Red Hat OpenShift Data Foundation 4.11

Deploying and managing OpenShift Data Foundation on single node OpenShift clusters

Instructions for deploying and managing OpenShift Data Foundation on single node OpenShift clusters.

Last Updated: 2023-02-01

Red Hat OpenShift Data Foundation 4.11 Deploying and managing OpenShift Data Foundation on single node OpenShift clusters

Instructions for deploying and managing OpenShift Data Foundation on single node OpenShift clusters.

Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at
http://creativecommons.org/licenses/by-sa/3.0/
. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Read this document for instructions regarding installing Red Hat OpenShift Data Foundation Logical Volume Manager Operator on single node OpenShift clusters. Deploying and managing OpenShift Data Foundation 4.11 on single node OpenShift clusters is a Technology Preview feature. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. A General Available version of this product (renamed as logical volume manager storage (LVM storage)) is now available. For information, see LVM Storage documentation.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message .

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Do let us know how we can make it better. To give feedback:

- For simple comments on specific passages:
  1. Make sure you are viewing the documentation in the *Multi-page HTML* format. In addition, ensure you see the **Feedback** button in the upper right corner of the document.
  2. Use your mouse cursor to highlight the part of text that you want to comment on.
  3. Click the **Add Feedback** pop-up that appears below the highlighted text.
  4. Follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
  1. Go to the Bugzilla website.
  2. In the **Component** section, choose **documentation**.
  3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
  4. Click **Submit Bug**.

PREFACE

Red Hat OpenShift Data Foundation supports deploying OpenShift Data Foundation using the Red Hat OpenShift Data Foundation Logical Volume Manager Operator on single node OpenShift (SNO) clusters. This operator uses the TopoLVM CSI driver to dynamically provision local storage.

Red Hat OpenShift Data Foundation Logical Volume Manager Operator creates thin-provisioned volumes using the Logical Volume Manager and provides dynamic provisioning of block storage on a single node, limited resources SNO cluster.

You can deploy the Red Hat OpenShift Data Foundation Logical Volume Manager Operator on a single node Openshift bare metal or user provisioned infrasturcture cluster and configure it to dynamically provision storage for your workloads.

The operator creates a volume group using all the available unused disks and creates a single thin pool with a size of 90% of the volume group. The remaining 10% of the volume group is left free to enable data recovery by expanding the thin pool when required. You might need to manually perform such recovery.

You can use persistent volume claims (PVCs) and volume snapshots provisioned by the Logical Volume Manager Operator to request storage and create volume snapshots.

The Red Hat OpenShift Data Foundation Logical Volume Manager Operator configures a default overprovisioning limit of 10 to take advantage of the thin-provisioning feature. The total size of the volumes and volume snapshots that can be created on the single node OpenShift clusters is 10 times the size of the thin pool.

You can deploy OpenShift Data Foundation on single node OpenShift clusters using one of the following:

- Red Hat Advanced Cluster Management for Kubernetes (RHACM)
- OpenShift Web Console

CHAPTER 1. DEPLOYING OPENSHIFT DATA FOUNDATION ON SINGLE NODE OPENSHIFT CLUSTERS USING RHACM

1.1. REQUIREMENTS FOR DEPLOYING USING RHACM

Before you begin deploying OpenShift Data Foundation Logical Volume Manager Operator on single node Openshift clusters, ensure that the following requirements are met:

1. You have installed Red Hat Advanced Cluster Management for Kubernetes (RHACM) on an OpenShift cluster. For information, see Red Hat Advanced Cluster Management for Kubernetes: Install.

2. Every managed SNO cluster has dedicated disks that are used to provision storage.

1.2. INSTALLING THE OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER OPERATOR USING RHACM

The OpenShift Data Foundation Logical Volume Manager Operator is deployed on single node OpenShift (SNO) clusters using Red Hat Advanced Cluster Management for Kubernetes (RHACM). You create a Policy on RHACM that deploys and configures the operator when it is applied to managed clusters which match the selector specified in the **PlacementRule**. The policy is also applied to clusters that are imported later and satisfy the **PlacementRule**.

Prerequisites

- Access to the RHACM cluster using an account with **cluster-admin** and operator installation permissions.

- Dedicated disks on each SNO cluster to be used by OpenShift Data Foundation Logical Volume Manager Operator.

NOTE

Ensure that no other storage provisioner is installed on the SNO cluster. OpenShift Data Foundation Logical Volume Manager Operator should be the only storage provisioner as it uses all available disks on the node.

Procedure

1. Log in to the RHACM CLI using your OpenShift credentials.
For more information, see Install Red Hat Advanced Cluster Management for Kubernetes .

2. Create a namespace in which you will create policies.

```
# oc create ns lvm-policy-ns
```

3. Save the following YAML to a file with a name such as **policy-lvm-operator.yaml** to create a policy.

```
# This policy verifies the installation of the official version of the {product-name-short}
# Logical Volume Manager Operator on the managed clusters.
# If set to "enforce" it installs the operator.
# Used APIs: OLM, ODF-LVMO #https://github.com/operator-framework/operator-lifecycle-manager
# https://github.com/red-hat-storage/lvm-operator

apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-install-odf-lvm-operator
spec:
  clusterConditions:
  - status: "True"
    type: ManagedClusterConditionAvailable
  clusterSelector:
    matchExpressions:
    - key: vendor
      operator: In
      values:
      - OpenShift
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-install-odf-lvm-operator
placementRef:
  apiGroup: apps.open-cluster-management.io
  kind: PlacementRule
  name: placement-install-odf-lvm-operator
subjects:
- apiGroup: policy.open-cluster-management.io
  kind: Policy
  name: install-odf-lvm-operator
---
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  annotations:
    policy.open-cluster-management.io/categories: CM Configuration Management
    policy.open-cluster-management.io/controls: CM-2 Baseline Configuration
    policy.open-cluster-management.io/standards: NIST SP 800-53
  name: install-odf-lvm-operator
spec:
  disabled: false
  remediationAction: enforce
  policy-templates:
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name: install-odf-lvm-operator
      spec:
        object-templates:
        - complianceType: musthave
          objectDefinition:
            apiVersion: v1
            kind: Namespace
            metadata:
```

```
      labels:
       openshift.io/cluster-monitoring: "true"
       pod-security.kubernetes.io/enforce: privileged
       pod-security.kubernetes.io/audit: privileged
       pod-security.kubernetes.io/warn: privileged
      name: openshift-storage
   - complianceType: musthave
    objectDefinition:
     apiVersion: operators.coreos.com/v1
     kind: OperatorGroup
     metadata:
      name: openshift-storage-operatorgroup
      namespace: openshift-storage
     spec:
      targetNamespaces:
      - openshift-storage
   - complianceType: musthave
    objectDefinition:
     apiVersion: operators.coreos.com/v1alpha1
     kind: Subscription
     metadata:
      name: odf-lvm-operator
      namespace: openshift-storage
     spec:
      installPlanApproval: Automatic
      name: odf-lvm-operator
      source: redhat-operators
      sourceNamespace: openshift-marketplace
   remediationAction: enforce
   severity: low
- objectDefinition:
  apiVersion: policy.open-cluster-management.io/v1
  kind: ConfigurationPolicy
  metadata:
   name: odf-lvmcluster
  spec:
   object-templates:
     - complianceType: musthave
      objectDefinition:
       apiVersion: lvm.topolvm.io/v1alpha1
       kind: LVMCluster
       metadata:
        name: odf-lvmcluster
        namespace: openshift-storage
       spec:
        storage:
         deviceClasses:
         - name: vg1
          thinPoolConfig:
           name: thin-pool-1
           sizePercent: 90
           overprovisionRatio: 10
   remediationAction: enforce
   severity: low
```

4. Create the policy in the namespace by running the following command:

```
# oc create -f policy-lvm-operator.yaml -n lvm-policy-ns
```

where, ***policy-lvm-operator.yaml*** is the name of the file to which the policy is saved.

This creates a **Policy**, a **PlacementRule**, and a **PlacementBinding** in the namespace, **lvm-policy-ns**. The **Policy** creates a **Namespace**, **OperatorGroup**, **Subscription**, and **LVMCluster** resource on the clusters matching the PlacementRule. This deploys the operator on the SNO clusters which match the selection criteria and configures it to set up the required resources to provision storage. The operator uses all the unused disks after installation.

1.3. UNINSTALLING OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER OPERATOR INSTALLED USING RHACM

To uninstall OpenShift Data Foundation Logical Volume Manager Operator when you have installed the operator using RHACM, you need to delete the ACM policy that you created for deploying and configuring the operator. However, when you delete the ACM policy, the resources that the policy has created are not removed. You need to create additional policies to remove the resources.

As the resources that are created are not removed when you delete the policy, you need to perform the following steps:

- Remove all the PVCs and volume snapshots provisioned by the Logical Volume Manager Operator.
- Remove the **LVMCluster** resources to clean up the Logical Volume Manager resources created on the disks.
- Create an additional policy to uninstall the operator.

Prerequisites

- Ensure that the following are deleted before deleting the policy:
  - All the applications on the managed clusters that are using the storage provisioned by the OpenShift Data Foundation Logical Volume Manager Operator.
  - Persistent volume claims (PVCs) and persistent volumes (PVs) provisioned using the OpenShift Data Foundation Logical Volume Manager Operator.
  - All volume snapshots provisioned by the OpenShift Data Foundation Logical Volume Manager Operator.
- Ensure that no logical volume resources exist by using the **oc get logicalvolume** command.
- Access to the RHACM cluster using an account with **cluster-admin** role.

Procedure

1. In the OpenShift command-line interface, delete the ACM policy that you created for deploying and configuring the OpenShift Data Foundation Logical Volume Manager Operator on the hub cluster by using the following command:

```
# oc delete -f policy-lvm-operator.yaml -n lvm-policy-ns
```

2. Save the following YAML to a file with a name such as **odf-lvmcluster-deletion.yaml** to create a policy for removing the **LVMCluster**. This enables the operator to clean up all the Logical Volume Manager resources that it created on the cluster.

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy-lvmcluster-delete
  annotations:
    policy.open-cluster-management.io/standards: NIST SP 800-53
    policy.open-cluster-management.io/categories: CM Configuration Management
    policy.open-cluster-management.io/controls: CM-2 Baseline Configuration
spec:
  remediationAction: enforce
  disabled: false
  policy-templates:
    - objectDefinition:
        apiVersion: policy.open-cluster-management.io/v1
        kind: ConfigurationPolicy
        metadata:
          name: policy-lvmcluster-removal
        spec:
          remediationAction: enforce # the policy-template spec.remediationAction is overridden
by the preceding parameter value for spec.remediationAction.
          severity: low
          object-templates:
            - complianceType: mustnothave
              objectDefinition:
                kind: LVMCluster
                apiVersion: lvm.topolvm.io/v1alpha1
                metadata:
                  name: odf-lvmcluster
                  namespace: openshift-storage # must have namespace 'openshift-storage'
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
 name: binding-policy-lvmcluster-delete
placementRef:
 apiGroup: apps.open-cluster-management.io
 kind: PlacementRule
 name: placement-policy-lvmcluster-delete
subjects:
 - apiGroup: policy.open-cluster-management.io
   kind: Policy
   name: policy-lvmcluster-delete
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
 name: placement-policy-lvmcluster-delete
spec:
 clusterConditions:
   - status: 'True'
     type: ManagedClusterConditionAvailable
 clusterSelector:
    matchExpressions:
      - key: vendor
        operator: In
        values:
          - OpenShift
```

3. Create the policy by running the following command:

```
# oc create -f odf-lvmcluster-deletion.yaml -n lvm-policy-ns
```

4. Save the following YAML to a file with a name such as **check-odf-lvmcluster-deletion.yaml** to create a policy to check if the **LVMCluster** CR has been removed.

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy-lvmcluster-inform
  annotations:
    policy.open-cluster-management.io/standards: NIST SP 800-53
    policy.open-cluster-management.io/categories: CM Configuration Management
    policy.open-cluster-management.io/controls: CM-2 Baseline Configuration
spec:
  remediationAction: inform
  disabled: false
  policy-templates:
    - objectDefinition:
        apiVersion: policy.open-cluster-management.io/v1
        kind: ConfigurationPolicy
        metadata:
          name: policy-lvmcluster-removal-inform
        spec:
          remediationAction: inform # the policy-template spec.remediationAction is overridden by the preceding parameter value for spec.remediationAction.
          severity: low
          object-templates:
            - complianceType: mustnothave
              objectDefinition:
                kind: LVMCluster
                apiVersion: lvm.topolvm.io/v1alpha1
                metadata:
                  name: odf-lvmcluster
                  namespace: openshift-storage # must have namespace 'openshift-storage'
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
 name: binding-policy-lvmcluster-check
placementRef:
 apiGroup: apps.open-cluster-management.io
 kind: PlacementRule
 name: placement-policy-lvmcluster-check
subjects:
 - apiGroup: policy.open-cluster-management.io
   kind: Policy
   name: policy-lvmcluster-inform
```

```
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
 name: placement-policy-lvmcluster-check
spec:
 clusterConditions:
   - status: 'True'
     type: ManagedClusterConditionAvailable
 clusterSelector:
   matchExpressions:
     - key: vendor
       operator: In
       values:
         - OpenShift
```

5. Create the policy by running the following command:

```
# oc create -f check-odf-lvmcluster-deletion.yaml -n lvm-policy-ns
```

6. Check the policy status.

```
# oc get policy -n lvm-policy-ns
NAME                      REMEDIATION ACTION   COMPLIANCE STATE   AGE
policy-lvmcluster-delete   enforce              Compliant          15m
policy-lvmcluster-inform   inform               Compliant          15m
```

7. After both the policies are compliant, save the following YAML to a file with a name such as **odf-lvm-operator-remove-policy.yaml** to create a policy to uninstall the OpenShift Data Foundation Logical Volume Manager Operator.

```
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-uninstall-odf-lvm-operator
spec:
  clusterConditions:
  - status: "True"
    type: ManagedClusterConditionAvailable
  clusterSelector:
    matchExpressions:
    - key: vendor
      operator: In
      values:
      - OpenShift
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-uininstall-odf-lvm-operator
placementRef:
  apiGroup: apps.open-cluster-management.io
  kind: PlacementRule
  name: placement-uninstall-odf-lvm-operator
subjects:
```

```
- apiGroup: policy.open-cluster-management.io
  kind: Policy
  name: uninstall-odf-lvm-operator
---
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  annotations:
    policy.open-cluster-management.io/categories: CM Configuration Management
    policy.open-cluster-management.io/controls: CM-2 Baseline Configuration
    policy.open-cluster-management.io/standards: NIST SP 800-53
  name: uninstall-odf-lvm-operator
spec:
  disabled: false
  policy-templates:
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name: uninstall-odf-lvm-operator
      spec:
        object-templates:
        - complianceType: mustnothave
          objectDefinition:
            apiVersion: v1
            kind: Namespace
            metadata:
              name: openshift-storage
        - complianceType: mustnothave
          objectDefinition:
            apiVersion: operators.coreos.com/v1
            kind: OperatorGroup
            metadata:
              name: openshift-storage-operatorgroup
              namespace: openshift-storage
            spec:
              targetNamespaces:
              - openshift-storage
        - complianceType: mustnothave
          objectDefinition:
            apiVersion: operators.coreos.com/v1alpha1
            kind: Subscription
            metadata:
              name: odf-lvm-operator
              namespace: openshift-storage
            spec:
              installPlanApproval: Automatic
              name: odf-lvm-operator
              source: redhat-operators
              sourceNamespace: openshift-marketplace
        remediationAction: enforce
        severity: low
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
```

```
  name: policy-remove-lvm-operator-crds
spec:
  object-templates:
  - complianceType: mustnothave
    objectDefinition:
      apiVersion: apiextensions.k8s.io/v1
      kind: CustomResourceDefinition
      metadata:
        name: logicalvolumes.topolvm.cybozu.com
  - complianceType: mustnothave
    objectDefinition:
      apiVersion: apiextensions.k8s.io/v1
      kind: CustomResourceDefinition
      metadata:
        name: lvmclusters.lvm.topolvm.io
  - complianceType: mustnothave
    objectDefinition:
      apiVersion: apiextensions.k8s.io/v1
      kind: CustomResourceDefinition
      metadata:
        name: lvmvolumegroupnodestatuses.lvm.topolvm.io
  - complianceType: mustnothave
    objectDefinition:
      apiVersion: apiextensions.k8s.io/v1
      kind: CustomResourceDefinition
      metadata:
        name: lvmvolumegroups.lvm.topolvm.io
  remediationAction: enforce
  severity: high
```

8. Create the policy by running the following command:

```
# oc create -f odf-lvm-operator-remove-policy.yaml -ns lvm-policy-ns
```

CHAPTER 2. DEPLOYING OPENSHIFT DATA FOUNDATION ON SINGLE NODE OPENSHIFT CLUSTERS USING OPENSHIFT WEB CONSOLE

2.1. INSTALLING RED HAT OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER OPERATOR USING OPENSHIFT WEB CONSOLE

You can install Red Hat OpenShift Data Foundation Logical Volume Manager Operator using the Red Hat OpenShift Container Platform Operator Hub.

Prerequisites

- Access to an OpenShift Container Platform cluster using an account with **cluster-admin** and Operator installation permissions.

> **NOTE**
>
> Ensure that no other storage provisioner is installed on the SNO cluster. OpenShift Data Foundation Logical Volume Manager Operator should be the only storage provisioner as it uses all available disks on the node.

Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators → OperatorHub**.
3. Scroll or type **ODF LVM Operator** into the **Filter by keyword** box to find the **ODF LVM Operator**.
4. Click **Install**.
5. Set the following options on the **Install Operator** page:
    a. Update Channel as **stable-4.11**.

    b. Installation Mode as **A specific namespace on the cluster**.

    c. Installed Namespace as **Operator recommended namespace openshift-storage**. If Namespace **openshift-storage** does not exist, it is created during the operator installation.

    d. Select Approval Strategy as **Automatic** or **Manual**.
    If you select **Automatic** updates, then the Operator Lifecycle Manager (OLM) automatically upgrades the running instance of your Operator without any intervention.

    If you select **Manual** updates, then the OLM creates an update request. As a cluster administrator, you must then manually approve that update request to update the Operator to a newer version.

    e. Click **Install**.

Verification steps

- Verify that the **ODF LVM Operator** shows a green tick indicating successful installation.

2.2. CREATING OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER CLUSTER

Create Logical Volume Manager cluster after you install the OpenShift Data Foundation Logical Volume Manager operator.

Prerequisites

- The OpenShift Data Foundation Logical Volume Manager operator must be installed from the Operator Hub.

Procedure

1. In the OpenShift Web Console, click **Operators → Installed Operators** to view all the installed operators.
   Ensure that the **Project** selected is **openshift-storage**.
2. Click on the **ODF LVM operator**, and then click **Create instance** under LVMCluster.
3. In the Create LVMCluster page, select either **Form view** or **YAML view**.
4. Enter a name for the cluster.
5. Click **Create**.

Verification Steps

1. Click **Storage → Storage Classes** from the left pane of the OpenShift Web Console.
2. Verify that the **odf-lvm-<device-class-name>** storage class is created with the ODF LVM cluster creation. By default, **vg1** is the device-class-name.

2.3. UNINSTALLING OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER OPERATOR INSTALLED USING OPENSHIFT WEB CONSOLE

Prerequisites

- Ensure that the following are deleted before deleting the policy:
  - All the applications on the managed clusters that are using the storage provisioned by the OpenShift Data Foundation Logical Volume Manager Operator.
  - Persistent volume claims (PVCs) and persistent volumes (PVs) provisioned using the OpenShift Data Foundation Logical Volume Manager Operator.
  - All volume snapshots provisioned by the OpenShift Data Foundation Logical Volume Manager Operator.
- Ensure that no logical volume resources exist by using the **oc get logicalvolume** command.

Procedure

1. Select the project **openshift-storage**.

```
$ oc project openshift-storage
```

2. Display the pods.

```
$ oc get pods
NAME                                         READY   STATUS    RESTARTS      AGE
lvm-operator-controller-manager-54df65b5c4-n7bhb   3/3     Running   1 (45h ago)   7d2h
topolvm-controller-645cb47cd4-kskqb               5/5     Running   5 (45h ago)   7d3h
topolvm-node-7bqxp                                4/4     Running   0             7d3h
vg-manager-hwmg2                                  1/1     Running   0             7d
```

3. Display the **lvmcluster**.

```
$ oc get lvmcluster
NAME             AGE
odf-lvmcluster   7d3h
```

4. Delete the **lvmcluster**.

```
$ oc delete lvmcluster odf-lvmcluster
lvmcluster.lvm.topolvm.io "odf-lvmcluster" deleted
```

5. Verify the deletion by displaying the **lvmcluster**.

```
$ oc get lvmcluster
No resources found in openshift-storage namespace.
```

6. Wait until there is only the **lvm-operator** pod running.

```
$ oc get pods
NAME                                         READY   STATUS    RESTARTS      AGE
lvm-operator-controller-manager-54df65b5c4-n7bhb   3/3     Running   1 (45h ago)   7d2h
```

7. Change the project to **default**.

```
$ oc project default
```

8. Delete the project **openshift-storage**.

```
$ oc delete project openshift-storage
```

CHAPTER 3. PROVISIONING STORAGE USING LOGICAL VOLUME MANAGER OPERATOR

You can provision persistent volume claims (PVCs) using the storage class that gets created during the operator installation. You can provision block and file PVCs, however, the storage is allocated only when a pod that uses the PVC is created.

NOTE

The Red Hat OpenShift Data Foundation Logical Volume Manager Operator provisions PVCs in units of 1 GiB. The requested storage is rounded up to the nearest GiB.

Procedure

1. Identify the StorageClass that is created when Red Hat OpenShift Data Foundation Logical Volume Manager Operator is deployed.
   The StorageClass name is in the format, **odf-lvm-<device-class-name>**. **device-class-name** is the name of the device class that you provided in the LVMCluster of the policy YAML. For example, if the deviceClass has the name as **vg1**, then the storageClass name is **odf-lvm-vg1**.

2. Save the following YAML to a file with a name such as **odf-lvm-storage-class.yaml** to create a PVC where the application requires storage.

```
# Sample YAML to create a PVC
# block pvc
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: lvm-block-1
  namespace: default
spec:
  accessModes:
    - ReadWriteOnce
  volumeMode: Block
  resources:
    requests:
      storage: 10Gi
  storageClassName: odf-lvm-vg1
---
# file pvc
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: lvm-file-1
  namespace: default
spec:
  accessModes:
    - ReadWriteOnce
  volumeMode: Filesystem
  resources:
    requests:
      storage: 10Gi
  storageClassName: odf-lvm-vg1
```

3. Create the policy by running the following command:

   ```
   # oc create -f odf-lvm-storage-class.yaml -ns lvm-policy-ns
   ```

   The PVCs that are created will remain in pending state until you deploy the pods that use them.

CHAPTER 4. MONITORING THE OPENSHIFT DATA FOUNDATION LOGICAL VOLUME MANAGER OPERATOR

When the OpenShift Data Foundation Logical Volume Manager Operator is installed using the OpenShift Web Console, you can monitor the cluster using the **Block and File** dashboard in the console by default. However, when you use RHACM to install the OpenShift Data Foundation Logical Volume Manager Operator, you need to configure the RHACM Observability to monitor the all the SNO cluster from one place.

You can monitor the OpenShift Data Foundation Logical Volume Manager Operator by viewing the metrics exported by the operator on the RHACM dashboards and the alerts that are triggered. Enable RHACM Observability as described in the Observability guide.

Metrics

- Add the following **topolvm** metrics to the allow list as specified in the Adding custom metrics section:

```
topolvm_thinpool_data_percent
topolvm_thinpool_metadata_percent
topolvm_thinpool_size_bytes
```

NOTE

Metrics are updated every 10 minutes or when there is a change in the thin-pool, such as a new logical volume creation.

Alerts

When the thin pool and volume group are filled up, further operations fail and might lead to data loss. The Logical Volume Manager Operator sends the following alerts the usage of the thin pool and volume group crosses certain value:

Table 4.1. Alerts for Logical Volume Manager cluster in Red Hat Advanced Cluster Management for Kubernetes

Alert	Description
VolumeGroupUsageAtThresholdNearFull	This alert is triggered when both the volume group and thin pool utilization cross 75% on nodes. Data deletion or volume group expansion is required.
VolumeGroupUsageAtThresholdCritical	This alert is triggered when both the volume group and thin pool utilization cross 85% on nodes.VolumeGroup is critically full. Data deletion or volume group expansion is required.
ThinPoolDataUsageAtThresholdNearFull	This alert is triggered when the thin pool data utilization in the volume group crosses 75% on nodes. Data deletion or thin pool expansion is required.

Alert	Description
ThinPoolDataUsageAtThresholdCritical	This alert is triggered when the thin pool data utilization in the volume group crosses 85% on nodes. Data deletion or thin pool expansion is required.
ThinPoolMetaDataUsageAtThresholdNearFull	This alert is triggered when the thin pool metadata utilization in the volume group crosses 75% on nodes. Data deletion or thin pool expansion is required.
ThinPoolMetaDataUsageAtThresholdCritical	This alert is triggered when the thin pool metadata utilization in the volume group crosses 85% on nodes. Data deletion or thin pool expansion is required.

CHAPTER 5. VOLUME SNAPSHOTS FOR SINGLE NODE OPENSHIFT

You can take volume snapshots of persistent volumes (PVs) that are provisioned by the OpenShift Data Foundation Logical Volume Manager Operator. You can also create volume snapshots of the cloned volumes. Volume snapshots help you to:

- Back up your application data (volume snapshots are not backups)
- Revert to a state at which the volume snapshot was taken

You can create volume snapshots based on the available capacity of the thin pool and overprovisioning limits. The Red Hat OpenShift Data Foundation Logical Volume Manager Operator creates a **VolumeSnapshotClass** with the name **odf-lvm-<deviceclass-name>**.

5.1. CREATING VOLUME SNAPSHOTS IN SINGLE NODE OPENSHIFT

Prerequisites

- For a consistent snapshot, ensure that the PVC is in Bound state. Also, ensure that all the I/O to the PVC is stopped before taking the snapshot.

Procedure

1. Save the following YAML to a file with a name such as **odf-lvm-vol-snapshot.yaml** to create a policy for volume snapshot.

```
# Sample YAML to create a volume snapshot

apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshot
metadata:
    name: lvm-block-1-snap
 spec:
     volumeSnapshotClassName: odf-lvm-vg1
     source:
        persistentVolumeClaimName: lvm-block-1
```

2. Create the policy by running the following command:

```
# oc create -f odf-lvm-vol-snapshot.yaml -ns lvm-policy-ns
```

A read only copy of the PVC is created as a volume snapshot.

5.2. RESTORING VOLUME SNAPSHOTS IN SINGLE NODE OPENSHIFT

When you restore a volume snapshot, a new Persistent Volume Claim (PVC) gets created. The restored PVC is independent of the volume snapshot and the source PVC.

Prerequisites

- The storage class must be the same as that of the source PVC.

- The size of the requested PVC must be the same as that of the source volume of the snapshot.

Procedure

1. Identify the storage class name of the source PVC and volume snapshot name.

2. Save the following YAML to a file with a name such as **odf-lvm-vol-restore.yaml** to restore the snapshot.

```
# Sample YAML to restore a PVC.

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: lvm-block-1-restore
spec:
  accessModes:
 - ReadWriteOnce
  volumeMode: Block
  Resources:
    Requests:
      storage: 2Gi
  storageClassName: odf-lvm-vg1
  dataSource:
    name: lvm-block-1-snap
    kind: VolumeSnapshot
    apiGroup: snapshot.storage.k8s.io
```

3. Create the policy by running the following command:

```
# oc create -f odf-lvm-vol-restore.yaml -ns lvm-policy-ns
```

5.3. DELETING VOLUME SNAPSHOTS IN SINGLE NODE OPENSHIFT

Procedure

- To delete the volume snapshot, delete the volume snapshot resource.

```
oc delete volumesnapshot <volume-snapshot-name> -n <namespace>
```

> **NOTE**
>
> When you delete a persistent volume claim (PVC), the snapshots of the PVC are not deleted.

- To delete the restored volume snapshot, delete the PVC that was created to restore the volume snapshot.

```
oc delete pvc <pvc-name> -n <namespace>
```

CHAPTER 6. VOLUME CLONING FOR SINGLE NODE OPENSHIFT

A clone is a duplicate of an existing storage volume that can be used like any standard volume. You create a clone of a volume to make a point in time copy of the data. A persistent volume claim (PVC) cannot be cloned with a different size.

6.1. CREATING VOLUME CLONES IN SINGLE NODE OPENSHIFT

Prerequisites

- Ensure that the source PVC is in Bound state and not in use.
- Ensure that the StorageClass is the same as that of the parent.

Procedure

1. Identify the storage class of the source PVC.
2. Save the following YAML to a file with a name such as **odf-lvm-vol-clone.yaml** to create a volume clone.

```
# Sample YAML to clone a volume
# pvc-clone.yaml
apiVersion: v1
kind: PersistentVolumeClaim
Metadata:
  name: lvm-block-1-clone
Spec:
  storageClassName: odf-lvm-vg1
  dataSource:
    name: lvm-block-1
    kind: PersistentVolumeClaim
  accessModes:
   - ReadWriteOnce
  volumeMode: Block
  Resources:
    Requests:
      storage: 2Gi
The cloned PVC has write access.
```

3. Create the policy by running the following command:

```
# oc create -f odf-lvm-vol-clone.yaml -ns lvm-policy-ns
```

6.2. DELETING CLONED VOLUMES IN SINGLE NODE OPENSHIFT

Procedure

- To delete the cloned volume, you can delete the cloned PVC.

```
oc delete pvc <clone-pvc-name> -n <namespace>
```

-