604539.68Mi",
 "capacityBytes": "604629.16Mi",
 "percentageUsed": "0.01%",
 "usedBytes": "73.48Mi"
 }
 },
 {
 "pvcName": "ip-x-x-x-x.ec2.internal-data-0",
 "stats": {
 "availableBytes": "4160957.62Mi",
 "capacityBytes": "4161107.91Mi",
 "percentageUsed": "0.00%",
 "usedBytes": "134.29Mi"
 }
 }
],
"ready": "True"
},
```

## Monitoring with Amazon CloudWatch

You can use Amazon CloudWatch to collect Prometheus metrics to monitor pods in MSDP-X cluster.

### To configure Amazon CloudWatch

- 1 Install the CloudWatch agent with Prometheus metrics collection on EKS.  
See [AWS documentation](#).
- 2 Install the CloudWatch agent on EKS clusters. Select the EC2 launch type, and download the template YAML file `Prometheus-eks.yaml`.

### 3 Add the YAML file with the following sample configuration.

```
create configmap for prometheus cwagent config
apiVersion: v1
data:
 # cwagent json config
 cwagentconfig.json: |
 {
 "logs": {
 "metrics_collected": {
 "prometheus": {
 "prometheus_config_path": "/etc/prometheusconfig/
prometheus.yaml",
 "emf_processor": {
 "metric_declaration": [
 {
 "source_labels": ["job"],
 "label_matcher": "^msdoperator-metrics",
 "dimensions": [
 ["ClusterName", "NameSpace"]
],
 "metric_selectors": [
 "msdoperator_reconcile_failed",
 "msdoperator_operator_run",
 "msdoperator_diskFreeLess5GBEngines_total",
 "msdoperator_diskFreeMiBytesInEngine",
 "msdoperator_diskFreeLess10GBClusters_total",
 "msdoperator_totalDiskFreePercentInCluster",
 "msdoperator_diskFreePercentInEngine",
 "msdoperator_pvcFreePercentInCluster",
 "msdoperator_unhealthyEngines_total",
 "msdoperator_createdPods_total"
]
 }
]
 }
 }
 }
 },
 "force_flush_interval": 5
 }
kind: ConfigMap
metadata:
```

```
name: prometheus-cwagentconfig
namespace: amazon-cloudwatch
```

```

```

```
create configmap for prometheus scrape config
```

```
apiVersion: v1
```

```
data:
```

```
 # prometheus config
```

```
 prometheus.yaml: |
```

```
 global:
```

```
 scrape_interval: 1m
```

```
 scrape_timeout: 10s
```

```
 scrape_configs:
```

```
 - job_name: 'msdpoperator-metrics'
```

```
 scheme: https
```

```
 tls_config:
```

```
 ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
```

```
 insecure_skip_verify: true
```

```
 bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/
 /token
```

```
 kubernetes_sd_configs:
```

```
 - role: pod
```

```
 relabel_configs:
```

```
 - source_labels: [__meta_kubernetes_pod_annotation_prometheus_io_
 _scrape]
```

```
 action: keep
```

```
 regex: true
```

```
 - source_labels: [__meta_kubernetes_pod_annotation_prometheus_io_
 _path]
```

```
 action: replace
```

```
 target_label: __metrics_path__
```

```
 regex: (.+)
```

```
 - source_labels: [__address__, __meta_kubernetes_pod_annotation_
 prometheus_io_port]
```

```
 action: replace
```

```
 regex: ([^:]+)(?::\d+)?;(\d+)
```

```
 replacement: $1:$2
```

```
 target_label: __address__
```

```
 - source_labels: [__meta_kubernetes_namespace]
```

```
 action: replace
```

```

 target_label: NameSpace
 - source_labels: [__meta_kubernetes_pod_name]
 action: replace
 target_label: PodName

kind: ConfigMap
metadata:
 name: prometheus-config
 namespace: amazon-cloudwatch

```

[Table 11-1](#) lists the Prometheus metrics that MSDP Scaleout supports.

#### 4 Apply the YAML file.

```
Kubectl apply -f Prometheus-eks.yaml
```

The default log groups name is

```
/aws/containerinsights/{cluster_name}/Prometheus.
```

#### 5 Create Amazon CloudWatch alarms.

See [Using Amazon CloudWatch alarms](#) in AWS documentation.

#### 6 In the CloudWatch console, add the related log query. In the navigation pane, select **Log Insights**.

For example, the free space size of the MSDP scaleout cluster engines is lower than 1 GB in past 5 minutes. Select the log group from the drop-down list, and select the time duration 5m on the time line.

Log query:

```

fields @timestamp, @message
| filter msdpoperator_diskFreeMiBytesInEngine <= 100000
| sort @timestamp desc

```

If multiple MSDP scaleout clusters are deployed in the same EKS cluster, use the filter to search the results. For example, search the MSDP engines with the free space size lower than 1GB in the namespace **sample-cr-namespace**.

Log query:

```

fields @timestamp, @message
| filter msdpscalout_ns == "sample-cr-namespace"
| filter msdpoperator_diskFreeMiBytesInEngine <= 100000
| sort @timestamp desc

```

MSDP Scaleout supports the following Prometheus metrics:

**Table 11-1** Supported Prometheus metrics list in MSDP Scaleout

| Metrics                                     | Type    | Filters                   | Description                                                                                      |
|---------------------------------------------|---------|---------------------------|--------------------------------------------------------------------------------------------------|
| msdpoperator_reconcile_total                | Counter | N/A                       | The total of the reconcile loops msdp-operator run.                                              |
| msdpoperator_reconcile_failed               | Counter | N/A                       | The total of the reconcile loops msdp-operator failed to run.                                    |
| msdpoperator_operator_run                   | Counter | N/A                       | The total of the running operator.                                                               |
| msdpoperator_diskFreeLess5GBEngines_total   | Gauge   | msdpscalout_ns            | The checked number of the engines which have free spaces lower than 5GB.                         |
| msdpoperator_diskFreeMiBytesInEngine        | Gauge   | msdpscalout_ns            | The free space of current engine in MiBytes.                                                     |
| msdpoperator_diskFreeLess10GBClusters_total | Gauge   | msdpscalout_ns            | The checked number of the msdp scaleout applications that have free spaces lower than 10GB.      |
| msdpoperator_totalDiskFreePercentInCluster  | Gauge   | msdpscalout_ns            | The percent of the msdp scaleout applications that have free spaces. For example, 0.95 means 95% |
| msdpoperator_diskFreePercentInEngine        | Gauge   | msdpscalout_ns            | The percent of the current engines, which have free spaces.                                      |
| msdpoperator_pvcFreePercentInCluster        | Gauge   | msdpscalout_ns, component | The percent of the used PVC, which have free spaces.                                             |
| msdpoperator_unhealthyEngines_total         | Gauge   | msdpscalout_ns            | The total of unhealthy engines.                                                                  |
| msdpoperator_createdPods_total              | Gauge   | msdpscalout_ns, component | The total of created msdp scaleout pods.                                                         |

## The Kubernetes resources for MSDP Scaleout and MSDP operator

Do not change or delete the Kubernetes resources that MSDP deployment has created.

- Run the following command to find all the namespaced resources:

```
kubectl api-resources --verbs=list --namespaced=true -o name |
xargs -n 1 -i bash -c 'if ! echo {} |grep -q events; then kubectl
```

```
get --show-kind --show-labels --ignore-not-found -n <cr or operator namespace> {}; fi'
```

- Run the following command to find commonly used namespace resources:

```
kubectl get pod,svc,deploy,rs,pvc -n <cr or operator namespace> -o wide
```

- Run the following command to find the Kubernetes cluster level resources that belong to the CR:

```
kubectl api-resources --verbs=list --namespaced=false -o name | xargs -n 1 -i bash -c 'kubectl get --show-kind --show-labels --ignore-not-found {} |grep [msdp-operator|<cr-name>]'
```

# Monitoring Snapshot Manager deployment

This chapter includes the following topics:

- [Overview](#)
- [Logs of Snapshot Manager](#)
- [Configuration parameters](#)

## Overview

The status of Snapshot Manager deployment can be verified by using the following command:

```
kubectl describe cpserver -n $ENVIRONMENT_NAMESPACE
```

This displays the status of deployment as follows:

| Status  | Description                |
|---------|----------------------------|
| Running | Deployment is in progress. |
| Failed  | Deployment has failed.     |
| Success | Deployment is successful.  |

## Logs of Snapshot Manager

Fluent log collector service collects the logs from various services in Snapshot Manager at one shared storage. To read these services, exec into `flexsnap-fluend-collector pod`.

Run the `kubectl` command as follows:

```
kubectl exec -it <flexsnap-fluend-collector pod_name> bash -n
$ENVIRONMENT_NAMESPACE
```

You can find the Snapshot Manager log files under `/cloudpoint/logs/` folder.

## Configuration parameters

- Any configuration related parameter that must be added in `/cloudpoint/flexsnap.conf` file can be added in `flexsnap-conf` configmap by editing it as follows:

```
kubectl edit configmap flexsnap-conf -n $ENVIRONMENT_NAMESPACE
```

For example, for changing the log level from info to debug, add the following:

```
[logging]
level = debug
```

- Any configuration related parameter which needs to be added in `/cloudpoint/openv/netbackup/bp.conf` file can be added in `nbuconf` configmap by editing it as follows:

```
kubectl edit configmap nbuconf -n $ENVIRONMENT_NAMESPACE
```

# Managing the Load Balancer service

This chapter includes the following topics:

- [About the Load Balancer service](#)
- [Notes for Load Balancer service](#)
- [Opening the ports from the Load Balancer service](#)

## About the Load Balancer service

Key features of the Load Balancer service:

- Load balancer services are created in primary server and media server deployment that allows you to access the NetBackup application from public domains.
- In primary server or media server CR spec, networkLoadBalancer section is used for handling the IP address and DNS name allocation for load balancer services. This section combines to sub fields **type**, **annotations**, and **ipList** whereas these fields are optional. If **ipList** is provided in CR spec, IP address count must match the replica count in case of media server CR whereas in case of primary server CR, only one IP address needs to be mentioned.
- NetBackup supports the network load balancer with AWS Load Balancer scheme as **internet-facing**.
- The networkLoadBalancer section can be used to provide static IP address and dns name allocation to the Load Balancer services.
- FQDN must be created before being used. Refer below sections for different allowed annotations to be used in CR spec.

- User must add the following annotations:

```
service.beta.kubernetes.io/aws-load-balancer-subnets: <subnet1 name>
```

In addition to the above annotations, if required user can add more annotations supported by AWS. For more information, see [AWS Load Balancer Controller Annotations](#).

Example: CR spec in primary server,

```
networkLoadBalancer:
 type: Private
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-subnets: <subnet1 name>
 ipList:
 "10.244.33.27: abc.vxindia.veritas.com"
```

CR spec in media server,

```
networkLoadBalancer:
 type: Private
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-subnets: <subnet1 name>
 ipList:
 "10.244.33.28: pqr.vxindia.veritas.com"
 "10.244.33.29: xyz.vxindia.veritas.com"
```

The IP address, the subnet provided in ipList and annotations in networkLoadBalancer section in CR spec must belong to same availability zone that of the node group.

---

**Note:** The subnet provided here should be same as the one given in node pool used for primary server and media server.

---

If NetBackup client is outside VPC or to access Web UI from outside VPC, then client CIDR must be added with all NetBackup ports in security group rule of cluster. Run the following command, to obtain the cluster security group:

```
aws eks describe-cluster --name <my-cluster> --query
cluster.resourcesVpcConfig.clusterSecurityGroupId
```

For more information on cluster security group, see [Amazon EKS security group requirements and considerations](#).

Add inbound rule to security group. For more information, see [Add rules to a security group](#).

## Default ports used in the Load Balancer service

- Primary server:
  - 1556  
Used as bidirectional port. Primary server to/from media servers and primary server to/from client require this TCP port for communication.
  - 8443  
Used to inbound to java nbwmc on the primary server.
  - 443  
Used to inbound to vnet proxy tunnel on the primary server. Also, this is used Nutanix workload, communication from primary server to the deduplication media server.
  - 13781  
The MQBroker is listening on TCP port 13781. NetBackup client hosts - typically located behind a NAT gateway - be able to connect to the message queue broker (MQBroker) on the primary server.
  - 13782  
Used by primary server for bpcd process.
  - Port 22  
Used by NetBackup IT Analytics data collector for data collection.
- Media server:
  - 1556  
Used as bidirectional port. Primary server to/from media servers and primary server to/from client require this TCP port for communication.
  - 13782  
Used by media server for bpcd process.
  - 443  
The Snapshot Manager user interface uses this port as the default HTTPS port.
  - 5671  
The Snapshot Manager RabbitMQ server uses this port for internal service communications. This port must be open to support multiple agents, extensions, backup from snapshot, and restore from backup jobs.
  - 2049  
It is used for Amazon EFS access.  
For more information, see [Source ports for working with EFS](#).

## Notes for Load Balancer service

Note the following points:

- After deployment of primary server or media server, updating the DNS name, IP address and subnet through CR is not allowed.
- If mistakenly user has added wrong values:
  - User wants to update IP address and subnet, you must delete the CR and update the CR yaml and reapply it.
  - User wants to update the DNS name, you must delete the respective CR and delete the respective PVC and PV as well.

---

**Note:** Be caution while performing this step, this may lead to data loss.

---

- Before using the DNS and its respective IP address in CR yaml, you can verify the IP address and its DNS resolution using nslookup.
- In case of media server scaleout, ensure that the number of IP addresses mentioned in IPList in networkLoadBalancer section matches the replica count.
- If nslookup is done for loadbalancer IP inside the container, it returns the DNS in the form of **<svc name>.<namespace\_name>.svc.cluster.local**. This is Kubernetes behavior. Outside the pod, the loadbalancer service IP address is resolved to the configured DNS. The `nbbptestconnection` command inside the pods can provide a mismatch in DNS names, which can be ignored.

## Opening the ports from the Load Balancer service

In this deployment, most of the required ports are already opened from the NetBackup primary and media server load balancer services by default.

- If you want to use a specific workload and that needs specific ports, you must add those ports in the port specification of the load balancer service.
- In case of media server, you must add custom ports in the load balancer service of all the replicas. In case of scaling up the media server, user needs to explicitly add newly added custom ports in respective newly created load balancer services.
- In case custom ports are added in the load balancer service and the same load balancer service is deleted or created again, you must add respective custom ports again in the load balancer service specification.

For all three scenarios, perform the steps given in this section.

### To open the ports from the Load Balancer service

- 1 Run the `kubectl get service -n <namespace>` command.

This command lists all the services available in given namespace.

- 2 Edit the required primary or media load balancer service using `kubectl edit service <service-name> -n <namespace>` command.

For example:

- For primary server load balancer service:
  - Service name starts with **Name** of primary server like **<Name>-primary**. Edit the service with the `kubectl edit service <Name>-primary -n <namespace>` command.
- For media server load balancer service:
  - Each replica of media server has its own load balancer service with name **<Name>-media-<ordinal number>**. For example, replica 2 of media server has a load balancer service with name **<Name>-media-1**.
  - You must modify service for specific replica with the `kubectl edit service <Name>-media-<replica-ordinal number> -n <namespace>` command.

---

**Note:** The load balancer service with name **Name** used in primary sever and media server specification must be unique.

---

- 3 Add entry for new port in ports array in specification field of the service. For example, if user want to add 111 port, then add the following entry in ports array in specification field.

```
name: custom-111

 port: 111

 protocol: TCP

 targetPort: 111
```

- 4 Save the changes.

The service is updated and the new port is listed in ports list of the respective service when you run the `kubectl get service -n <namespace>` command.

# Performing catalog backup and recovery

This chapter includes the following topics:

- [Backing up a catalog](#)
- [Restoring a catalog](#)

## Backing up a catalog

You can backup a catalog.

### To backup a catalog

- 1 Exec into the primary server pod using the following command:  

```
kubectl exec -it -n <namespace> <primary-pod-name> -- /bin/bash
```
- 2 Create a directory **DRPackages** at persisted location using `mkdir`  
`/mnt/nblogs/DRPackages`.
- 3 Change ownership of **DRPackages** folder to service user using `chown`  
`nbsvcusr:nbsvcusr /mnt/nblogs/DRPackages`.
- 4 Set the passphrase to be used at time of catalog recovery.
  - Open NetBackup Administrator Console (Java UI).
  - Navigate to **Security Management > Global Security Setting > Disaster Recovery**.
  - In **Encryption for Disaster Recovery** section, add the passphrase, confirm passphrase, and save it.

- 5 Add respective external media server entry in host properties through **NetBackup Management > Host properties > Master Server > Add Additional server.**

---

**Note:** It is recommended to use an external media server for catalog backup and recovery.

---

- 6 Exec into the primary server pod using the following command:  

```
kubectl exec -it -n <namespace> <primaryserver pod name> -- bash
```

Set the **KMS\_CONFIG\_IN\_CATALOG\_BKUP** configuration option to 1 in `/usr/openv/netbackup/bp.conf` file of primary server to include the KMS configuration as part of the disaster recovery package during catalog backup.

- 7 Restart the NetBackup services in primary and external media server.

- Exec into the primary server pod using the following command:  

```
kubectl exec -it -n <namespace> <primary-pod-name> -- /bin/bash
```
- Deactivate NetBackup health probes using the `/opt/veritas/vxapp-manage/nbu-health deactivate` command.
- Run the `/usr/openv/netbackup/bin/bp.kill_all` command. After stopping all services restart the services using the `/usr/openv/netbackup/bin/bp.start_all` command.
- Activate NetBackup health probes using the `/opt/veritas/vxapp-manage/nbu-health activate` command.
- Run the `/usr/openv/netbackup/bin/bp.kill_all` command. After stopping all services restart the services using the `/usr/openv/netbackup/bin/bp.start_all` command on the external media server.

- 8 Configure storage unit on earlier added external media server.

For more information, refer to the *NetBackup™ Administrator's Guide, Volume 1*

---

**Note:** It is recommended to use AdvancedDisk or BasicDisk storage unit.

---

- 9 Configure NetBackup catalog backup policy.  
Add package path as `/mnt/nblogs/DRPackages` while configuring the catalog backup policy.
- 10 Run the catalog backup job.

## Restoring a catalog

You can restore a catalog.

### To restore a catalog

- 1 Copy DRPackages files (packages) located at `/mnt/nblogs/DRPackages/` from the pod to the host machine from where Amazon Elastic Kubernetes Service cluster is accessed.  
  
Run the `kubectl cp`  
`<primary-pod-namespace>/<primary-pod-name>:/mnt/nblogs/DRPackages`  
`<Path_where_to_copy_on_host_machine>` command.
- 2 Preserve the data of `/mnt/nbdata` and `/mnt/nblogs` on host machine by creating tar and copying it using the `kubectl cp`  
`<primary-pod-namespace>/<primary-pod-name>:<tar_file_name>`  
`<path_on_host_machine_where_to_preserve_the_data>` command.
- 3 Change CR spec from *paused: false* to *paused: true* in primary, mediaServers, and msdpScaleouts sections in `environment.yaml` and re-apply yaml using the `kubectl apply -f environment.yaml -n <namespace>` command.
- 4 Change replica count to 0 in primary server's statefulset using the `kubectl edit statefulset <primary-server-statefulset-name> -n <namespace>` command.
- 5 Get names of PV attached to primary server PVC (catalog, log and data) using the `kubectl get pvc -n <namespace> -o wide` command.
- 6 Delete primary server PVC (catalog, log and data) using the `kubectl delete pvc <pvc-name> -n <namespace>` command.
- 7 Delete the PV linked to primary server PVC using the `kubectl delete pv <pv-name>` command.
- 8 Navigate to mounted EFS directory and delete the content from **primary\_catalog** folder by running the `rm -rf /efs/*` command.
- 9 Change CR spec *paused: true* to *paused: false* in primary server section in `environment.yaml` and reapply yaml with the `kubectl apply -f environment.yaml -n <namespace>` command.

- 10 After the primary server pod is in ready state, change CR spec from *paused: false* to *paused: true* in primary server section in `environment.yaml` and reapply yaml with the `kubectl apply -f environment.yaml -n <namespace>` command.
- 11 Execute the `kubectl exec -it -n <namespace> <primary-pod-name> -- /bin/bash` command in the primary server pod.
  - Increase the debug logs level on primary server.
  - Create a directory `DRPackages` at persisted location using `mkdir /mnt/nblogs/DRPackages`.
  - Change ownership of the `DRPackages` folder to service user using the `chown nbsvcusr:nbsvcusr /mnt/nblogs/DRPackages` command.
- 12 Copy earlier copied DR files to primary pod at `/mnt/nblogs/DRPackages` using the `kubectl cp <Path_of_DRPackages_on_host_machine> <primary-pod-namespace>/<primary-pod-name>:/mnt/nblogs/DRPackages` command.
- 13 Execute the following steps in the primary server pod.
  - Change ownership of files in `/mnt/nblogs/DRPackages` using the `chown nbsvcusr:nbsvcusr <file-name>` command.
  - Deactivate NetBackup health probes using the `/opt/veritas/vxapp-manage/nbu-health deactivate` command.
  - Stop the NetBackup services using `/usr/opensv/netbackup/bin/bp.kill_all`.
  - Execute the `nghostidentity -import -infile /mnt/nblogs/DRPackages/<filename>.drpkg` command.
  - Restart all the NetBackup services using `/usr/opensv/netbackup/bin/bp.start_all`.
- 14 Verify security settings are back.
- 15 Add respective media server entry in host properties using NetBackup Administration Console.
  - Navigate to **NetBackup Management > Host properties > Master Server > Add Additional server** and add media server.
- 16 Restart the NetBackup services in primary server pod and external media server.
  - Execute the following command in the primary server pod:  
`kubectl exec -it -n <namespace> <primary-pod-name> -- /bin/bash`

- Run the `/usr/opensv/netbackup/bin/bp.kill_all` command. After stopping all services restart the same using the `/usr/opensv/netbackup/bin/bp.start_all` command.
  - Run the `/usr/opensv/netbackup/bin/bp.kill_all` command. After stopping all services restart the services using the `/usr/opensv/netbackup/bin/bp.start_all` command on the external media server.
- 17** Configure a storage unit on external media server that is used during catalog backup.
- 18** Perform catalog recovery from NetBackup Administration Console.  
For more information, refer to the [NetBackup Troubleshooting Guide](#).
- 19** Execute the `kubectl exec -it -n <namespace> <primary-pod-name> -- /bin/bash` command in the primary server pod.
- Stop the NetBackup services using the `/usr/opensv/netbackup/bin/bp.kill_all` command.
  - Start NetBackup services using the `/usr/opensv/netbackup/bin/bp.start_all` command.
  - Activate NetBackup health probes using the `/opt/veritas/vxapp-manage/nbu-health activate` command.
- 20** Change CR spec from `paused: true` to `paused: false` in `primary`, `mediaServers`, and `msdpScaleouts` sections in `environment.yaml` and re-apply yaml using the `kubectl apply -f environment.yaml -n <namespace>` command.
- 21** To configure NetBackup IT Analytics refer to the following topic.  
See [“Configuring NetBackup IT Analytics for NetBackup deployment”](#) on page 81.

# Managing MSDP Scaleout

This chapter includes the following topics:

- [Adding MSDP engines](#)
- [Adding data volumes](#)
- [Expanding existing data or catalog volumes](#)
- [MSDP Scaleout scaling recommendations](#)
- [MSDP Cloud backup and disaster recovery](#)
- [MSDP multi-domain support](#)
- [Configuring Auto Image Replication](#)
- [About MSDP Scaleout logging and troubleshooting](#)

## Adding MSDP engines

You can add new MSDP engines by updating the CR. You can add maximum 16 MSDP engines.

Prerequisites:

- Allocate new static IP/FQDN pairs in the same node resource group.
- The node number must not be less than the MSDP Scaleout size that you plan to change.
- CR YAML file of MSDP Scaleout

**To add the MSDP engines by updating the CR YAML file**

- 1 Open the CR YAML file to edit.
- 2 Append the new IP/FQDN pairs in the `spec.serviceIPFQDNs` field.

- 3 Update the **spec.size** field to increase the cluster size accordingly.
- 4 Apply new CR YAML to update the CR in the Kubernetes environment.

```
kubectl apply -f <your-cr-yaml>
```

#### To add the MSDP engines using the kubectl command directly

- ◆ Run the following command to append the IP/FQDN pairs in the **spec.serviceIPFQDNs** field and increase the cluster size in **spec.size** field.

```
kubectl -n <sample-namespace> edit msdp-scaleout <your-cr-name>
[-o json | yaml]
```

The MSDP Scaleout services are not interrupted when MSDP engines are added.

## Adding data volumes

You can add the data volumes by updating the CR.

#### To add the data volumes by updating the CR YAML file

- 1 Open the CR YAML file to edit.
- 2 Append the new data volume specifications in the **spec.dataVolumes** field.
- 3 Apply new CR YAML to update the CR in the Kubernetes environment.

```
kubectl apply -f <your-cr-yaml>
```

#### To add the MSDP engine using the kubectl command directly

- ◆ Run the following command to append new data volume specifications in the **spec.dataVolumes** field.

```
kubectl -n <sample-namespace> edit msdp-scaleout <your-cr-name>
[-o json | yaml]
```

In the MSDP engine pod, the first data volume is mounted on `/msdp/data/dp1/pdvol`. Nth data volume is mounted on `/msdp/data/dp1/${N-1}pdvol`. For example, 2nd data volume is mounted on `/msdp/data/dp1/1pdvol`.

Each MSDP engine can support up to 16 data volumes.

It is recommended that you use the same data volume size if you add multiple volumes.

---

**Note:** Due to some Kubernetes restrictions, MSDP operator restarts the engine pods for attaching the existing and new volumes, which can cause the short downtime of the services.

---

# Expanding existing data or catalog volumes

You can expand the existing data or catalog volumes by updating the CR.

## To expand the data or catalog volumes by updating the CR YAML file

- 1 Open the CR YAML file to edit.
- 2 Increase the requested storage size in the `spec.dataVolumes` field or in the `spec.catalogVolume` field.
- 3 Apply new CR YAML to update the CR in the Kubernetes environment.

```
kubectl apply -f <your-cr-yaml>
```

## To expand the data or catalog volumes using the kubectl command directly

- ◆ Run the following command to increase the requested storage size in the `spec.dataVolumes` field or in the `spec.catalogVolume` field..

```
kubectl -n <sample-namespace> edit msdpyscaleout <your-cr-name>
[-o json | yaml]
```

Sometimes Amazon EBS CSI driver may not respond the volume expansion request promptly. In this case, the operator retries the request by adding 1 byte to the requested volume size to trigger the volume expansion again. If it is successful, the actual volume capacity could be slightly larger than the requested size.

Due to the limitation of Amazon EBS CSI driver, the engine pods need to be restarted for resizing the existing volumes. This can cause the short downtime of the services.

MSDP Scaleout does not support the following:

- Cannot shrink the volume size.
- Cannot change the existing data volumes other than for storage expansion.
- Cannot expand the log volume size. You can do it manually. See [“Manual storage expansion”](#) on page 158.
- Cannot expand the data volume size for MDS pods. You can do it manually. See [“Manual storage expansion”](#) on page 158.

## Manual storage expansion

You also can manually expand storage size by expanding PVC size.

### To expand the data or catalog volumes

- 1 Open the CR YAML file to edit.
- 2 Configure `spec.paused: true`.

- 3 Apply new CR YAML to stop MSDP operator from reconciling and repairing the pods automatically.

```
kubectl apply -f <your-cr-yaml>
```

- 4 Patch the corresponding PVCs.

```
kubectl patch pvc <pvc-name> --type merge --patch '{"spec":
{"resources": {"requests": {"storage": "<requested-size>"}}}'
-n <sample-namespace>
```

- 5 Specify `spec.paused: false` in the CR.
- 6 Apply new CR YAML to recover MSDP operator to continue to reconcile and repair the pods automatically.

```
kubectl apply -f <your-cr-yaml>
```

---

**Note:** If you add new MSDP Engines later, the new Engines will respect the CR specification only. Your manual changes would not be respected by the new Engines.

---

## MSDP Scaleout scaling recommendations

Following are the scaling recommendations for the MSDP Scaleout:

- Allocate the data volumes of the similar sizes for MSDP to have better load balancing performance.
- Each data volume size is more than 4 TB.
- Have multiple data volumes for each engine to gain better throughput.
- Split a bigger backup policy to smaller ones  
In most cases, one backup job goes to one MSDP engine at the same time even if multistream is enabled for the backup policy. If the current MSDP engine, which is taking a backup job hits the high space watermark, the following backup data would be sent to a second MSDP engine. If the backup data is too big for up to 2 MSDP engines to persist, the backup job fails. When more MSDP engines are added, the backup jobs may not be evenly balanced on each MSDP engine at the first a few hours or days. If the situation keeps longer beyond your expectation, consider to re-plan the backup policies, by splitting a bigger backup policy to two smaller ones, to help MSDP Scaleout to balance the new backup jobs more faster.
- After scaling up, the memory and CPU of the existing node group may not meet the performance requirements anymore. In this case, you can add more memory and CPU by upgrading to the higher instance type to improve the existing node group performance or create another node group with higher instance type and

update the node-selector for the CR accordingly. If you create another node group, the new node-selector does not take effect until you manually delete the pods and deployments from the old node group, or delete the old node group directly to have the pods re-scheduled to the new node group.

- Ensure that each EKS node supports mounting the number of data volumes plus 5 of the data disks.  
For example, if you have 16 data volumes for each engine, then each your EKS node should support mounting at least 21 data disks. The additional 5 data disks are for the potential MDS pod, Controller pod or MSDP operator pod to run on the same node with MSDP engine.

## MSDP Cloud backup and disaster recovery

For information about MSDP cloud backup and disaster recovery, see MSDP Cloud section of the *NetBackup Deduplication Guide*.

---

**Note:** In case of disaster recovery of NetBackup environment (that is, primary, media and MSDP), perform the primary catalog recovery first and then proceed with MSDP disaster recovery steps. See [“Backing up a catalog”](#) on page 151.

---

### About the reserved storage space

About 1 TB storage space is reserved by default on each MSDP engine for each cloud LSU.

The 1 TB storage space is selected from one of the data volumes of every engine. It requires each engine at least has one data volume, which has more than 1 TB available storage space, when a cloud LSU is to be configured. Otherwise, the configuration of the cloud LSU fails.

## Cloud LSU disaster recovery

### **Scenario 1: MSDP Scaleout and its data is lost and the NetBackup primary server remains unchanged and works well**

- 1 Redeploy MSDP Scaleout on a EKS cluster by using the same CR parameters and NetBackup re-issue token.
- 2 When MSDP Scaleout is up and running, re-use the cloud LSU on NetBackup primary server.

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <STORAGESEVERNAME> -stype PureDisk -configlist
<configuration file>
```

Credentials, bucket name, and sub bucket name must be the same as the recovered Cloud LSU configuration in the previous MSDP Scaleout deployment.

Configuration file template:

```
V7.5 "operation" "reuse-lsu-cloud" string
V7.5 "lsuName" "LSUNAME" string
V7.5 "lsuCloudUser" "XXX" string
V7.5 "lsuCloudPassword" "XXX" string
V7.5 "lsuCloudAlias" "<STORAGESEVERNAME_LSUNAME>" string
V7.5 "lsuCloudBucketName" "XXX" string
V7.5 "lsuCloudBucketSubName" "XXX" string
V7.5 "lsuKmsServerName" "XXX" string
```

If the LSU cloud alias does not exist, you can use the following command to add it.

```
/usr/opensv/netbackup/bin/admincmd/csconfig cldinstance -as -in
<instance-name> -sts <storage-server-name> -lsu_name <lsu-name>
```

- 3** On the first MSDP Engine of MSDP Scaleout, run the following command for each cloud LSU:

```
sudo -E -u msdpvc /usr/opencv/pdde/pdcr/bin/cacontrol --catalog
clouddr <LSUNAME>
```

- 4** Restart the MSDP services in the MSDP Scaleout.

Option 1: Manually delete all the MSDP engine pods.

```
kubectl delete pod <sample-engine-pod> -n <sample-cr-namespace>
```

Option 2: Stop MSDP services in each MSDP engine pod. MSDP service starts automatically.

```
kubectl exec <sample-engine-pod> -n <sample-cr-namespace> -c
uss-engine -- /usr/opencv/pdde/pdconfigure/pdde stop
```

**Scenario 2: MSDP Scaleout and its data is lost and the NetBackup primary server was destroyed and is re-installed**

- 1 Redeploy MSDP Scaleout on an EKS cluster by using the same CR parameters and new NetBackup token.
- 2 When MSDP Scaleout is up and running, reuse the cloud LSU on NetBackup primary server.

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -setconfig
-storage_server <STORAGESEVERNAME> -stype PureDisk -configlist
<configuration file>
```

Credentials, bucket name, and sub bucket name must be the same as the recovered Cloud LSU configuration in previous MSDP Scaleout deployment.

Configuration file template:

```
V7.5 "operation" "reuse-lsu-cloud" string
V7.5 "lsuName" "LSUNAME" string
V7.5 "lsuCloudUser" "XXX" string
V7.5 "lsuCloudPassword" "XXX" string
V7.5 "lsuCloudAlias" "<STORAGESEVERNAME_LSUNAME>" string
V7.5 "lsuCloudBucketName" "XXX" string
V7.5 "lsuCloudBucketSubName" "XXX" string
V7.5 "lsuKmsServerName" "XXX" string
```

If KMS is enabled, setup KMS server and import the KMS keys.

If the LSU cloud alias does not exist, you can use the following command to add it.

```
/usr/opensv/netbackup/bin/admincmd/csconfig cldinstance -as -in
<instance-name> -sts <storage-server-name> -lsu_name <lsu-name>
```

- 3 On the first MSDP Engine of MSDP Scaleout, run the following command for each cloud LSU:

```
sudo -E -u msdpvc /usr/opensv/pdde/pdcr/bin/cacontrol --catalog
cloudrr <LSUNAME>
```

- 4 Restart the MSDP services in the MSDP Scaleout.

Option 1: Manually delete all the MSDP engine pods.

```
kubectl delete <sample-engine-pod> -n <sample-cr-namespace>
```

Option 2: Stop MSDP services in each MSDP engine pod.

```
kubectl exec <sample-engine-pod> -n <sample-cr-namespace> -c
uss-engine -- /usr/opensv/pdde/pdconfigure/pdde stop
```

- 5 Create disk pool for the cloud LSU on NetBackup server.
- 6 Do two-phase image importing.

See the *NetBackup Administrator's Guide, Volume I*

For information about other DR scenarios, see *NetBackup Deduplication Guide*.

## MSDP multi-domain support

An MSDP storage server is configured in a NetBackup media server. The NetBackup media servers and clients in the NetBackup domain can use this storage server. By default, the NetBackup media servers and clients cannot directly use an MSDP storage server from another NetBackup domain. For example, NetBackup media servers or clients cannot backup data to an MSDP storage server from another NetBackup domain.

To use an MSDP storage server from another NetBackup domain, the MSDP storage server must have multiple MSDP users. Then NetBackup media servers or clients can use the MSDP storage server from another NetBackup domain by using a different MSDP user. Multiple NetBackup domains can use the same MSDP storage server but each NetBackup domain must use a different MSDP user to access that MSDP storage server.

For more information, See *NetBackup Deduplication Guide*.

When you add a new MSDP user, the command `spause` must be executed in the first MSDP engine of MSDP Scaleout, not on any of the NetBackup servers.

Ensure that you run MSDP commands with non-root user `msdpsvc` after logging into an engine pod.

For example, `sudo -E -u msdpsvc /usr/opensv/pdde/pdcr/bin/spause`

## Configuring Auto Image Replication

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication (A.I.R.).

You can configure Auto Image Replication in NetBackup, which is using MSDP Scaleout storage servers.

### To configure Auto Image Replication

- 1 Logon to the NetBackup Web UI of both replication source and target domain.
- 2 Add each other NetBackup's primary server as trusted primary server.  
 For more information, see the *NetBackup Web UI Administrator's Guide*.
- 3 In the replication source domain, get the MSDP\_SERVER name from the NetBackup Web UI.  
 Navigate to **Storage > Storage configuration > Storage servers**.
- 4 Add MSDP\_SERVER in the primary server of replication target domain. Login to the target primary server and run the following command:
 

```
echo "MSDP_SERVER = <Source MSDP server name>" >>
/usr/openv/netbackup/bp.conf
```
- 5 Get the token from the target domain NetBackup Web UI.  
 Navigate to **Security > Token**. In the **Create token** window, enter the token name and other required details. Click **Create**.  
 For more information, see the *NetBackup Web UI Administrator's Guide*.
- 6 Add replication targets for the disk pool in replication source domain.  
 In the **Disk pools** tab, click on the disk pool link.  
 Click **Add** to add the replication target.
- 7 In the **Add replication targets** window:
  - Select the replication target primary server.
  - Provide the target domain token.
  - Select the target volume.
  - Provide the target storage credentials.
 Click **Add**.

## About MSDP Scaleout logging and troubleshooting

- EKS troubleshooting  
 See [Amazon EKS troubleshooting](#) page of the *AWS documentation*.
- Logs and core dumps files in MSDP Scaleout  
 MSDP Operator, Controller, and MDS pod logs are stored in `/log` location of the pods.

- Collect the logs and inspection information  
 You can collect the logs and inspection information for MSDP Scaleout for troubleshooting purpose.  
 See [“Collecting the logs and the inspection information”](#) on page 166.

## Collecting the logs and the inspection information

You can collect the logs and inspection information for MSDP Scaleout for troubleshooting purpose.

Run the command `kubect1 msdp collect-logs`

For example, `kubect1 msdp collect-logs -o <output path> [-n <MSDP operator namespace>] [-c <MSDP applications namespace(s)>]`

**Table 15-1** collect-logs command options

| Option | Description                                                                                                                                                                                                                                                                                           |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -c     | Comma-separated namespaces of MSDP applications.<br><b>Note:</b> If not specified, it collects MSDP applications of all namespaces.                                                                                                                                                                   |
| -f     | Output format of logs/core files/MSDP history files.<br>Available options:<br>targz: Copy logs/core files/MSDP history files from containers and compress them by tar/gzip.<br>raw: Copy logs/core files/MSDP history files from containers as same format in the containers.<br>Default value: targz |
| -n     | Namespace of MSDP operator.<br>Default value: msdp-operator-system                                                                                                                                                                                                                                    |
| -o     | Output path of the log file.                                                                                                                                                                                                                                                                          |

# About MSDP Scaleout maintenance

This chapter includes the following topics:

- [Pausing the MSDP Scaleout operator for maintenance](#)
- [Logging in to the pods](#)
- [Reinstalling MSDP Scaleout operator](#)
- [Migrating the MSDP Scaleout to another node group](#)

## Pausing the MSDP Scaleout operator for maintenance

For maintenance purpose, if you want the operator to stop reconciling the resources of one CR but do not affect the resources of the other CRs, you can pause the MSDP Scaleout operator.

### To pause the MSDP Scaleout operator

- 1 Specify `spec.paused: true` in the CR.
- 2 Run `kubectl apply -f <sample CR YAML>`.

Do not forcibly delete the deployment resource of MSDP Scaleout operator.

## Logging in to the pods

You can log in to the pods for the maintenance purpose.

To log in to the pod, run the `kubectl` executable file.

Run MSDP commands with non-root user **msdpsvc** after logging in to an engine pod.

For example, `sudo -E -u msdpsvc <command>`

The MSDP Scaleout services in an engine pods are running with non-root user **msdpsvc**. If you run the MSDP Scaleout services or commands with the root user, MSDP Scaleout may stop working due to file permissions issues.

## Reinstalling MSDP Scaleout operator

When you undeploy MSDP Scaleout operator, the MSDP Scaleout CRD is removed from the EKS cluster. It also deletes all the existing MSDP Scaleout on the EKS cluster. The PVC for the operator logs is also deleted. However, the MSDP Scaleout critical data and metadata is not deleted.

### To reinstall MSDP Scaleout operator

- 1 Run the following command to delete the MSDP Scaleout operator:

```
kubectl msdp delete [-k] [-n <sample-operator-namespace>]
```

- 2 Run the following command to redeploy the operator.

```
kubectl msdp init -i <your-acr-url>/msdp-operator:<version> -s
<storage-class-name> -l agentpool=<nodegroup-name> [-n
<sample-operator-namespace>]
```

- 3 If the reclaim policy of the storage class is **Retain**, run the following command to restart the existing MSDP Scaleout. MSDP Scaleout starts with the existing data/metadata.

```
kubectl apply -f <your-cr-yaml>
```

## Migrating the MSDP Scaleout to another node group

You can migrate an existing MSDP Scaleout on another node group in case of the Kubernetes infrastructure issues.

### To migrate the MSDP Scaleout to another node group

- 1 Ensure that no job running related to MSDP Scaleout that is going to migrate.
- 2 Update the node selector value **spec.nodeSelector** to the new node in the CR YAML file.

- 3 Apply new CR YAML to update the CR in the Kubernetes environment.

```
kubect1 apply -f <your-cr-yaml>
```

---

**Note:** All affected pods or other Kubernetes workload objects must be restarted for the change to take effect.

---

- 4 After the CR YAML file update, existing pods are terminated and restarted one at a time, and the pods are re-scheduled for the new node group automatically.

---

**Note:** Controller pods are temporarily unavailable when the MDS pod restarts. Do not delete pods manually.

---

- 5 Run the following command to change MSDP Scaleout operator to the new node group:

```
kubect1 msdp init -i <your-acr-url>/msdp-operator:<version> -s
<storage-class-name> -l agentpool=<new-nodegroup-name>
```

- 6 If node selector does not match any existing nodes at the time of change, you see the message on the console.

If auto scaling for node is enabled, it may resolve automatically as the new nodes are made available to the cluster. If invalid node selector is provided, pods may go in the pending state after the update. In that case, run the command above again.

Do not delete the pods manually.

# Uninstalling MSDP Scaleout from EKS

This chapter includes the following topics:

- [Cleaning up MSDP Scaleout](#)
- [Cleaning up the MSDP Scaleout operator](#)

## Cleaning up MSDP Scaleout

When you uninstall the MSDP Scaleout deployment from EKS, the MSDP engines, MSDP MDS servers, and the data is deleted from the cluster. The data is lost and cannot be recovered.

## To clean up MSDP Scaleout from EKS

- 1 Delete the MSDP Scaleout CR.

```
kubectl delete -f <sample-cr-yaml>
```

When an MSDP Scaleout CR is deleted, the critical MSDP data and metadata is not deleted. You must delete it manually. If you delete the CR without cleaning up the data and metadata, you can re-apply the same CR YAML file to restart MSDP Scaleout again by reusing the existing data.

- 2 If your storage class is with the **Retain** policy, you must write down the PVs that are associated with the CR PVCs for deletion in the Kubernetes cluster level.

```
kubectl get
pod,svc,deploy,rs,ds,pvc,secrets,certificates,issuers,cm,sa,role,rolebinding
-n <sample-namespace> -o wide

kubectl get clusterroles,clusterrolebindings,pv -o wide
--show-labels|grep <sample-cr-name>
```

- 3 Delete all resources under the namespace where MSDP CR is deployed.

```
kubectl delete namespace <namespace>
```

- 4 If your storage class is with the **Retain** policy, you must delete the EBS volumes on Amazon console or using the AWS CLI.

```
aws ec2 delete-volume --volume-id <value>
```

See [“Deploying MSDP Scaleout”](#) on page 114.

See [“Reinstalling MSDP Scaleout operator”](#) on page 168.

## Cleaning up the MSDP Scaleout operator

You can delete the MSDP Scaleout operator to remove all related resources about MSDP Scaleout operator. The MSDP Scaleout operator and logs are deleted.

## To clean up MSDP Scaleout operator from EKS

- 1 If your storage class is with **Retain** policy, write down the PVs that are associated with the Operator PVCs for deletion in the Kubernetes cluster level.

```
kubectl get
pod,svc,deploy,rs,ds,pvc,secrets,certificates,issuers,cn,sa,role,rolebinding
-n <sample-operator-namespace> -o wide

kubectl get clusterroles,clusterrolebindings,pv -o wide
--show-labels
```

- 2 Delete the MSDP Scaleout operator.

```
kubectl msdp delete [-n <sample-operator-namespace>].
```

- `-k`: Delete all resources of MSDP Scaleout operator except the namespace.
- `-n`: Namespace scope for this request.  
Default value: `msdp-operator-system`

- 3 If your storage class is with the **Retain** policy, you must delete the EBS volumes on Amazon console or using the AWS CLI.

```
aws ec2 delete-volume --volume-id <value>
```

See [“Deploying MSDP Scaleout”](#) on page 114.

See [“Reinstalling MSDP Scaleout operator”](#) on page 168.

# Uninstalling Snapshot Manager

This chapter includes the following topics:

- [Uninstalling Snapshot Manager from EKS](#)

## Uninstalling Snapshot Manager from EKS

When you uninstall Snapshot Manager from EKS, the Snapshot Manager related services are deleted from the cluster.

1. Delete cpServer related parameters from `environment.yaml` file and apply it.

Following commands can be used to remove and disable the Snapshot Manager from NetBackup:

```
ENVIRONMENT_NAMESPACE="netbackup-environment"
Make sure the flexsnap-operator pod is running and ready.
Comment out / remove cpServer part from environment.yaml then apply it.

kubectl apply -f environment.yaml -n $ENVIRONMENT_NAMESPACE sleep
10s
```

2. Ensure that you get the uninstall message in `flexsnap-operator operator` log.
3. To clean-up cpServer component, delete flexsnap specific persistent volumes (PVs), persistent volume claims (PVCs) and config maps. Note that these resources contain metadata of current cpServer installation and would be deleted.

Use the following respective commands to delete these resources:

- **To remove** `certauth-pvc`, `cloudpoint-pvc` **and** `mongodb-pvc`: # `kubectl delete pvc certauth-pvc cloudpoint-pvc mongodb-pvc -n <nb-namespace>`
- **To remove associated PVs of** `certauth-pvc`, `cloudpoint-pvc` **and** `mongodb-pvc`: # `kubectl delete pv <pv-name> -n <nb-namespace>`
- **To remove** `nbuconf` **and** `flexsnap-conf`: # `kubectl delete cm nbuconf flexsnap-conf -n <nb-namespace>`

# Troubleshooting

This chapter includes the following topics:

- [View the list of operator resources](#)
- [View the list of product resources](#)
- [View operator logs](#)
- [View primary logs](#)
- [Pod restart failure due to liveness probe time-out](#)
- [Socket connection failure](#)
- [Resolving an invalid license key issue](#)
- [Resolving an issue where external IP address is not assigned to a NetBackup server's load balancer services](#)
- [Resolving the issue where the NetBackup server pod is not scheduled for long time](#)
- [Resolving an issue where the Storage class does not exist](#)
- [Resolving an issue where the primary server or media server deployment does not proceed](#)
- [Resolving an issue of failed probes](#)
- [Resolving token issues](#)
- [Resolving an issue related to insufficient storage](#)
- [Resolving an issue related to invalid nodepool](#)
- [Resolving a token expiry issue](#)

- [Resolve an issue related to KMS database](#)
- [Resolve an issue related to pulling an image from the container registry](#)
- [Resolving an issue related to recovery of data](#)
- [Check primary server status](#)
- [Pod status field shows as pending](#)
- [Ensure that the container is running the patched image](#)
- [Getting EEB information from an image, a running container, or persistent data](#)
- [Resolving the certificate error issue in NetBackup operator pod logs](#)
- [Resolving the primary server connection issue](#)
- [Primary pod is in pending state for a long duration](#)
- [Host mapping conflict in NetBackup](#)
- [NetBackup messaging queue broker take more time to start](#)
- [Local connection is getting treated as insecure connection](#)
- [Issue with capacity licensing reporting which takes longer time](#)
- [Backing up data from Primary server's /mnt/nbdata/ directory fails with primary server as a client](#)
- [Wrong EFS ID is provided in environment.yaml file](#)
- [Primary pod is in ContainerCreating state](#)
- [Webhook displays an error for PV not found](#)

## View the list of operator resources

To view all the operator resources, execute the following command on Kubernetes cluster:

```
$ kubectl get all -n netbackup-operator-system
```

The output should be something like this:

| NAME                                | READY | STATUS | RESTARTS | AGE  |
|-------------------------------------|-------|--------|----------|------|
| pod/msdp-operator-                  |       | 2/2    | Running  | 0    |
| controller-manager-65d8fd7c4d-whqpm |       |        |          | 3h6m |

```

pod/netbackup-operator- 2/2 Running 0 93m
controller-manager-55d6bf59c8-vltmp

NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S) AGE
service/msdp-operator- ClusterIP 10.96.144.99 <none> 8443/TCP 3h6m
controller-manager-
metrics-service

service/msdp-operator- ClusterIP 10.96.74.75 <none> 443/TCP 3h6m
webhook-service

service/netbackup- ClusterIP 10.96.104.94 <none> 8443/TCP 93m
operator-controller
-manager-metrics-service

service/netbackup- ClusterIP 10.96.210.26 <none> 443/TCP 93m
operator-webhook-service

NAME READY UP-TO-DATE AVAILABLE AGE
deployment.apps/msdp- 1/1 1 1 3h6m
operator-controller-manager

deployment.apps/netbackup 1/1 1 1 93m
-operator-controller-manager

NAME DESIRED CURRENT READY AGE
replicaset.apps/msdp- 1 1 1 3h6m
operator-controller-
manager-65d8fd7c4d
replicaset.apps/netbackup- 1 1 1 93m
operator-controller-manager-
55d6bf59c8

```

Verify that both pods display **Running** in the Status column and both deployments display **2/2** in the **Ready** column.

## View the list of product resources

To view the list of product resources run the following command:

```

$ kubectl get --namespace <namespace>
all,environments,primaryservers,mediaservers,msdpscaleouts

```

The output should look like the following:

| NAME                                        | READY | STATUS  | RESTARTS | AGE |
|---------------------------------------------|-------|---------|----------|-----|
| pod/dedupe1-uss-controller-79d554f8cc-598pr | 1/1   | Running | 0        | 68m |
| pod/dedupe1-uss-mds-1                       | 1/1   | Running | 0        | 75m |
| pod/dedupe1-uss-mds-2                       | 1/1   | Running | 0        | 74m |
| pod/dedupe1-uss-mds-3                       | 1/1   | Running | 0        | 71m |
| pod/media1-media-0                          | 1/1   | Running | 0        | 53m |
| pod/environment-sample-primary-0            | 1/1   | Running | 0        | 86m |
| pod/x10-240-0-12.veritas.internal           | 1/1   | Running | 0        | 68m |
| pod/x10-240-0-13.veritas.internal           | 2/2   | Running | 0        | 64m |
| pod/x10-240-0-14.veritas.internal           | 2/2   | Running | 0        | 61m |
| pod/x10-240-0-15.veritas.internal           | 2/2   | Running | 0        | 59m |

| NAME                               | TYPE         | CLUSTER-IP   | PORT(S)                                                                                                      | AGE |
|------------------------------------|--------------|--------------|--------------------------------------------------------------------------------------------------------------|-----|
| service/dedupe1-uss-controller     | ClusterIP    | 10.1.109.118 | 10100/TCP                                                                                                    | 68m |
| service/dedupe1-uss-mds            | ClusterIP    | None         | 2379/TCP,2380/TCP                                                                                            | 75m |
| service/dedupe1-uss-mds-client     | ClusterIP    | 10.1.5.208   | 2379/TCP                                                                                                     | 75m |
| service/media1-media-0             | LoadBalancer | 10.1.121.115 | 13782:30648/TCP,<br>1556:30248/TCP                                                                           | 54m |
| service/environment-sample-primary | LoadBalancer | 10.1.206.39  | 13781:30246/TCP,<br>13782:30498/TCP,<br>1556:31872/TCP,<br>443:30049/TCP,<br>8443:32032/TCP,<br>22:31511/TCP | 87m |
| service/x10-240-0-12               |              |              |                                                                                                              |     |

```

-veritas-internal LoadBalancer 10.1.44.188 10082:31199/TCP 68m
service/
x10-240-0-13
-veritas-internal LoadBalancer 10.1.21.176 10082:32439/TCP, 68m
service/
x10-240-0-14 10102:30284/TCP
-veritas-internal LoadBalancer 10.1.25.99 10082:31810/TCP, 68m
service/
x10-240-0-15 10102:31755/TCP
-veritas-internal LoadBalancer 10.1.185.135 10082:31664/TCP, 68m
 10102:31811/TCP

```

```

NAME READY UP-TO-DATE AVAILABLE AGE
deployment.apps/dedupe1
-uss-controller 1/1 1 1 68m

```

```

NAME DESIRED CURRENT READY AGE
replicaset.apps/dedupe1-uss
-controller-79d554f8cc 1 1 1 68m

```

```

NAME READY AGE
statefulset.apps/medial-media 1/1 53m
statefulset.apps/environment
-sample-primary 1/1 86m

```

```

NAME TAG AGE STATUS
primaryserver.netbackup
.veritas.com/environment
-sample 10.1 88m Success

```

```

NAME TAG AGE PRIMARY SERVER STATUS
mediaserver.netbackup.
veritas.com/medial 10.1 54m .veritas.internal Success

```

```

NAME AGE TAG SIZE READY
msdpscaleout.msdp.
veritas.com/dedupe1 75m 17.0 4 4

```

```

NAME READY AGE STATUS
environment.netbackup.
veritas.com/
environment-sample 3/3 88m Success

```

An environment is deployed successfully if all pods and environment CR display status as "Success".

## View operator logs

If environment deployment status is not successful, check operator logs for errors.

Command for MSDP Scaleout operator logs

```
$ kubectl logs pod/msdp-operator-controller-manager-65d8fd7c4d-whqpm
manager -n netbackup-operator-system-c manager
```

Command for NetBackup operator logs

```
$ kubectl logs
pod/netbackup-operator-controller-manager-55d6bf59c8-vltmp
netbackup-operator -n netbackup-operator-system
```

## View primary logs

To view primary server logs execute the following command to get a shell to the running container.

```
$ kubectl exec --stdin --tty pod/<primary-server-pod-name> -n
<namespace> -- /bin/bash
```

Once in the primary server shell prompt, to see the list of logs, run:

```
ls /usr/opensv/logs/
```

## Pod restart failure due to liveness probe time-out

As part of liveness probe for primary and media pods, a health script runs inside the container to check the NetBackup health status.

When there is an issue with a container related to a full disk, CPU, or memory pressure, the liveness probe gets timed out because of no response from the health script. As a result, the Pod does not restart.

To resolve this issue, restart the Pod manually. Delete the Pod using the `kubectl delete pod/<podname> -n <namespace>` command.

The Pod is deleted and Kubernetes creates another Pod.

# Socket connection failure

Socket connection failure can happen because of the following reasons:

- Long processing delays
- AWS connection reset (default 4 minutes)
- Load on CPU or Memory pressure
- IO saturation and throttling under load

If there are problems with the TCP stacks on the hosts, network between the hosts, or unusual long processing delays, then the connection may drop and the TCP stack on the host is unaware of the situation.

The following error is displayed in the web UI under job details:

```
db_FLISTsend failed: unexpected message received (43)
*** - Error bptm (pid=14599) get_string() failed,
Connection reset by peer (104), network read error
*** - Info bptm (pid=14599) EXITING with status 42 <-----
*** - Info nbux-systest-media-1 (pid=14599)
StorageServer=PureDisk:nbux-systest-media-1;
Report=PDDO Stats for (nbux-systest-media-1):
scanned: 4195521 KB, CR sent: 171002 KB, CR sent over FC: 0 KB,
dedup: 95.9%, cache disabled, where dedup space saving:6.6%,
compression space saving:89.3%
*** - Info bpbkar (pid=19109) done. status: 42: network read failed
```

To resolve this issue, update the `sysctl.conf` values for NetBackup servers deployed on the EKS cluster.

NetBackup image sets following values in `sysctl.conf` during EKS deployment:

- `net.ipv4.tcp_keepalive_time = 180`
- `net.ipv4.tcp_keepalive_intvl = 10`
- `net.ipv4.tcp_keepalive_probes = 20`
- `net.ipv4.ip_local_port_range = 14000 65535`

These settings are persisted at the location `/mnt/nbdata/etc/sysctl.conf`.

There are two ways to modify these values:

- Modify the value in both `/etc/sysctl.conf` and `/mnt/nbdata/etc/sysctl.conf` and run the `sysctl -p` command to load the modified values.

- Modify the values in `/mnt/nbdata/etc/sysctl.conf` and restart the pod. The new values are reflected after the pod restart.

If external media servers are used, perform the steps in the following order:

1. Add the following in `/usr/opensv/netbackup/bp.conf`:
 

```
HOST_HAS_NAT_ENDPOINTS = YES
```
2. Add the following `sysctl` configuration values in `etc/sysctl.conf` on external media servers to avoid any socket connection issues:
  - `net.ipv4.tcp_keepalive_time = 180`
  - `net.ipv4.tcp_keepalive_intvl = 10`
  - `net.ipv4.tcp_keepalive_probes = 20`
  - `net.ipv4.ip_local_port_range = 14000 65535`
  - `net.core.somaxconn = 4096`
3. Save the setting using the `sysctl -p` command.

## Resolving an invalid license key issue

The NetBackup is not installed because the license key is invalid.

Pod remains in running state for long time and the installation log at `/mnt/nblogs/setup-server.log` displays the following error:

```
ERROR: No valid license key for NetBackup Server or Enterprise Server
```

When you deploy NetBackup for the first time, perform the steps for primary CR and media CR.

### To resolve an invalid license key issue for Primary CR

- 1 Get the configmap name created for primary CR or media CR using the following command:
 

```
kubectl get configmap -n <namespace>
```
- 2 Edit the license key stored in configmap using the following command:
 

```
kubectl edit configmap <primary-configmap-name> -n <namespace>
```
- 3 Update value for **ENV\_NB\_LICKEY** key in the configmap with correct license key and save.

- 4 Delete respective primary/media pod using the following command:

```
kubectl delete pod<primary-pod-name> -n <namespace>
```

New pod is auto created with updated license key value.

- 5 Edit environment CR with updated license key and save using the following command:

```
kubectl edit environments.netbackup.veritas.com -n <namespace>
```

## Resolving an issue where external IP address is not assigned to a NetBackup server's load balancer services

The issue can be because of one of the following reasons:

- The **resourcePrefixName** mentioned in custom resource is not unique and valid.
- The static IP is provided in **networkLoadBalancer** section in CR but it is not created in EKS.
- The subnet or resource group is mentioned in annotations of **networkLoadBalancer** section in CR spec, the IP address is not available in given subnet or resource group.
- The RBAC permissions in your cluster for the given subnet or resource group are not assigned properly for allocating IP addresses.

**To resolve an issue where external IP address is not assigned to a NetBackup server's load balancer services**

- 1 Check the event logs of load balancer service using the `kubectl describe service <svc-name> -n <namespace>` command for detailed error information.
- 2 Run the `kubectl edit Environment <environmentCR-name> -n <namespace>` command.
- 3 Depending on the output of the command and the reason for the issue, perform the required steps and update the environment CR to resolve the issue.

## Resolving the issue where the NetBackup server pod is not scheduled for long time

The NetBackup server (primary server and media server) pods are stuck in Pending state. The issue can be because of one of the following reasons:

- Insufficient resource allocation.
- Persistent volume claims are not bound to persistent volume.
- NetBackup server pods have the anti-affinity rule added.

As a result, primary server and media server pods are scheduled on different nodes. If nodes are not available, pod remains in pending state with event logs indicating nodes are scaling up, if auto scaling is configured in cluster.

### To resolve the issue where the NetBackup server pod is not scheduled for long time

- 1 Check the pod event details for more information about the error using `kubectl describe <PrimaryServer/MediaServer_Pod_Name> -n <namespace>` command.
- 2 Depending on the output of the command and the reason for the issue, perform the required steps and update the environment CR to resolve the issue.

## Resolving an issue where the Storage class does not exist

The Config-Checker checks if the storage class name given in primary server/media server CR is available in the cluster.

The following error is displayed:

```
Error: ERROR Storage class with the <storageClassName> name does not exist.
```

After fixing this error, primary server or media server CR does not require any changes. In this case, NetBackup operator reconciler loop is invoked after every 10 hours. If you want to reflect the changes and invoke the NetBackup operator reconciler loop immediately, delete and reapply the primary server or media server CR.

**To resolve an issue where the Storage class does not exist**

- 1 Create storage class with the same name given in primary server or media server CR with ReclaimPolicy as **Retain** in the cluster.

To create storage class, refer to the following link:

[Amazon Elastic Kubernetes Service storage classes](#)

In this scenario, no change in primary server or media server CR is required. As a result, reconciler loop is not invoked immediately.

- 2 To invoke the reconciler loop again, delete the respective CR.

If it is primary server CR, use the `kubectl delete -f <environment.yaml>` command, or if it is media server CR, edit the Environment CR by removing the media server section in the `environment.yaml`.

---

**Note:** To reuse the `mediaServer` section information, you must save it and apply the yaml again with the new changes using the `kubectl apply -f <environment.yaml>` command.

---

- 3 Apply the CR again.

If it is primary server CR then reapply it using the `kubectl apply -f <environment.yaml>` command or if it is media server CR, add `mediaServer` section that was deleted earlier with required data in `environment.yaml` and reapply it using the `kubectl apply -f <environment.yaml>` command.

## Resolving an issue where the primary server or media server deployment does not proceed

primary server or media server deployment does not proceed even if `configcheckmode = default` in primary server or media server CR spec and no other child resources are created. It is possible that the Config-Checker job has failed some of the configuration checks.

**To resolve an issue where the primary server or media server deployment does not proceed**

- 1 Check the status of Config-Checker **Configcheckerstatus** mentioned in primary server or media server CR status using the `kubectl describe <PrimaryServer/MediaServer> <CR name> -n <namespace>` command.

If the state is **failed**, check the Config-Checker pod logs.

- 2 Retrieve the Config-Checker pod logs using the `kubectl logs <config-checker-pod-name> -n <operator-namespace>` command.

Config-Checker pod name can be in the following format:

`<serverType>-configchecker-<configcheckermode>-randomID`, for example if its Config-Checker for primary server with **configcheckermode = default**, pod name is `primary-configcehcker-default-dhg34`.

- 3 Depending on the error in the pod logs, perform the required steps and edit the environment CR to resolve the issue.
- 4 Data migration jobs create the pods that run before deployment of primary server. Data migration pod exist after migration for one hour only if data migration job failed. The logs for data migration execution can be checked using the following command:

```
kubectl logs <migration-pod-name> -n
<netbackup-environment-namespace>
```

User can copy the logs to retain them even after job pod deletion using the following command:

```
kubectl logs <migration-pod-name> -n
<netbackup-environment-namespace> > jobpod.log
```

## Resolving an issue of failed probes

If pod is not in ready state for log time, the `kubectl describe pod/<podname> -n <namespace>` command displays the following errors:

- Readiness probe failed: The readiness of the external dependencies is not set.  
  
Server setup is still in progress.
- Liveness probe failed: bpps command did not list nbwmc process. nbwmc is not alive.  
  
The Primary server is unhealthy.

### To resolve an issue of failed probes

- 1 If you are deploying NetBackup on Amazon Elastic Kubernetes Service Cluster for the first time, check the installation logs for detailed error.

Use any of the following methods:

- Execute the following command in the respective primary server or media server pod and check the logs in `/mnt/nblogs/setup-server.logs`:

```
kubectl exec -it -n <namespace> <pod-name> -- /bin/bash
```

- Run the `kubectl logs pod/<podname> -n <namespace>` command.

- 2 Check pod events for obtaining more details for probe failure using the following command:

```
kubectl describe pod/<podname> -n <namespace>
```

Kubernetes will automatically try to resolve the issue by restarting the pod after liveness probe times out.

- 3 Depending on the error in the pod logs, perform the required steps or contact technical support.

## Resolving token issues

Media server installation log displays the following error in `/mnt/nblogs/setup-server.logs`:

```
nbcertcmd: The -getCertificate operation
failed for server <primaryServerName>,
EXIT STATUS 5940: Reissue token is mandatory,
please provide a reissue token
```

NetBackup media server and NetBackup primary server were in running state. Media server persistent volume claim or media server pod is deleted. In this case, reinstallation of respective media server can cause the issue.

### To resolve token issues

- 1 Open the NetBackup web UI using primary server hostname given in the primary server CR status.
- 2 Navigate to **Security > Host Mappings**.
- 3 Click **Actions > Allow auto reissue certificate** for the respective media server name.

- 4 Delete data and logs PVC for respective media server only using the `kubectl delete pvc <pvc-name> -n <namespace>` command.

The persisted data is deleted.

- 5 Delete respective media server pod using `kubectl delete <pod-name> -n <namespace>` command.

New media server pod and new PVCs for the same media server are created.

## Resolving an issue related to insufficient storage

Setup-server.logs of NetBackup primary server displays an error.

Insufficient storage on the node can cause this issue. Minimum hardware requirements for NetBackup may not be completed. During fresh deployment of primary server, the following error is displayed in `/mnt/nblogs/setup-server.logs`:

```
DBSPAWN ERROR: -86
Not enough memory to start
```

### To resolve an issue related to insufficient storage

- 1 Create a new nodepool with the hardware specifications as mentioned in the NetBackup Deployment on EKS Administrator's Guide.
- 2 Run the `kubectl get nodes` command to ensure that the nodes from the newly created nodepool are used in your cluster.
- 3 Delete the primary server CR using the `kubectl delete -f <environment.yaml>` command.
- 4 Update **nodeSelector** spec in primary section and apply the `environment.yaml` again using the `kubectl apply -f <environment.yaml>` command.

## Resolving an issue related to invalid nodepool

Invalid nodepool is mentioned in primary server or media server CR nodeSelector spec. Due to this, primary server or media server pod fails to schedule.

The following error is displayed:

```
Error: Did not match Pod's node affinity/selector.
```

### To resolve an issue related to invalid nodepool

- 1 If you are deploying NetBackup on Amazon Elastic Kubernetes Service Cluster for the first time, delete the respective CR.

If it is primary server CR:

- Delete it using the `kubectl delete -f <environment.yaml>` command.
  - Update the node selector in primary server section in `environment.yaml` and apply it again using the `kubectl apply -f <environment.yaml>` command.
- 2 For media server CR: Delete the media server CR by removing the **mediaServer** section in the `environment.yaml` and save the changes.

---

**Note:** Ensure that you copy spec information of the media server CR. The spec information is used to reapply the media server CR.

---

- 3 Apply the new changes using the `kubectl apply -f <environment.yaml>` command.
- 4 Add the **mediaServer** section, update the `nodeSelector`, and reapply the `environment.yaml` using the `kubectl apply -f <environment.yaml>` command.

## Resolving a token expiry issue

While creating a new media pod, token may expire, and installation of media server is not completed. The installation logs at `/mnt/nblogs/setup-server.logs` display an error on the respective media server.

```
EXIT STATUS 5934: The token has expired.
```

### To resolve a token expiry issue

- 1 Edit the environment server CR using the `kubectl edit environment <environment-CR-name> -n <namespace>` command.
- 2 In the **mediaServer** section, reduce the replica count.

For example, if media pod with name `xyz-media-2` has the token expired issue and the replica was originally 3, then change the replica count to 2. Save the changes. The extra pods are deleted and `statefulset` displays new replica count in ready state (2/2).

**3** Edit the environment server CR using the `kubectl edit environment <environment-CR-name> -n <namespace>` command.

**4** Increase replica count to original replica count.

As given in the example, change the replica count to 3. This creates additional media pods and reissues the token for newly added media server.

## Resolve an issue related to KMS database

Installation logs at `/mnt/nblogs/setup-server.logs` display an error message with other details. In this scenario, you must configure KMS manually.

```
Error: Failed to create KMS database
```

To resolve this issue, execute the following command in the primary server pod:

```
kubectl exec -it -n <namespace> <primary-server-pod-name> -- /bin/bash
```

Refer the [NetBackup Security and Encryption Guide](#) for configure KMS manually:

For other troubleshooting issue related to KMS, refer the [NetBackup Troubleshooting Guide](#).

## Resolve an issue related to pulling an image from the container registry

Primary or media server failed to deploy with **ImagePullBackOff** error. If the pod Status field displays **ImagePullBackOff**, it means that the pod could not start because Kubernetes cannot pull a container image. A misspelled registry or image name or image registry being not reachable can cause a **ImagePullBackOff** status.

Run the `$ k get all -n netbackup-operator-system` command.

The output should look like:

| NAME                                                               | READY | STATUS           | RESTARTS | AGE  |
|--------------------------------------------------------------------|-------|------------------|----------|------|
| pod/msdp-operator<br>-controller-manager-<br>65d8fd7c4d-bsgms      | 2/2   | Running          | 0        | 7m9s |
| pod/netbackup-operator<br>-controller-manager-<br>5df6f58b9b-6ftt9 | 1/2   | ImagePullBackOff | 0        | 13s  |

For additional details, use the following command:

```
$ kubectl describe pod/<pod_name> -n netbackup-operator-system
```

Resolve this issue using any of the following methods:

- Check if image name and tag are correct. If not, edit and update the environment CR using the `kubectl edit environment <environment CR-name> -n <namespace>` command with correct image name and tag, and then save the changes.
- Check if the user is authorized and has permissions to access the AWS container registry.

## Resolving an issue related to recovery of data

If a PVC is deleted or the namespace where primary or media server is deployed, is deleted or deployment setup is uninstalled, and you want to recover the previous data, attach the primary server and media server PVs to its corresponding PVCs.

In case of recovering data from PV, you must use the same environment CR specs that are used at the time of previous deployment. If any spec field is modified, data recovery may not be possible.

### To resolve an issue related to recovery of data

- 1 Run the `kubectl get PV` command.
- 2 From the output list, note down PV names and its corresponding claim (PVC name and namespace) that are relevant from previous deployment point of view.
- 3 Set claim ref for the PV to null using the `kubectl patch pv <pv name> -p '{"spec":{"claimRef": null}}'` command.

For example, `kubectl patch pv`

```
pvc-4df282e2-b65b-49b8-8d90-049a27e60953 -p '{"spec":{"claimRef": null}}'
```

- 4 Run the `kubectl get PV` command and verify bound state of PVs is **Available**.

- For the PV to be claimed by specific PVC, add the **claimref spec** field with PVC name and namespace using the `kubectl patch pv <pv-name> -p '{"spec":{"claimRef": {"apiVersion": "v1", "kind": "PersistentVolumeClaim", "name": "<Name of claim i.e. PVC name>", "namespace": "<namespace of pvc>"}}}'` command.

For example,

```
kubectl patch pv <pv-name> -p '{"spec":{"claimRef": {"apiVersion": "v1", "kind": "PersistentVolumeClaim", "name": "data-testmedia-media-0", "namespace": "test"}}}'
```

While adding claimRef add correct PVC names and namespace to respective PV. Mapping should be as it was before deletion of the namespace or deletion of PVC.

- Deploy environment CR that deploys the primary server and media server CR internally.

## Check primary server status

Check the primary server custom resource status, with the command:

```
$ kubectl get primaryserver.netbackup.veritas.com/environment-sample -n <namespace>
```

| NAME               | TAG  | AGE  | STATUS |
|--------------------|------|------|--------|
| environment-sample | 10.0 | 3h1m | Failed |

If the output shows STATUS as *Failed* as in the example above, check the primary pod log for errors with the command:

```
$ kubectl logs pod/environment-sample-primary-0 -n <namespace>
```

## Pod status field shows as pending

If the pod Status field shows Pending state, it indicates that Kubernetes is not able to schedule the pod. To check use the following command:

```
$ kubectl get all -n netbackup-operator-system
```

The output is something like:

| NAME                                                  | READY | STATUS  | RESTARTS | AGE |
|-------------------------------------------------------|-------|---------|----------|-----|
| pod/msdp-operator-controller-manager-65d8fd7c4d-bsgms | 2/2   | Running | 0        | 12m |

```
pod/netbackup-
operator-controller-
manager-6c9dc8d87f
-pq8mr 0/2 Pending 0 15s
```

For more details use the following pod describe command:

```
$ kubectl describe
pod/netbackup-operator-controller-manager-6c9dc8d87f-pq8mr -n
netbackup-operator-system
```

The output is something like:

| Type    | Reason           | Age                 | Message                                                                                                                                                            |
|---------|------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ----    | -----            | ----                | -----                                                                                                                                                              |
| Warning | FailedScheduling | 56s (x3 over 2m24s) | 0/4 nodes are available:1 node(s) had taint {node-role.kubernetes.io/master: }, that the pod didn't tolerate, 3 node(s) didn't match Pod's node affinity/selector. |

To resolve this issue verify the nodeSelector settings in the operator/patch/operator\_patch.yaml file.

## Ensure that the container is running the patched image

There are three copies of the container image present in the Kubernetes environment during deployment or patching.

The first image copy is created on a local docker instance during image load operation. To check this copy, perform the following:

**1 Run:**

```
$ docker load -i images/pdk8soptr-17.0.tar.gz
```

Sample output:

```
Loaded image: msdp-operator:17.0
```

**2 Taking the image name from step 1, run:**

```
$ docker image ls | grep msdp-operator
```

Sample output:

```
msdp-operator 17.0 353d2bd50105 2 days ago 480 MB
```

**3 Taking the value from step 2, run:**

```
$ docker inspect 353d2bd50105 | jq .[].Id
```

```
"sha256:353d2bd50105cbc3c61540e10cf32a152432d5173bb6318b8e"
```

The second copy is created in Amazon Elastic Container Registry (ECR). To check this copy, perform the following:

**1 Keep the image name and version same as original, run:**

```
$ docker image tag msdp-operator:17.0
```

```
046777922665.dkr.ecr.us-east-2.amazonaws.com/nbuxeksdeploy.aws.io/msdp-operator:17.0
```

**2 Run:**

```
$ docker image ls | grep msdp-operator
```

Sample output:

```
msdp-operator 17.0 353d2bd50105 2 days ago 480 MB
```

```
msdp-operator 17.0 046777922665.dkr.ecr.us-east-2.
```

```
amazonaws.com/nbuxeksdeploy.aws.io/msdp-operator
```

```
17.0 353d2bd50105 2 days ago 480 MB /msdp-operator 17.0
```

```
353d2bd50105 2 days ago 480 MB
```

**3** To push the image to the registry, run:

```
$ docker push testregistry.<account id>.dkr.ecr.<region>.amazonaws.com/<registry>:<tag>.io/msdp-operator
```

The push refers to a repository

```
[046777922665.dkr.ecr.us-east-2.amazonaws.com/hbuxeksdeploy.aws.io/msdp-operator]
```

0a504041c925: Layer already exists

```
17.0: digest:
sha256:d294f260813599562eb5ace9e0acd91d61b7dbc53c3 size:
2622
```

**4** To verify local image digest after the push operation, run:

```
$ docker inspect 353d2bd50105 | jq .[].RepoDigests
```

Sample output:

```
[
 "testregistry.<account id>.dkr.ecr.<region>.amazonaws.com/<registry>:<tag>.io/msdp-operator@sha256: d294f260813599562eb5ace9e0acd91d61b7dbc53c3"
]
```

**5** To verify image presence in the registry, run:

```
$ aws ecr describe-repositories --repository-names
"veritas/main_test1"
```

**Sample output:**

```
"repositories": [
 {
 "repositoryArn": "arn:aws:ecr:us-east-2:046777922665:
repository/veritas/main_test1",
 "registryId": "046777922665",
 "repositoryName": "veritas/main_test1",
 "repositoryUri": "046777922665.dkr.ecr.us-east-2.
amazonaws.com/veritas/main_test1",
 "createdAt": "2022-04-13T07:27:52+00:00",
 "imageTagMutability": "MUTABLE",
 "imageScanningConfiguration": {
 "scanOnPush": false
 },
 "encryptionConfiguration": {
 "encryptionType": "AES256"
 }
 }
]
```

**6** To verify image digest in registry, run:

```
$ aws ecr describe-images --registry-id 046777922665
--repository-name "veritas/main_test1" --image-ids
imageTag=latestTest5
```

**Sample output:**

```
"imageDetails": [
 {
 "registryId": "046777922665",
 "repositoryName": "veritas/main_test1",
 "imageDigest":
"sha256:d0095074286a50c6bca3daeddbaf264cf4006a92fa3a074daa4739cc995b36f8"
 "imageTags": [
 "latestTest5"
],
 "imageSizeInBytes": 38995046,
 "imagePushedAt": "2022-04-13T15:56:07+00:00",
 "imageManifestMediaType": "application/vnd.docker.
distribution.manifest.v2+json",
 "artifactMediaType": "application/vnd.docker.container.image.
 }
]
```

The third copy is located on a Kubernetes node running the container after it is pulled from the registry. To check this copy, do the following:

**1** Run;

```
$ kubectl get nodes -o wide

NAME STATUS VERSION INTERNAL-IP OS-IMAGE
eks-agentpool-7601-vmss000 Ready v1.21.7 10.240.0.4 Ubuntu 18.04.6 LTS
```

**2** Use kubectl debug to run a container on the node:

```
$ kubectl debug node/eks-nodepool11-7601-vmss000-it
--image=mcr.microsoft.com/eks/fundamental/base-ubuntu:v0.0.11
root@eks-agentpool-7601-vmss000:/#
```

**3** You can interact with the node session from the privileged container:

```
chroot /host
```

**4** Verify the presence of the image:

```
/usr/local/bin/crictl image | grep msdp
```

Sample output:

```
<account id>.dkr.ecr.<region>.amazonaws.com/msdp-operator 17.0 353d2bd5
```

**5** Verify the image ID on the Kubernetes node, run:

```
/usr/local/bin/crictl inspecti 353d2bd50105c | jq .[].id
```

Sample output

```
"sha256:353d2bd50105cbc3c61540e10cf32a152432d5173bb6318b8e"
null
```

**6** Verify the image digest on the Kubernetes node, run:

```
/usr/local/bin/crictl inspecti 353d2bd50105c | jq .[].repoDigests
```

Sample output

```
[
 "<account id>.dkr.ecr.<region>.amazonaws.com/msdp-operator@sha256:
d294f260813599562eb5ace9e0acd91d61b7dbc53c3"
]
null
```

## How to make sure that you are running the correct image

Use the steps given above to identify image ID and Digest and compare with values obtained from the registry and the Kubernetes node running the container.

---

**Note:** MSDP Scaleout images (uss-engine, uss-mds, uss-controller, msdp-operator) use IfNotPresent imagePullPolicy. A unique image tag is required in order for a Kubernetes node to pull an updated image.

---

# Getting EEB information from an image, a running container, or persistent data

To view the list of installed EEBs, run the `nbbuilder` script provided in the EEB file archive.

**Getting EEB information from an image, a running container, or persistent data**

```
$ bash nbbuilder.sh -registry_name <account
id>.dkr.ecr.<region>.amazonaws.com -list_installed_eebs
-nb_src_tag=10.1-2 -msdp_src_tag=17.0-2
```

**Sample output:**

```
Wed Feb 2 20:48:13 UTC 2022: Listing strings for EEBs
installed in <account id>.dkr.ecr.<region>.amazonaws.com/netbackup/main:10.1-
EEB_NetBackup_10.1Beta6_PET3980928_SET3992004_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992021_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992022_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992023_EEB1
EEB_NetBackup_10.1Beta6_PET3992020_SET3992019_EEB2
EEB_NetBackup_10.1Beta6_PET3980928_SET3992009_EEB2
EEB_NetBackup_10.1Beta6_PET3980928_SET3992016_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992017_EEB1
Wed Feb 2 20:48:13 UTC 2022: End
Wed Feb 2 20:48:13 UTC 2022: Listing strings for EEBs
installed in <account id>.dkr.ecr.<region>.amazonaws.com/uss-controller:17.0-
EEB_MSDP_17.0_PET3980928_SET3992007_EEB1
EEB_MSDP_17.0_PET3992020_SET3992019_EEB2
EEB_MSDP_17.0_PET3980928_SET3992010_EEB2
Wed Feb 2 20:48:14 UTC 2022: End
Wed Feb 2 20:48:14 UTC 2022: Listing strings for EEBs
installed in <account id>.dkr.ecr.<region>.amazonaws.com/uss-engine:17.0-2.
EEB_MSDP_17.0_PET3980928_SET3992006_EEB1
EEB_MSDP_17.0_PET3980928_SET3992023_EEB1
EEB_MSDP_17.0_PET3992020_SET3992019_EEB2
EEB_MSDP_17.0_PET3980928_SET3992009_EEB2
EEB_MSDP_17.0_PET3980928_SET3992010_EEB2
EEB_MSDP_17.0_PET3980928_SET3992018_EEB1
Wed Feb 2 20:48:14 UTC 2022: End
Wed Feb 2 20:48:14 UTC 2022: Listing strings for EEBs
installed in <account id>.dkr.ecr.<region>.amazonaws.com/uss-mds:17.0-2.
EEB_MSDP_17.0_PET3980928_SET3992008_EEB1
EEB_MSDP_17.0_PET3992020_SET3992019_EEB2
EEB_MSDP_17.0_PET3980928_SET3992010_EEB2
Wed Feb 2 20:48:15 UTC 2022: End
```

Alternatively, if the `nbbuilder` script is not available, you can view the installed EEBs by executing the following command:

```
$ docker run --rm <image_name>:<image_tag> cat
/usr/opencv/pack/pack.summary
```

Sample output:

```
EEB_NetBackup_10.1Beta6_PET3980928_SET3992004_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992021_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992022_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992023_EEB1
EEB_NetBackup_10.1Beta6_PET3992020_SET3992019_EEB2
EEB_NetBackup_10.1Beta6_PET3980928_SET3992009_EEB2
EEB_NetBackup_10.1Beta6_PET3980928_SET3992016_EEB1
EEB_NetBackup_10.1Beta6_PET3980928_SET3992017_EEB1
```

To view all EEBs installed in a running container, run:

```
$ kubectl exec --stdin --tty <primary-pod-name> -n <namespace> --
cat /usr/opencv/pack/pack.summary
```

---

**Note:** The pack directory may be located in different locations in the `uss-*` containers. For example: `/uss-controller/pack` , `/uss-mds/pack`, `/uss-proxy/pack`.

---

To view a list of installed data EEBs from a running container, run:

```
$ kubectl exec --stdin --tty <primary-pod-name> -n <namespace> --
cat /mnt/nbdata/usr/opencv/pack/pack.summary
```

## Resolving the certificate error issue in NetBackup operator pod logs

Following error is displayed in NetBackup operator pod logs when the primary server certificate is changed:

```
ERROR controller-runtime.manager.controller.environment
Error defining desired resource {"reconciler group": "netbackup.veritas.com",
"reconciler kind": "Environment", "name": "test-delete", "namespace": "netbac
>Type": "MSDPScaleout", "Resource": "dedupe1", "error": "Unable to get primar
Get \"https://nbux-10-244-33-24.vxindia.veritas.com:1556/netbackup/config/hos
x509: certificate signed by unknown authority (possibly because of \"crypto/r
verification error\" while trying to verify candidate authority certificate \"
```

To resolve this issue, restart the NetBackup operator by deleting the NetBackup operator pod using the following command:

```
kubectl delete <Netbackup-operator-pod-name> -n <namespace>
```

# Resolving the primary server connection issue

NetBackup operator pod logs displays the following error:

```
Failed to connect the Primary server. "error":
"Get \"https://abc.xyz.com:*/netbackup/security/cacert\":
dial tcp: i/o timeout
```

The above error appears due to the following reasons:

- Operator, primary server and media server pod must have been started in different availability zones.
- Load balancer IP address and node on which primary server pod have started are in different availability zones.
- Load balancer created for NetBackup load balancer service is in failed state.
- FQDN to IP address given in networkLoadBalancer section in CR spec is not DNS resolvable.

**To resolve the primary server connection issue, perform the following**

- 1 Delete the primary server and media server CR instance using the following command:

```
kubectl delete -f <environment.yaml>
```

---

**Caution:** Be cautious while performing this step.

---

- 2 Fix the issue and provide appropriate details in CR specs in `environment.yaml` file.
- 3 Redeploy the NetBackup by reapplying the `environment.yaml` file using the following command:

```
kubectl apply -f environment.yaml
```

## Primary pod is in pending state for a long duration

Primary pod and Operator pod have pod anti-affinity set.

**To resolve the issue of long duration pending state of primary pod**

- 1 Verify if Operator pod and Primary pod are scheduled to same node using the following commands:

```
kubectl get all -n <primary pod namespace> -o wide
```

```
kubectl get all -n <operator pod namespace> -o wide
```

- 2 If it is allocated to same node then create new node with same node selector given in CR for primary server.
- 3 Delete the Primary pod which is in pending state.

The newly created Primary pod must not be in pending state now.

## Host mapping conflict in NetBackup

The following error message is displayed due to host mapping conflict in NetBackup:

```
..exited with status 7659: Connection cannot be established because
the host validation failed on the target host
```

In kubernetes deployment, communication to pod goes through multiple layers that is, load balancer, nodes and pod. In certain setups during communication host may get associated with certain IP and would be changed. That IP may get associated with some different pod and which causes conflict. The host mapping entries is in the form of "::

**To resolve the issue of host mapping conflict in NetBackup**

- 1 To resolve the conflict issue, refer to [Mappings for Approval tab](#) section of the *NetBackup Security and Encryption Guide*.
- 2 To remove the entries that are not valid, refer to [Removing host ID to host name mappings](#) section of the *NetBackup Security and Encryption Guide*.

## NetBackup messaging queue broker take more time to start

This issue is due to **nbmqbroker** service taking more time to start.

**To resolve this issue, perform the following steps**

- 1 Exec into the respective Primary Server pod using the following command:
 

```
kubectl exec -it <pod-name> -n <namespace> -- /bin/bash
```
- 2 Check the nbmqbroker service logs which are in /usr/opensv/mqbroker/logs/ folder.
 

If the value of **checking service start status count**: is more than the 75 then **nbmqbroker** would take more time to start.
- 3 Stop the nbmqbroker service using the following command:
 

```
/usr/opensv/mqbroker/bin/nbmqbroker stop
```
- 4 Open the /usr/opensv/mqbroker/bin/nbmqbroker file.
- 5 Increase the value of **total\_time** and **sleep\_duration** and save the file.
- 6 Start the mqbroker service using the following command:
 

```
/usr/opensv/mqbroker/bin/nbmqbroker start
```

If the Primary Server pod gets restarted then the user must perform the same above steps to increase the values of **total\_time** and **sleep\_duration**, as these values will not get persisted after pod restart.

## Local connection is getting treated as insecure connection

The following error message is displayed when un-necessary audit events are get logged only when reverse dns lookup is enabled for primary and media Load Balancer service:

```
Host 'eaebbef2-57bc-483b-8146-1f6616622276' is trying to connect to
host '<serverName>.abc.com'. The connection is dropped, because the
host '<serverName>.abc.com' now appears to be NetBackup 8.0 or earlier
```

Primary and media servers are referred with multiple IP's inside the pod (pod IP/LoadBalancer IP). With reverse name lookup of IP enabled, NetBackup treats the local connection as remote insecure connection.

To resolve the audit events issue, disable the reverse name lookup of primary and media Load Balancer IP.

## Issue with capacity licensing reporting which takes longer time

The `nbdeployutil` utility does not perform well on EFS or Azure files based volumes. Specify different block storage based volume to get good performance.

**To resolve the issue, perform the following:**

- 1 For running report manually, pass `--parentdir /mnt/nbdb/<FOLDER_NAME>` to `nbdeployutil` command.
- 2 For changing `parentdir` to scheduled capacity reporting, provide a custom value in `nbdeployutilconfig.txt` file.
- 3 Create/Edit the `nbdeployutilconfig.txt` file located at `/usr/opensv/var/global/` by adding the following entry:

```
[NBDEPLOYUTIL_INCREMENTAL]
PARENTDIR=/mnt/nbdb/<FOLDER_NAME>
```

## Backing up data from Primary server's `/mnt/nbdata/` directory fails with primary server as a client

Backing up data from Primary server's `/mnt/nbdata/` directory fails with primary server as a client.

From NetBackup version 10.1 onwards, the `/mnt/nbdata` directory would be on EFS. The `/mnt/nbdata` directory of primary server is mounted using the NFS protocol.

Hence for backing up the data from `/mnt/nbdata` directory, change the policy attribute for such policies.

To resolve this issue, enable the **Follow NFS** check box in policy attribute.

## Wrong EFS ID is provided in `environment.yaml` file

- Following error message is displayed when wrong EFS ID is provided in `environment.yaml` file:

```
"samples/environment.yaml": admission webhook
"environment2-validating-webhook.netbackup.veritas.com" denied the
request: Environment change rejected by validating webhook: EFS ID
provided for Catalog storage is not matching with EFS ID of already
created persistent volume for Primary servers Catalog volume. Old
EFS ID fs-0bf084568203f1c27 : New EFS ID fs-0bf084568203f1c29
```

Above error can appear due to the following reasons:

- During upgrade, if EFS ID provided is different from the EFS ID used in the previous version deployment.
- During fresh deployment, if user manually creates PV and PVC with EFS ID and provides different EFS ID in `environment.yaml` file.

**To resolve this issue, perform the following:**

- 1 Identify the correct EFS ID used for PV and PVC.
  - Previously used EFS ID can be retrieved from PV and PVC by using the following steps:
 

```
kubectl get pvc -n <namespace>
```
  - From the output, copy the name of catalog PVC which is of the following format:
 

```
catalog-<resource name prefix>-primary-0
```
  - Describe catalog PVC using the following command:
 

```
kubectl describe pvc <pvc name> -n <namespace>
```

 Note down the value of **Volume** field from the output.
  - Describe PV using the following command:
 

```
kubectl describe pv <value of Volume obtained from above step>
```

 Note down the value of **VolumeHandle** field from the output. This is the EFS ID used previously.
- 2 Provide correct EFS ID in the `environment.yaml` file and apply the `environment.yaml` using the following command:
 

```
kubectl apply -f environment.yaml
```

## Primary pod is in ContainerCreating state

Primary pod is in ContainerCreating state for a long period due to the following reasons:

- Wrong EFS ID
- Format of the EFS ID is incorrect

**To resolve this issue, perform the following:**

- 1 Describe primary pod using the following command:

```
kubectl describe <name of primary server pod> -n <namespace>
```

- 2 Depending on the following appropriate scenario, fix the error from the output under the Event section:

- If the event log has an error related to incorrect EFS ID or incorrect format, then update the `environment.yaml` file with the correct EFS ID and perform the below steps.

Or

- If the event log has an error other than the error related to incorrect EFS ID, then analyze and fix the error and perform the below steps.

- 3 After fixing the error, clean the environment using the following command:

```
kubectl delete -k operator/
```

- 4 Delete PV and PVC created for primary server only by using the following command:

```
Kubectl get pvc -n <namespace>
```

Describe the PVC for primary server which has the following format and obtain the corresponding PV name:

```
catalog-<resource name prefix of primary>-primary-0
data-<resource name prefix of primary>-primary-0
logs-<resource name prefix of primary>-primary-0
```

Delete PVC and PV names using the following commands: For PVC: `kubectl delete pvc <pvc name> -n <namespace>` For PV: `kubectl delete pv <pv name>`

- PVC: `kubectl delete pvc <pvc name> -n <namespace>`
- PV: `kubectl delete pv <pv name>`

- 5 Deploy NetBackup operator again and then apply the `environment.yaml` file.

## Webhook displays an error for PV not found

Following error message is displayed when the user creates PVC with `claimRef` and fails to create PV during deployment:

```
error when creating "samples/environment.yaml": admission webhook
"environment2-validating-webhook.netbackup.veritas.com" denied the
request: Environment change rejected by validating webhook: PVC exist
```

for Catalog Volume. PersistentVolume pv-catalog-nbuxsystest-primary-0 does not exist for primary server's Catalog volume : PersistentVolume "pv-catalog-nbuxsystest-primary-0" not found.

This issue can be resolved by creating PV and apply `environment.yaml` file again.

# CR template

This appendix includes the following topics:

- [Secret](#)
- [MSDP Scaleout CR](#)

## Secret

The Secret is the Kubernetes security component that stores the MSDP credentials that are required by the CR YAML.

```
The secret is used to store the MSDP credential, which is required
by the CR YAML as follows.
This part should be created separately and not be part of CR Template.
The secret should have a "username" and a "password" key-pairs with the
corresponding username and password values.
Please follow MSDP guide for the rules of the credential.
https://www.veritas.com/content/support/en_US/article.100048511
The pattern is "^[\w!$+\\-,.:;=?@[\\`{}|\\~]{1,62}$"
We can create the secret directly via kubectl command:
kubectl create secret generic sample-secret --namespace
sample-namespace \
--from-literal=username=<username> --from-literal=password=<password>
Alternatively, we can create the secret with a YAML file in the
following format.
apiVersion: v1
kind: Secret
metadata:
 name: sample-secret
 # The namespace needs to be present.
 namespace: sample-namespace
```

```
stringData:
 # Please follow MSDP guide for the credential characters and length.
 # https://www.veritas.com/content/support/en_US/article.100048511
 # The pattern is "^[\w!$+\-.,:;=?@[\\]\`{}|\~]{1,62}$"
 username: xxxx
 password: xxxxxx
```

## MSDP Scaleout CR

- The CR name must be less than 40 characters.
- The MSDP credentials stored in the Secret must match MSDP credential rules. See [Deduplication Engine credentials for NetBackup](#)
- MSDP CR cannot be deployed in the namespace of MSDP operator. It must be in a separate namespace.
- You cannot reorder the IP/FQDN list. You can update the list by appending the information.
- You cannot change the storage class name. The storage class must be backed with Amazon EBS CSI driver "ebs.csi.aws.com".
- You cannot change the data volume list other than for storage expansion. It is append-only and storage expansion only. Up to 16 data volumes are supported.
- Like the data volumes, the catalog volume can be changed for storage expansion only.
- You cannot change or expand the size of the log volume by changing the MSDP CR.
- You cannot enable NBCA after the configuration.
- Once KMS and the OST registration parameters set, you cannot change them.
- You cannot change the core pattern.

MSDP Scaleout CR template:

```
The MSDPScaleout CR YAML
notes:
The CR name should be <= 40 characters.
The MSDP credential stored in the Secret should match MSDP credential
rules defined in https://www.veritas.com/content/support/en_US/article.
100048511
apiVersion: msdp.veritas.com/v1
kind: MSDPScaleout
metadata:
```

```
The CR name should not be longer than 40 characters.
name: sample-app
The namespace needs to be present for the CR to be created in.
It is not allowed to deploy the CR in the same namespace with MSDP
operator.
 namespace: sample-namespace
spec:
 # Your Container Registry(ECR for AWS EKS) URL where
the docker images can be pulled from the k8s cluster on demand
 # The allowed length is in range 1-255
 # It is optional for BYO. The code does not check the presence or
validation.
 # User needs to specify it correctly if it is needed.
 containerRegistry: sample.url
 #
 # The MSDP version string. It is the tag of the MSDP docker images.
 # The allowed length is in range 1-64
 version: "sample-version-string"
 #
 # Size defines the number of Engine instances in the MSDP-X cluster.
 # The allowed size is between 1-16
 size: 4
 #
 # The IP and FQDN pairs are used by the Engine Pods to expose the
MSDP services.
 # The IP and FQDN in one pair should match each other correctly.
 # They must be pre-allocated.
 # The item number should match the number of Engine instances.
 # They are not allowed to be changed or re-ordered. New items can be
appended for scaling out.
 # The first FQDN is used to configure the storage server in NetBackup,
automatically if autoRegisterOST is enabled,
 # or manually by the user if not.
 serviceIPFQDNs:
 # The pattern is IPv4 or IPv6 format
 - ipAddr: "sample-ip1"
 # The pattern is FQDN format.
 fqdn: "sample-fqdn1"
 - ipAddr: "sample-ip2"
 fqdn: "sample-fqdn2"
 - ipAddr: "sample-ip3"
 fqdn: "sample-fqdn3"
 - ipAddr: "sample-ip4"
```

```
fqdn: "sample-fqdn4"
#
Optional annotations to be added in the LoadBalancer services for the
Engine IPs.
In case we run the Engines on private IPs, we need to add some
customized annotations to the LoadBalancer services.
loadBalancerAnnotations:
If it's an EKS environment, specify the following annotation
to use the internal IPs.
see https://docs.microsoft.com/en-us/amazon/aws/internal-lb
service.beta.kubernetes.io/aws-load-balancer: "true"
If the internal IPs are in a different subnet as the EKS cluster,
the following annotation should be
specified as well. The subnet specified must be in the same virtual
network as the EKS cluster.
service.beta.kubernetes.io/aws-load-balancer-internal-subnet:
"apps-subnet"
#
If your cluster is EKS, the following annotation item is required.
The subnet specified must be in the same VPC as your EKS.
service.beta.kubernetes.io/aws-load-balancer-subnets: "subnet-04c47
28ec4d0ecb90"
#
SecretName is the name of the secret which stores the MSDP credential.
AutoDelete, when true, will automatically delete the secret specified
by SecretName after the
initial configuration. If unspecified, AutoDelete defaults to true.
When true, SkipPrecheck will skip webhook validation of the MSDP
credential. It is only used in data re-use
scenario (delete CR and re-apply with pre-existing data) as the secret
will not take effect in this scenario. It
cannot be used in other scenarios. If unspecified, SkipPrecheck defaults
to false.
credential:
The secret should be pre-created in the same namespace which has the
MSDP credential stored.
The secret should have a "username" and a "password" key-pairs with
the corresponding username and password values.
Please follow MSDP guide for the rules of the credential.
https://www.veritas.com/content/support/en_US/article.100048511
A secret can be created directly via kubectl command or with the
equivalent YAML file:
kubectl create secret generic sample-secret --namespace sample-
```

```
namespace \
 # --from-literal=username=<username> --from-literal=password=
<password>
 secretName: sample-secret
 # Optional
 # Default is true
 autoDelete: true
 # Optional
 # Default is false.
 # Should be specified only in data re-use scenario (aka delete and
re-apply CR with pre-existing data)
 skipPrecheck: false
 #
 # Paused is used for maintenance only. In most cases you do not need
to specify it.
 #
 # When it is specified, MSDP operator stops reconciling the corresponding
MSDP-X cluster (aka the CR).
 # Optional.
 # Default is false
 # paused: false
 #
 # The storage classes for logVolume, catalogVolume and dataVolumes
should be:
 # - Backed with AWS disk CSI driver "disk.csi.aws.com" with the
managed disks, and allow volume
 # expansion.
 # - The AWS in-tree storage driver "kubernetes.io/aws-disk" is not
supported. You need to explicitly
 # enable the AWS disk CSI driver when configuring your EKS cluster,
or use k8s version v1.21.x which
 # has the AWS disk CSI driver built-in.
 # - In LRS category.
 # - At least Standard SSD for dev/test, and Premium SSD or Ultra Disk
for production.
 # - The same storage class can be used for all the volumes.
 #
 # LogVolume is the volume specification which is used to provision a
volume of an MDS or Controller
 # Pod to store the log files and core dump files.
 # It is not allowed to be changed.
 # In most cases, 5-10 GiB capacity should be big enough for one MDS or
Controller Pod to use.
```

```
logVolume:
 storageClassName: sample-AWS-disk-sc1
 resources:
 requests:
 storage: xGi
#
CatalogVolume is the volume specification which is used to provision a
volume of an MDS or Engine
Pod to store the catalog and metadata. It is not allowed to be changed
unless for capacity expansion.
Expanding the existing catalog volumes expects short downtime of the
Engines.
Please note the MDS Pods do not respect the storage request in
CatalogVolume, instead they provision the
volumes with the minimal capacity request of 500MiB.
catalogVolume:
 storageClassName: sample-AWS-disk-sc2
 resources:
 requests:
 storage: xxxGi
#
DataVolumes is a list of volume specifications which are used to
provision the volumes of
an Engine Pod to store the MSDP data.
The items are not allowed to be changed or re-ordered unless for
capacity expansion.
New items can be appended for adding more data volumes to each
Engine Pod.
Appending new data volumes or expanding the existing data volumes
expects short downtime of the Engines.
The allowed item number is in range 1-16. To allow the other MSDP-X
Pods (e.g. Controller, MDS) running
on the same node, the item number should be no more than "<the
maximum allowed volumes on the node> - 5".
The additional 5 data disks are for the potential one MDS Pod, one
Controller Pod or one MSDP operator Pod
to run on the same node with one MSDP Engine.
dataVolumes:
- storageClassName: sample-aws-disk-sc3
 resources:
 requests:
 storage: xxTi
- storageClassName: sample-aws-disk-sc3
```

```
resources:
 requests:
 storage: xxTi
#
NodeSelector is used to schedule the MSDPScaleout Pods on the
specified nodes.
Optional.
Default is empty (aka all available nodes)
nodeSelector:
 # e.g.
 # agentpool: nodegroup2
 sample-node-label1: sampel-label-value1
 sample-node-label2: sampel-label-value2
#
NBCA is the specification for the MSDP-X cluster to enable NBCA
SecComm for the Engines.
Optional.
nbca:
 # The master server name
 # The allowed length is in range 1-255
 masterServer: sample-master-server-name
 # The CA SHA256 fingerprint
 # The allowed length is 95
 cafp: sample-ca-fp
 # The NBCA authentication/reissue token
 # The allowed length is 16
 # For security consideration, a token with maximum 1 user allowed
and valid for 1 day should be sufficient.
 token: sample-auth-token
#
KMS includes the parameters to enable KMS for the Engines.
We support to enable KMS in init or post configuration.
We do not support to change the parameters once they have been set.
Optional.
kms:
 # As either the NetBackup KMS or external KMS (EKMS) is configured
or registered on NetBackup master server, then used by
 # MSDP by calling the NetBackup API, kmsServer is the NetBackup master
server name.
 kmsServer: sample-master-server-name
 keyGroup: sample-key-group-name
#
autoRegisterOST includes the parameter to enable or disable the
```

```
automatic registration of
the storage server, the default disk pool and storage unit when
MSDP-X configuration finishes.
We do not support to change autoRegisterOST.
autoRegisterOST:
If it is true, and NBCA is enabled, the operator would register
the storage server,
disk pool and storage unit on the NetBackup primary server, when
the MSDP CR is deployed.
The first Engine FQDN is the storage server name.
The default disk pool is in format "default_dp_<firstEngineFQDN>".
The default storage unit is in format "default_stu_<firstEngineFQDN>".
The default maximum number of concurrent jobs for the STU is 240.
In the CR status, field "ostAutoRegisterStatus.registered" with
value True, False or Unknown indicates the registration state.
It is false by default.
enabled: true
#
CorePattern is the core pattern of the nodes where the MSDPScaleout
Pods are running.
It is path-based. A default core path "/core/core.%e.%p.%t" will be
used if not specified.
In most cases, you do not need to specify it.
It is not allowed to be changed.
Optional.
corePattern: /sample/core/pattern/path
#
tcpKeepAliveTime sets the namespaced sysctl parameter net.ipv4.tcp_
keepalive_time in Engine Pods.
It is in seconds.
The minimal allowed value is 60 and the maximum allowed value is 1800.
A default value 120 is used if not specified. Set it to 0 to disable
the option.
It is not allowed to change unless in maintenance mode (paused=true),
and the change will not apply until the Engine Pods get restarted.
For EKS deployment in 10.1 release, please leave it unspecified or
specify it with a value smaller than 240.
tcpKeepAliveTime: 120
#
TCPIIdleTimeout is used to change the default value for AWS Load
Balancer rules and Inbound NAT rules.
It is in minutes.
The minimal allowed value is 4 and the maximum allowed value is 30.
```

```
A default value 30 minutes is used if not specified. Set it to 0 to
disable the option.
It is not allowed to change unless in maintenance mode (paused=true),
and the change will not apply until the Engine Pods and the LoadBalancer
services get recreated.
For EKS deployment in 10.1 release, please leave it unspecified or
specify it with a value larger than 4.
tcpIdleTimeout: 30
```