



Red Hat Advanced Cluster Management for Kubernetes 2.7

Governance

Read more to learn about the governance policy framework, which helps harden cluster security by using policies.

Red Hat Advanced Cluster Management for Kubernetes 2.7 Governance

Read more to learn about the governance policy framework, which helps harden cluster security by using policies.

Legal Notice

Copyright © 2024 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 more to learn about the governance policy framework, which helps harden cluster security by using policies.

Table of Contents

| | |
|---|-----------|
| CHAPTER 1. RISK AND COMPLIANCE | 6 |
| 1.1. CERTIFICATES | 6 |
| 1.1.1. Red Hat Advanced Cluster Management hub cluster certificates | 7 |
| 1.1.1.1. Observability certificates | 7 |
| 1.1.1.2. Bring Your Own (BYO) observability certificate authority (CA) certificates | 8 |
| 1.1.1.2.1. OpenSSL commands to generate CA certificate | 8 |
| 1.1.1.2.2. Create the secrets associated with the BYO observability CA certificates | 8 |
| 1.1.1.2.3. Replacing certificates for alertmanager route | 8 |
| 1.1.2. Red Hat Advanced Cluster Management component certificates | 9 |
| 1.1.2.1. List hub cluster managed certificates | 9 |
| 1.1.2.2. Refresh hub cluster managed certificates | 9 |
| 1.1.2.3. Refresh a Red Hat Advanced Cluster Management webhook certificate | 9 |
| 1.1.3. Red Hat Advanced Cluster Management managed certificates | 10 |
| 1.1.3.1. Channel certificates | 10 |
| 1.1.3.2. Managed cluster certificates | 10 |
| 1.1.4. Third-party certificates | 11 |
| 1.1.4.1. Rotating the gatekeeper webhook certificate | 11 |
| CHAPTER 2. GOVERNANCE | 12 |
| 2.1. GOVERNANCE ARCHITECTURE | 12 |
| 2.2. POLICY OVERVIEW | 14 |
| 2.2.1. Policy YAML structure | 15 |
| 2.2.2. Policy YAML table | 16 |
| 2.2.3. Policy sample file | 18 |
| 2.2.4. Placement YAML sample file | 19 |
| 2.3. POLICY CONTROLLERS | 20 |
| 2.3.1. Kubernetes configuration policy controller | 20 |
| 2.3.1.1. Configuration policy sample | 21 |
| 2.3.1.2. Configuration policy YAML table | 22 |
| 2.3.1.3. Additional resources | 26 |
| 2.3.2. Certificate policy controller | 26 |
| 2.3.2.1. Certificate policy controller YAML structure | 27 |
| 2.3.2.1.1. Certificate policy controller YAML table | 27 |
| 2.3.2.2. Certificate policy sample | 30 |
| 2.3.3. IAM policy controller | 30 |
| 2.3.3.1. IAM policy YAML structure | 30 |
| 2.3.3.2. IAM policy YAML table | 31 |
| 2.3.3.3. IAM policy sample | 32 |
| 2.3.4. Policy set controller | 32 |
| 2.3.4.1. Policy set YAML structure | 32 |
| 2.3.4.2. Policy set table | 33 |
| 2.3.4.3. Policy set sample | 34 |
| 2.4. POLICY CONTROLLER ADVANCED CONFIGURATION | 34 |
| 2.4.1. Configure the concurrency | 34 |
| 2.4.2. Configure debug log | 35 |
| 2.4.3. Governance metric | 35 |
| 2.4.3.1. Metric: policy_governance_info | 35 |
| 2.4.3.2. Metric: config_policies_evaluation_duration_seconds | 36 |
| 2.4.4. Verify configuration changes | 36 |
| 2.5. SUPPORTED POLICIES | 37 |
| 2.5.1. Table of sample configuration policies | 37 |

| | |
|--|----|
| 2.5.2. Support matrix for out-of-box policies | 38 |
| 2.5.3. Memory usage policy | 39 |
| 2.5.3.1. Memory usage policy YAML structure | 40 |
| 2.5.3.2. Memory usage policy table | 40 |
| 2.5.3.3. Memory usage policy sample | 41 |
| 2.5.4. Namespace policy | 41 |
| 2.5.4.1. Namespace policy YAML structure | 41 |
| 2.5.4.2. Namespace policy YAML table | 42 |
| 2.5.4.3. Namespace policy sample | 43 |
| 2.5.5. Image vulnerability policy | 43 |
| 2.5.5.1. Image vulnerability policy YAML structure | 43 |
| 2.5.5.2. Image vulnerability policy sample | 45 |
| 2.5.6. Pod policy | 45 |
| 2.5.6.1. Pod policy YAML structure | 45 |
| 2.5.6.2. Pod policy table | 46 |
| 2.5.6.3. Pod policy sample | 46 |
| 2.5.7. Pod security policy (Deprecated) | 47 |
| 2.5.7.1. Pod security policy YAML structure | 47 |
| 2.5.7.2. Pod security policy table | 48 |
| 2.5.7.3. Pod security policy sample | 48 |
| 2.5.8. Role policy | 49 |
| 2.5.8.1. Role policy YAML structure | 49 |
| 2.5.8.2. Role policy table | 50 |
| 2.5.8.3. Role policy sample | 51 |
| 2.5.9. Role binding policy | 51 |
| 2.5.9.1. Role binding policy YAML structure | 51 |
| 2.5.9.2. Role binding policy table | 52 |
| 2.5.9.3. Role binding policy sample | 53 |
| 2.5.10. Security Context Constraints policy | 53 |
| 2.5.10.1. SCC policy YAML structure | 53 |
| 2.5.10.2. SCC policy table | 54 |
| 2.5.10.3. SCC policy sample | 55 |
| 2.5.11. ETCD encryption policy | 55 |
| 2.5.11.1. ETCD encryption policy YAML structure | 56 |
| 2.5.11.2. ETCD encryption policy table | 56 |
| 2.5.11.3. ETCD encryption policy sample | 57 |
| 2.5.12. Compliance Operator policy | 57 |
| 2.5.12.1. Compliance Operator policy overview | 57 |
| 2.5.12.2. Compliance operator resources | 58 |
| 2.5.12.3. Additional resources | 59 |
| 2.5.13. E8 scan policy | 59 |
| 2.5.13.1. E8 scan policy resources | 59 |
| 2.5.14. OpenShift CIS scan policy | 61 |
| 2.5.14.1. OpenShift CIS resources | 61 |
| 2.6. MANAGE SECURITY POLICIES | 62 |
| 2.6.1. Governance page | 63 |
| 2.6.2. Governance automation configuration | 63 |
| 2.6.3. Configuring Ansible Automation Platform for governance | 64 |
| 2.6.3.1. Prerequisites | 64 |
| 2.6.3.2. Creating a policy violation automation from the console | 65 |
| 2.6.3.3. Creating a policy violation automation from the CLI | 65 |
| 2.6.4. Deploying policies using GitOps | 66 |
| 2.6.4.1. Customizing your local repository | 67 |

| | |
|---|----|
| 2.6.4.2. Committing to your local repository | 67 |
| 2.6.4.3. Deploying policies to your cluster | 68 |
| 2.6.4.4. Verifying GitOps policy deployments from the console | 70 |
| 2.6.4.4.1. Verifying GitOps policy deployments from the CLI | 70 |
| 2.6.5. Support for templates in configuration policies | 71 |
| 2.6.5.1. Prerequisite | 71 |
| 2.6.5.2. Template functions | 71 |
| 2.6.5.2.1. fromSecret function | 72 |
| 2.6.5.2.2. fromConfigMap function | 73 |
| 2.6.5.2.3. fromClusterClaim function | 74 |
| 2.6.5.2.4. lookup function | 74 |
| 2.6.5.2.5. base64enc function | 75 |
| 2.6.5.2.6. base64dec function | 76 |
| 2.6.5.2.7. indent function | 76 |
| 2.6.5.2.8. autoindent function | 77 |
| 2.6.5.2.9. toInt function | 77 |
| 2.6.5.2.10. toBool function | 78 |
| 2.6.5.2.11. protect function | 78 |
| 2.6.5.2.12. toLiteral function | 79 |
| 2.6.5.2.13. Open source community functions | 79 |
| 2.6.5.3. Support for hub cluster templates in configuration policies | 80 |
| 2.6.5.3.1. Template processing | 80 |
| 2.6.5.3.2. Special annotation for reprocessing | 81 |
| 2.6.5.4. Object template processing | 81 |
| 2.6.5.4.1. Bypass template processing | 81 |
| 2.6.5.4.2. Comparison of hub cluster and managed cluster templates | 82 |
| 2.6.6. Managing security policies | 83 |
| 2.6.6.1. Creating a security policy | 84 |
| 2.6.6.1.1. Creating a security policy from the command line interface | 84 |
| 2.6.6.1.1.1. Viewing your security policy from the CLI | 86 |
| 2.6.6.1.2. Creating a cluster security policy from the console | 86 |
| 2.6.6.1.2.1. Viewing your security policy from the console | 87 |
| 2.6.6.1.3. Creating policy sets from the CLI | 87 |
| 2.6.6.1.4. Creating policy sets from the console | 87 |
| 2.6.6.2. Updating security policies | 88 |
| 2.6.6.2.1. Adding a policy to a policy set from the CLI | 88 |
| 2.6.6.2.2. Adding a policy to a policy set from the console | 88 |
| 2.6.6.2.3. Disabling security policies | 88 |
| 2.6.6.3. Deleting a security policy | 88 |
| 2.6.6.3.1. Deleting policy sets from the console | 89 |
| 2.6.6.4. Cleaning up resources that are created by policies | 89 |
| 2.6.7. Managing configuration policies | 89 |
| 2.6.7.1. Creating a configuration policy | 90 |
| 2.6.7.1.1. Creating a configuration policy from the CLI | 90 |
| 2.6.7.1.2. Viewing your configuration policy from the CLI | 90 |
| 2.6.7.1.3. Creating a configuration policy from the console | 91 |
| 2.6.7.1.4. Viewing your configuration policy from the console | 91 |
| 2.6.7.2. Updating configuration policies | 91 |
| 2.6.7.2.1. Disabling configuration policies | 91 |
| 2.6.7.3. Deleting a configuration policy | 92 |
| 2.6.8. Managing gatekeeper operator policies | 92 |
| 2.6.8.1. Installing gatekeeper using a gatekeeper operator policy | 93 |
| 2.6.8.2. Creating a gatekeeper policy from the console | 93 |

| | |
|---|-----|
| 2.6.8.2.1. Gatekeeper operator CR | 93 |
| 2.6.8.3. Upgrading gatekeeper and the gatekeeper operator | 94 |
| 2.6.8.4. Updating gatekeeper operator policy | 94 |
| 2.6.8.4.1. Viewing gatekeeper operator policy from the console | 94 |
| 2.6.8.4.2. Disabling gatekeeper operator policy | 95 |
| 2.6.8.5. Deleting gatekeeper operator policy | 95 |
| 2.6.8.6. Uninstalling gatekeeper policy, gatekeeper, and gatekeeper operator policy | 95 |
| 2.6.9. Managing operator policies in disconnected environments | 96 |
| 2.7. POLICY DEPENDENCIES | 96 |
| 2.8. SECURE THE HUB CLUSTER | 98 |
| 2.9. INTEGRATE THIRD-PARTY POLICY CONTROLLERS | 98 |
| 2.9.1. Integrating gatekeeper constraints and constraint templates | 98 |
| 2.9.2. Policy Generator | 100 |
| 2.9.2.1. Policy Generator capabilities | 101 |
| 2.9.2.2. Policy Generator configuration structure | 101 |
| 2.9.2.3. Policy Generator configuration reference table | 103 |
| 2.9.2.4. Related resources | 112 |
| 2.9.3. Generating a policy to install an Operator | 112 |
| 2.9.3.1. Generating a policy that installs OpenShift GitOps | 112 |
| 2.9.3.2. Generating a policy that installs the Compliance Operator | 115 |
| 2.9.3.3. Using policy dependencies with OperatorGroups | 117 |
| 2.9.3.4. Related resources | 117 |
| 2.9.4. Managing policy definitions with OpenShift GitOps (ArgoCD) | 118 |
| 2.9.4.1. Integrating the Policy Generator with OpenShift GitOps (ArgoCD) | 119 |
| 2.9.4.2. Related resource | 121 |

CHAPTER 1. RISK AND COMPLIANCE

Manage your security of Red Hat Advanced Cluster Management for Kubernetes components. Govern your cluster with defined policies and processes to identify and minimize risks. Use policies to define rules and set controls.

Prerequisite: You must configure authentication service requirements for Red Hat Advanced Cluster Management for Kubernetes. See [Access control](#) for more information.

Review the following topics to learn more about securing your cluster:

- [Role-based access control](#)
- [Managing credentials overview](#)
- [Certificates](#)
- [Governance](#)

1.1. CERTIFICATES

Various certificates are created and used throughout Red Hat Advanced Cluster Management for Kubernetes.

You can bring your own certificates. You must create a Kubernetes TLS Secret for your certificate. After you create your certificates, you can replace certain certificates that are created by the Red Hat Advanced Cluster Management installer.

Required access: Cluster administrator

All certificates required by services that run on Red Hat Advanced Cluster Management are created during the installation of Red Hat Advanced Cluster Management. Certificates are created and managed by the following components:

- [OpenShift Service Serving Certificates](#)
- Red Hat Advanced Cluster Management webhook controllers
- Kubernetes Certificates API
- OpenShift default ingress

Continue reading to learn more about certificate management:

[Red Hat Advanced Cluster Management hub cluster certificates](#)

- [Replacing the OpenShift default ingress certificate](#)
- [Observability certificates](#)
 - [Bring Your Own \(BYO\) observability certificate authority \(CA\) certificates](#)
 - [OpenSSL commands to generate CA certificate](#)
 - [Create the secrets associated with the BYO observability CA certificates](#)
 - [Replacing certificates for alertmanager route](#)

Red Hat Advanced Cluster Management component certificates

- [List hub cluster managed certificates](#)
- [Refresh hub cluster managed certificates](#)
- [Refresh a Red Hat Advanced Cluster Management webhook certificate](#)

Red Hat Advanced Cluster Management managed certificates

- [Channel certificates](#)
- [Managed cluster certificates](#)

Third-party certificates

- [Rotating the gatekeeper webhook certificate](#)

Note: Users are responsible for certificate rotations and updates.

1.1.1. Red Hat Advanced Cluster Management hub cluster certificates

1.1.1.1. Observability certificates

After Red Hat Advanced Cluster Management is installed, observability certificates are created and used by the observability components, to provide mutual TLS on the traffic between the hub cluster and managed cluster. The Kubernetes secrets that are associated with the observability certificates.

The **open-cluster-management-observability** namespace contain the following certificates:

- **observability-server-ca-certs:** Has the CA certificate to sign server-side certificates
- **observability-client-ca-certs:** Has the CA certificate to sign client-side certificates
- **observability-server-certs:** Has the server certificate used by the **observability-observatorium-api** deployment
- **observability-grafana-certs:** Has the client certificate used by the **observability-rbac-query-proxy** deployment

The **open-cluster-management-addon-observability** namespace contain the following certificates on managed clusters:

- **observability-managed-cluster-certs:** Has the same server CA certificate as **observability-server-ca-certs** in the hub server
- **observability-controller-open-cluster-management.io-observability-signer-client-cert:** Has the client certificate used by the **metrics-collector-deployment**

The CA certificates are valid for five years and other certificates are valid for one year. All observability certificates are automatically refreshed upon expiration.

View the following list to understand the effects when certificates are automatically renewed:

- Non-CA certificates are renewed automatically when the remaining valid time is no more than 73 days. After the certificate is renewed, the pods in the related deployments restart automatically to use the renewed certificates.

- CA certificates are renewed automatically when the remaining valid time is no more than one year. After the certificate is renewed, the old CA is not deleted but co-exist with the renewed ones. Both old and renewed certificates are used by related deployments, and continue to work. The old CA certificates are deleted when they expire.
- When a certificate is renewed, the traffic between the hub cluster and managed cluster is not interrupted.

1.1.1.2. Bring Your Own (BYO) observability certificate authority (CA) certificates

If you do not want to use the default observability CA certificates generated by Red Hat Advanced Cluster Management, you can choose to use the BYO observability CA certificates before you enable observability.

1.1.1.2.1. OpenSSL commands to generate CA certificate

Observability requires two CA certificates; one is for the server-side and the other is for the client-side.

- Generate your CA RSA private keys with the following commands:

```
openssl genrsa -out serverCAKey.pem 2048
```

```
openssl genrsa -out clientCAKey.pem 2048
```

- Generate the self-signed CA certificates using the private keys. Run the following commands:

```
openssl req -x509 -sha256 -new -nodes -key serverCAKey.pem -days 1825 -out serverCACert.pem
```

```
openssl req -x509 -sha256 -new -nodes -key clientCAKey.pem -days 1825 -out clientCACert.pem
```

1.1.1.2.2. Create the secrets associated with the BYO observability CA certificates

Complete the following steps to create the secrets:

1. Create the **observability-server-ca-certs** secret by using your certificate and private key. Run the following command:

```
oc -n open-cluster-management-observability create secret tls observability-server-ca-certs -cert ./serverCACert.pem --key ./serverCAKey.pem
```

2. Create the **observability-client-ca-certs** secret by using your certificate and private key. Run the following command:

```
oc -n open-cluster-management-observability create secret tls observability-client-ca-certs --cert ./clientCACert.pem --key ./clientCAKey.pem
```

1.1.1.2.3. Replacing certificates for alertmanager route

You can replace alertmanager certificates by updating the alertmanager route, if you do not want to use the OpenShift default ingress certificate. Complete the following steps:

1. Examine the observability certificate with the following command:

```
openssl x509 -noout -text -in ./observability.crt
```

2. Change the common name (**CN**) on the certificate to **alertmanager**.
3. Change the SAN in the **csr.cnf** configuration file with the hostname for your alertmanager route.
4. Create the two following secrets in the **open-cluster-management-observability** namespace. Run the following command:

```
oc -n open-cluster-management-observability create secret tls alertmanager-byo-ca --cert ./ca.crt --key ./ca.key
```

```
oc -n open-cluster-management-observability create secret tls alertmanager-byo-cert --cert ./ingress.crt --key ./ingress.key
```

1.1.2. Red Hat Advanced Cluster Management component certificates

1.1.2.1. List hub cluster managed certificates

You can view a list of hub cluster managed certificates that use [OpenShift Service Serving Certificates](#) service internally. Run the following command to list the certificates:

```
for ns in multicluster-engine open-cluster-management ; do echo "$ns:" ; oc get secret -n $ns -o custom-  
columns=Name:.metadata.name,Expiration:.metadata.annotations.service\\.beta\\.openshift\\.io/expiry  
| grep -v '<none>' ; echo "" ; done
```

Note: If observability is enabled, there are additional namespaces where certificates are created.

1.1.2.2. Refresh hub cluster managed certificates

You can refresh a hub cluster managed certificate by running the **delete secret** command in the [List hub cluster managed certificates](#) section. When you identify the certificate that you need to refresh, delete the secret that is associated with the certificate. For example, you can delete a secret by running the following command:

```
oc delete secret grc-0c925-grc-secrets -n open-cluster-management
```

Note: After you delete the secret, a new one is created. However, you must restart pods that use the secret manually so they can begin to use the new certificate.

1.1.2.3. Refresh a Red Hat Advanced Cluster Management webhook certificate

You can refresh OpenShift Container Platform managed certificates, which are certificates that are used by Red Hat Advanced Cluster Management webhooks.

Complete the following steps to refresh Red Hat Advanced Cluster Management webhook certificate:

1. Delete the secret that is associated with the OpenShift Container Platform managed certificate by running the following command:

```
oc delete secret -n open-cluster-management ocm-webhook-secret
```

Note: Some services might not have a secret that needs to be deleted.

- Restart the services that are associated with the OpenShift Container Platform managed certificate(s) by running the following command:

```
oc delete po -n open-cluster-management ocm-webhook-679444669c-5cg76
```

Important: There are replicas of many services; each service must be restarted.

View the following table for a summarized list of the pods that contain certificates and whether a secret needs to be deleted prior to restarting the pod:

Table 1.1. Pods that contain OpenShift Container Platform managed certificates

| Service name | Namespace | Sample pod name | Secret name (if applicable) |
|---|-----------------------------|--|--|
| channels-apps-open-cluster-management-webhook-svc | open-cluster-management | multicluster-operators-application-8c446664c-5lbfk | channels-apps-open-cluster-management-webhook-svc-ca |
| multicluster-operators-application-svc | open-cluster-management | multicluster-operators-application-8c446664c-5lbfk | multicluster-operators-application-svc-ca |
| cluster-manager-registration-webhook | open-cluster-management-hub | cluster-manager-registration-webhook-fb7b99c-d8wfc | registration-webhook-serving-cert |
| cluster-manager-work-webhook | open-cluster-management-hub | cluster-manager-work-webhook-89b8d7fc-f4pv8 | work-webhook-serving-cert |

1.1.3. Red Hat Advanced Cluster Management managed certificates

1.1.3.1. Channel certificates

CA certificates can be associated with Git channel that are a part of the Red Hat Advanced Cluster Management application management. See [Using custom CA certificates for a secure HTTPS connection](#) for more details.

Helm channels allow you to disable certificate validation. Helm channels where certificate validation is disabled, must be configured in development environments. Disabling certificate validation introduces security risks.

1.1.3.2. Managed cluster certificates

Certificates are used to authenticate managed clusters with the hub. Therefore, it is important to be aware of troubleshooting scenarios associated with these certificates. View [Troubleshooting imported clusters offline after certificate change](#) for more details.

The managed cluster certificates are refreshed automatically.

1.1.4. Third-party certificates

1.1.4.1. Rotating the gatekeeper webhook certificate

Complete the following steps to rotate the gatekeeper webhook certificate:

1. Edit the secret that contains the certificate with the following command:

```
oc edit secret -n openshift-gatekeeper-system gatekeeper-webhook-server-cert
```

2. Delete the following content in the **data** section: **ca.crt**, **ca.key**, **tls.crt**, and **tls.key**.
3. Restart the gatekeeper webhook service by deleting the **gatekeeper-controller-manager** pods with the following command:

```
oc delete po -n openshift-gatekeeper-system -l control-plane=controller-manager
```

The gatekeeper webhook certificate is rotated.

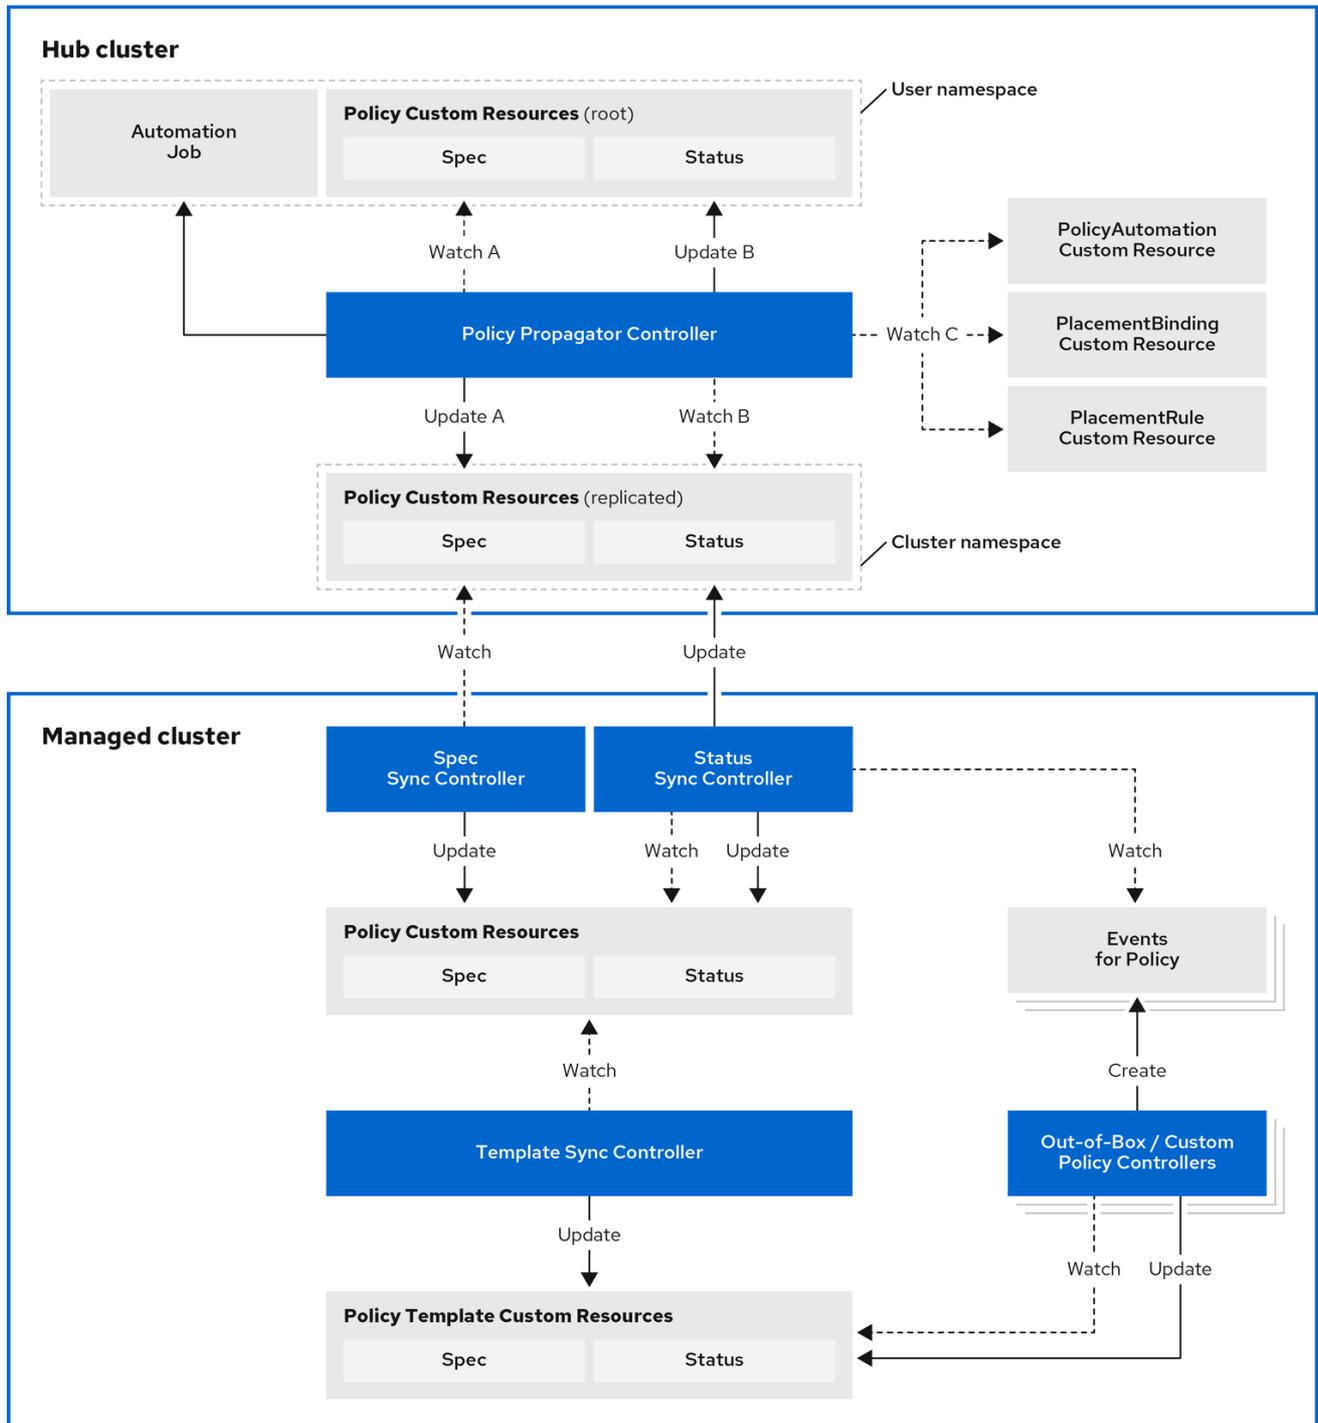
Use the certificate policy controller to create and manage certificate policies on managed clusters. See [Policy controllers](#) to learn more about controllers. Return to the [Risk and compliance](#) page for more information.

CHAPTER 2. GOVERNANCE

Enterprises must meet internal standards for software engineering, secure engineering, resiliency, security, and regulatory compliance for workloads hosted on private, multi and hybrid clouds. Red Hat Advanced Cluster Management for Kubernetes governance provides an extensible policy framework for enterprises to introduce their own security policies.

2.1. GOVERNANCE ARCHITECTURE

Enhance the security for your cluster with the Red Hat Advanced Cluster Management for Kubernetes governance lifecycle. The product governance lifecycle is based on defined policies, processes, and procedures to manage security and compliance from a central interface page. View the following diagram of the governance architecture:



186_RHACM_I221

The governance architecture is composed of the following components:

- **Governance dashboard:** Provides a summary of your cloud governance and risk details, which include policy and cluster violations. Refer to the *Manage security policies* section to learn about the structure of an Red Hat Advanced Cluster Management for Kubernetes policy framework, and how to use the Red Hat Advanced Cluster Management for Kubernetes *Governance* dashboard.

Notes:

- When a policy is propagated to a managed cluster, it is first replicated to the cluster namespace on the hub cluster, and is named and labeled using **namespaceName.policyName**. When you create a policy, make sure that the length of the

namespaceName.policyName does not exceed 63 characters due to the Kubernetes length limit for label values.

- When you search for a policy in the hub cluster, you might also receive the name of the replicated policy in the managed cluster namespace. For example, if you search for **policy-dhaz-cert** in the **default** namespace, the following policy name from the hub cluster might also appear in the managed cluster namespace: **default.policy-dhaz-cert**.
- **Policy-based governance framework:** Supports policy creation and deployment to various managed clusters based on attributes associated with clusters, such as a geographical region. There are examples of the predefined policies and instructions on deploying policies to your cluster in the *open source community*. Additionally, when policies are violated, automations can be configured to run and take any action that the user chooses.
- **Policy controller:** Evaluates one or more policies on the managed cluster against your specified control and generates Kubernetes events for violations. Violations are propagated to the hub cluster. Policy controllers that are included in your installation are the following: Kubernetes configuration, Certificate, and IAM. Customize the your policy controllers using advanced configurations.
- **Open source community:** Supports community contributions with a foundation of the Red Hat Advanced Cluster Management policy framework. Policy controllers and third-party policies are also a part of the [stolostron/policy-collection repository](#). You can contribute and deploy policies using GitOps. For more information, see *Deploying policies using GitOps* in the *Manage security policies* section. Learn how to integrate third-party policies with Red Hat Advanced Cluster Management for Kubernetes.

Continue reading the related topics:

- [Policy overview](#)
- [Policy controllers](#)
- [Policy controller advanced configuration](#)
- [Supported policies](#)
- [Policy dependencies](#)
- [Manage security policies](#)
- [Secure the hub cluster](#)
- [Integrate third-party policy controllers](#)

2.2. POLICY OVERVIEW

Use the Red Hat Advanced Cluster Management for Kubernetes security policy framework to create and manage policies. Kubernetes custom resource definition instances are used to create policies.

Each Red Hat Advanced Cluster Management policy can have at least one or more templates. For more details about the policy elements, view the [Policy YAML table](#) section on this page.

The policy requires a *PlacementRule* or *Placement* that defines the clusters that the policy document is applied to, and a *PlacementBinding* that binds the Red Hat Advanced Cluster Management for Kubernetes policy to the placement rule. For more on how to define a **PlacementRule**, see [Placement](#)

[rules](#) in the Application lifecycle documentation. For more on how to define a **Placement** see [Placement overview](#) in the Cluster lifecycle documentation.

Important:

- You must create the **PlacementBinding** to bind your policy with either a **PlacementRule** or a **Placement** in order to propagate the policy to the managed clusters.
Best practice: Use the command line interface (CLI) to make updates to the policies when you use the **Placement** resource.
- You can create a policy in any namespace on the hub cluster except the managed cluster namespaces. Every managed cluster has a namespace on the hub cluster that matches the name of the managed cluster. If you create a policy in a managed cluster namespace, it is deleted by Red Hat Advanced Cluster Management for Kubernetes. Instead, use a namespace such as **policy**.
- Each client and provider is responsible for ensuring that their managed cloud environment meets internal enterprise security standards for software engineering, secure engineering, resiliency, security, and regulatory compliance for workloads hosted on Kubernetes clusters. Use the governance and security capability to gain visibility and remediate configurations to meet standards.

Learn more details about the policy components in the following sections:

- [Policy YAML structure](#)
- [Policy YAML table](#)
- [Policy sample file](#)
- [Placement YAML sample file](#)

2.2.1. Policy YAML structure

When you create a policy, you must include required parameter fields and values. Depending on your policy controller, you might need to include other optional fields and values. View the following YAML structure for the explained parameter fields:

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  dependencies:
  - apiVersion: policy.open-cluster-management.io/v1
  compliance:
  kind: Policy
  name:
  namespace:
  policy-templates:
  - objectDefinition:
    apiVersion:
```

```

    kind:
    metadata:
      name:
    spec:
    remediationAction:
    disabled:
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name:
placementRef:
  name:
  kind:
  apiGroup:
subjects:
- name:
  kind:
  apiGroup:
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name:
spec:
  clusterConditions:
  - type:
  clusterLabels:
  matchLabels:
    cloud:

```

2.2.2. Policy YAML table

Table 2.1. Parameter table

| Field | Optional or required | Description |
|----------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |

| Field | Optional or required | Description |
|---|----------------------|---|
| metadata.annotations | Optional | Used to specify a set of security details that describes the set of standards the policy is trying to validate. All annotations documented here are represented as a string that contains a comma-separated list. Note: You can view policy violations based on the standards and categories that you define for your policy on the <i>Policies</i> page, from the console. |
| annotations.policy.open-cluster-management.io/standards | Optional | The name or names of security standards the policy is related to. For example, National Institute of Standards and Technology (NIST) and Payment Card Industry (PCI). |
| annotations.policy.open-cluster-management.io/categories | Optional | A security control category represent specific requirements for one or more standards. For example, a System and Information Integrity category might indicate that your policy contains a data transfer protocol to protect personal information, as required by the HIPAA and PCI standards. |
| annotations.policy.open-cluster-management.io/controls | Optional | The name of the security control that is being checked. For example, Access Control or System and Information Integrity. |
| spec.dependencies | Optional | Used to create a list of dependency objects detailed with extra considerations for compliance. |
| spec.policy-templates | Required | Used to create one or more policies to apply to a managed cluster. |

| Field | Optional or required | Description |
|--|----------------------|---|
| spec.policy-templates.extraDependencies | Optional | For policy templates, this is used to create a list of dependency objects detailed with extra considerations for compliance. |
| spec.policy-templates.ignorePending | Optional | Used to mark a policy template as compliant until the dependency criteria is verified. |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.remediationAction | Optional. | Specifies the remediation of your policy. The parameter values are enforce and inform . If specified, the spec.remediationAction value that is defined overrides any remediationAction parameter defined in the child policies in the policy-templates section. For example, if the spec.remediationAction value is set to enforce , then the remediationAction in the policy-templates section is set to enforce during runtime. Important: Some policy kinds might not support the enforce feature. |

2.2.3. Policy sample file

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy-role
  annotations:
    policy.open-cluster-management.io/standards: NIST SP 800-53
    policy.open-cluster-management.io/categories: AC Access Control
    policy.open-cluster-management.io/controls: AC-3 Access Enforcement
spec:
  remediationAction: inform
  disabled: false
  policy-templates:
    - objectDefinition:

```

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-role-example
spec:
  remediationAction: inform # the policy-template spec.remediationAction is overridden by the
preceding parameter value for spec.remediationAction.
  severity: high
  namespaceSelector:
    include: ["default"]
  object-templates:
    - complianceType: mustonlyhave # role definition should exact match
      objectDefinition:
        apiVersion: rbac.authorization.k8s.io/v1
        kind: Role
        metadata:
          name: sample-role
        rules:
          - apiGroups: ["extensions", "apps"]
            resources: ["deployments"]
            verbs: ["get", "list", "watch", "delete", "patch"]
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-policy-role
placementRef:
  name: placement-policy-role
  kind: PlacementRule
  apiGroup: apps.open-cluster-management.io
subjects:
  - name: policy-role
    kind: Policy
    apiGroup: policy.open-cluster-management.io
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-policy-role
spec:
  clusterConditions:
    - status: "True"
      type: ManagedClusterConditionAvailable
  clusterSelector:
    matchExpressions:
      - {key: environment, operator: In, values: ["dev"]}

```

2.2.4. Placement YAML sample file

The **PlacementBinding** and **Placement** resources can be combined with the previous policy example to deploy the policy using the cluster **Placement** API instead of the **PlacementRule** API.

```

---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding

```

```

metadata:
  name: binding-policy-role
placementRef:
  name: placement-policy-role
  kind: Placement
  apiGroup: cluster.open-cluster-management.io
subjects:
- name: policy-role
  kind: Policy
  apiGroup: policy.open-cluster-management.io
---
//Depends on if governance would like to use v1beta1
apiVersion: cluster.open-cluster-management.io/v1beta1
kind: Placement
metadata:
  name: placement-policy-role
spec:
  predicates:
  - requiredClusterSelector:
    labelSelector:
      matchExpressions:
      - {key: environment, operator: In, values: ["dev"]}

```

- Refer to [Policy controllers](#).
- See [Managing security policies](#) to create and update a policy. You can also enable and update Red Hat Advanced Cluster Management policy controllers to validate the compliance of your policies.
- Return to the [Governance](#) documentation.

2.3. POLICY CONTROLLERS

Policy controllers monitor and report whether your cluster is compliant with a policy. Use the Red Hat Advanced Cluster Management for Kubernetes policy framework by using the out-of-the-box policy templates to apply policies managed by these controllers. The policy controllers manage Kubernetes custom resource definition (CRD) instances.

Policy controllers monitor for policy violations, and can make the cluster status compliant if the controller supports the enforcement feature. View the following topics to learn more about the following Red Hat Advanced Cluster Management for Kubernetes policy controllers:

- [Kubernetes configuration policy controller](#)
- [Certificate policy controller](#)
- [IAM policy controller](#)
- [Policy set controller](#)

Important: Only the configuration policy controller policies support the **enforce** feature. You must manually remediate policies, where the policy controller does not support the **enforce** feature.

Refer to [Governance](#) for more topics about managing your policies.

2.3.1. Kubernetes configuration policy controller

The configuration policy controller can be used to configure any Kubernetes resource and apply security policies across your clusters. The configuration policy is provided in the **policy-templates** field of the policy on the hub cluster, and is propagated to the selected managed clusters by the governance framework. See the [Policy overview](#) for more details on the hub cluster policy.

A Kubernetes object is defined (in whole or in part) in the **object-templates** array in the configuration policy, indicating to the configuration policy controller of the fields to compare with objects on the managed cluster. The configuration policy controller communicates with the local Kubernetes API server to get the list of your configurations that are in your cluster.

The configuration policy controller is created on the managed cluster during installation. The configuration policy controller supports the **enforce** feature to remediate when the configuration policy is non-compliant. When the **remediationAction** for the configuration policy is set to **enforce**, the controller applies the specified configuration to the target managed cluster. **Note:** Configuration policies that specify an object without a name can only be **inform**.

You can also use templated values within configuration policies. For more information, see [Support for templates in configuration policies](#).

If you have existing Kubernetes manifests that you want to put in a policy, the [Policy Generator](#) is a useful tool to accomplish this.

Continue reading to learn more about the configuration policy controller:

- [Configuration policy sample](#)
- [Configuration policy YAML table](#)

2.3.1.1. Configuration policy sample

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-config
spec:
  namespaceSelector:
    include: ["default"]
    exclude: []
    matchExpressions: []
    matchLabels: {}
  remediationAction: inform
  severity: low
  evaluationInterval:
    compliant:
    noncompliant:
  object-templates:
  - complianceType: musthave
    objectDefinition:
      apiVersion: v1
      kind: Pod
      metadata:
        name: pod
      spec:
        containers:
        - image: pod-image
          name: pod-name

```

```

    ports:
      - containerPort: 80
  object-templates-raw: |
  - complianceType: musthave
    objectDefinition:
      apiVersion: v1
      kind: ConfigMap
      metadata:
        name: myconfig
        namespace: default
      data:
        testData: hello
    spec:
  ...

```

2.3.1.2. Configuration policy YAML table

Table 2.2. Parameter table

| Field | Optional or required | Description |
|-------------------------------|--|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to ConfigurationPolicy to indicate the type of policy. |
| metadata.name | Required | The name of the policy. |
| spec.namespaceSelector | Required for namespaced objects that do not have a namespace specified | Determines namespaces in the managed cluster that the object is applied to. The include and exclude parameters accept file path expressions to include and exclude namespaces by name. The matchExpressions and matchLabels parameters specify namespaces to include by label. See the Kubernetes labels and selectors documentation. The resulting list is compiled by using the intersection of results from all parameters. |
| spec.remediationAction | Required | Specifies the action to take when the policy is non-compliant. Use the following parameter values: inform or enforce . |

| Field | Optional or required | Description |
|---|----------------------|---|
| spec.severity | Required | Specifies the severity when the policy is non-compliant. Use the following parameter values: low , medium , high , or critical . |
| spec.evaluationInterval.compliant | Optional | <p>Used to define how often the policy is evaluated when it is in the compliant state. The values must be in the format of a duration which is a sequence of numbers with time unit suffixes. For example, 12h30m5s represents 12 hours, 30 minutes, and 5 seconds. It can also be set to never so that the policy is not reevaluated on the compliant cluster, unless the policy spec is updated.</p> <p>By default, the minimum time between evaluations for configuration policies is approximately 10 seconds when the evaluationInterval.compliant is not set or empty. This can be longer if the configuration policy controller is saturated on the managed cluster.</p> |
| spec.evaluationInterval.noncompliant | Optional | Used to define how often the policy is evaluated when it is in the non-compliant state. Similar to the evaluationInterval.compliant parameter, the values must be in the format of a duration which is a sequence of numbers with time unit suffixes. It can also be set to never so that the policy is not reevaluated on the non-compliant cluster, unless the policy spec is updated. |

| Field | Optional or required | Description |
|---------------------------------------|----------------------|--|
| spec.object-templates | Optional | The array of Kubernetes objects (either fully defined or containing a subset of fields) for the controller to compare with objects on the managed cluster. Note: While spec.object-templates and spec.object-templates-raw are optional, at least one of the two parameter fields must be set. |
| spec.object-templates-raw | Optional | Used to set object templates with a raw YAML string. Note: While spec.object-templates and spec.object-templates-raw are optional, at least one of the two parameter fields must be set. |
| spec.object-templates-raw.spec | Optional | Used to specify conditions for the object templates. If else statements are supported values. For example, add the following value to avoid duplication in your object-templates definition: <pre> {{- if eq .metadata.name "policy-grc-your-meta-data-name" }} replicas: 2 {{- else }} replicas: 1 {{- end }} </pre> |

| Field | Optional or required | Description |
|---|----------------------|---|
| spec.object-templates[].complianceType | Required | <p>Used to define the desired state of the Kubernetes object on the managed clusters. You must use one of the following verbs as the parameter value:</p> <p>mustonlyhave: Indicates that an object must exist with the exact fields and values as defined in the objectDefinition.</p> <p>musthave: Indicates an object must exist with the same fields as specified in the objectDefinition. Any existing fields on the object that are not specified in the object-template are ignored. In general, array values are appended. The exception for the array to be patched is when the item contains a name key with a value that matches an existing item. Use a fully defined objectDefinition using the mustonlyhave compliance type, if you want to replace the array.</p> <p>mustnothave: Indicates that an object with the same fields as specified in the objectDefinition cannot exist.</p> |
| spec.object-templates[].metadataComplianceType | Optional | <p>Overrides spec.object-templates[].complianceType when comparing the manifest's metadata section to objects on the cluster ("musthave", "mustonlyhave"). Default is unset to not override complianceType for metadata.</p> |
| spec.object-templates[].objectDefinition | Required | <p>A Kubernetes object (either fully defined or containing a subset of fields) for the controller to compare with objects on the managed cluster.</p> |
| spec.pruneObjectBehavior | Optional | <p>Determines whether to clean up resources related to the policy when the policy is removed from a managed cluster.</p> |

2.3.1.3. Additional resources

See the following topics for more information:

- See the [Policy overview](#) for more details on the hub cluster policy.
- See the policy samples that use [NIST Special Publication 800-53 \(Rev. 4\)](#) , and are supported by Red Hat Advanced Cluster Management from the [CM-Configuration-Management](#) folder.
- Learn about how policies are applied on your hub cluster, see [Supported policies](#) for more details.
- Refer to [Policy controllers](#) for more details about controllers.
- Customize your policy controller configuration. See [Policy controller advanced configuration](#) .
- Learn about cleaning up resources and other topics in the [Cleaning up resources that are created by policies](#) documentation.
- Refer to [Policy Generator](#) .
- Learn about how to create and customize policies, see [Manage security policies](#) .
- See [Support for templates in configuration policies](#) .

2.3.2. Certificate policy controller

Certificate policy controller can be used to detect certificates that are close to expiring, time durations (hours) that are too long, or contain DNS names that fail to match specified patterns. The certificate policy is provided in the **policy-templates** field of the policy on the hub cluster and is propagated to the selected managed clusters by the governance framework. See the [Policy overview](#) documentation for more details on the hub cluster policy.

Configure and customize the certificate policy controller by updating the following parameters in your controller policy:

- **minimumDuration**
- **minimumCADuration**
- **maximumDuration**
- **maximumCADuration**
- **allowedSANPattern**
- **disallowedSANPattern**

Your policy might become non-compliant due to either of the following scenarios:

- When a certificate expires in less than the minimum duration of time or exceeds the maximum time.
- When DNS names fail to match the specified pattern.

The certificate policy controller is created on your managed cluster. The controller communicates with the local Kubernetes API server to get the list of secrets that contain certificates and determine all non-compliant certificates.

Certificate policy controller does not support the **enforce** feature.

Note: The certificate policy controller automatically looks for a certificate in a secret in only the **tls.crt** key. If a secret is stored under a different key, add a label named **certificate_key_name** with a value set to the key to let the certificate policy controller know to look in a different key. For example, if a secret contains a certificate stored in the key named **sensor-cert.pem**, add the following label to the secret: **certificate_key_name: sensor-cert.pem**.

2.3.2.1. Certificate policy controller YAML structure

View the following example of a certificate policy and review the element in the YAML table:

```
apiVersion: policy.open-cluster-management.io/v1
kind: CertificatePolicy
metadata:
  name: certificate-policy-example
spec:
  namespaceSelector:
    include: ["default"]
    exclude: []
    matchExpressions: []
    matchLabels: {}
  labelSelector:
    myLabelKey: myLabelValue
  remediationAction:
  severity:
  minimumDuration:
  minimumCADuration:
  maximumDuration:
  maximumCADuration:
  allowedSANPattern:
  disallowedSANPattern:
```

2.3.2.1.1. Certificate policy controller YAML table

Table 2.3. Parameter table

| Field | Optional or required | Description |
|----------------------|----------------------|---|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to CertificatePolicy to indicate the type of policy. |
| metadata.name | Required | The name to identify the policy. |

| Field | Optional or required | Description |
|-------------------------------|----------------------|--|
| metadata.labels | Optional | In a certificate policy, the category=system-and-information-integrity label categorizes the policy and facilitates querying the certificate policies. If there is a different value for the category key in your certificate policy, the value is overridden by the certificate controller. |
| spec.namespaceSelector | Required | Determines namespaces in the managed cluster where secrets are monitored. The include and exclude parameters accept file path expressions to include and exclude namespaces by name. The matchExpressions and matchLabels parameters specify namespaces to be included by label. See the Kubernetes labels and selectors documentation. The resulting list is compiled by using the intersection of results from all parameters. Note: If the namespaceSelector for the certificate policy controller does not match any namespace, the policy is considered compliant. |
| spec.labelSelector | Optional | Specifies identifying attributes of objects. See the Kubernetes labels and selectors documentation. |
| spec.remediationAction | Required | Specifies the remediation of your policy. Set the parameter value to inform . Certificate policy controller only supports inform feature. |
| spec.severity | Optional | Informs the user of the severity when the policy is non-compliant. Use the following parameter values: low , medium , high , or critical . |

| Field | Optional or required | Description |
|-------------------------------|----------------------|---|
| spec.minimumDuration | Required | When a value is not specified, the default value is 100h . This parameter specifies the smallest duration (in hours) before a certificate is considered non-compliant. The parameter value uses the time duration format from Golang. See Golang Parse Duration for more information. |
| spec.minimumCADuration | Optional | Set a value to identify signing certificates that might expire soon with a different value from other certificates. If the parameter value is not specified, the CA certificate expiration is the value used for the minimumDuration . See Golang Parse Duration for more information. |
| spec.maximumDuration | Optional | Set a value to identify certificates that have been created with a duration that exceeds your desired limit. The parameter uses the time duration format from Golang. See Golang Parse Duration for more information. |
| spec.maximumCADuration | Optional | Set a value to identify signing certificates that have been created with a duration that exceeds your defined limit. The parameter uses the time duration format from Golang. See Golang Parse Duration for more information. |
| spec.allowedSANPattern | Optional | A regular expression that must match every SAN entry that you have defined in your certificates. This parameter checks DNS names against patterns. See the Golang Regular Expression syntax for more information. |

| Field | Optional or required | Description |
|----------------------------------|----------------------|--|
| spec.disallowedSANPattern | Optional | <p>A regular expression that must not match any SAN entries you have defined in your certificates. This parameter checks DNS names against patterns.</p> <p>Note: To detect wild-card certificate, use the following SAN pattern: disallowedSANPattern: "[*]"</p> <p>See the Golang Regular Expression syntax for more information.</p> |

2.3.2.2. Certificate policy sample

When your certificate policy controller is created on your hub cluster, a replicated policy is created on your managed cluster. See [policy-certificate.yaml](#) to view the certificate policy sample.

Learn how to manage a certificate policy, see [Managing security policies](#) for more details. Refer to [Policy controllers](#) for more topics.

2.3.3. IAM policy controller

The Identity and Access Management (IAM) policy controller can be used to receive notifications about IAM policies that are non-compliant. The compliance check is based on the parameters that you configure in the IAM policy. The IAM policy is provided in the **policy-templates** field of the policy on the hub cluster and is propagated to the selected managed clusters by the governance framework. See the [Policy YAML structure](#) documentation for more details on the hub cluster policy.

The IAM policy controller monitors for the desired maximum number of users with a particular cluster role (i.e. **ClusterRole**) in your cluster. The default cluster role to monitor is **cluster-admin**. The IAM policy controller communicates with the local Kubernetes API server.

The IAM policy controller runs on your managed cluster. View the following sections to learn more:

- [IAM policy YAML structure](#)
- [IAM policy YAML table](#)
- [IAM policy sample](#)

2.3.3.1. IAM policy YAML structure

View the following example of an IAM policy and review the parameters in the YAML table:

```
apiVersion: policy.open-cluster-management.io/v1
kind: iamPolicy
metadata:
  name:
spec:
```

```

clusterRole:
severity:
remediationAction:
maxClusterRoleBindingUsers:
ignoreClusterRoleBindings:

```

2.3.3.2. IAM policy YAML table

View the following parameter table for descriptions:

Table 2.4. Parameter table

| Field | Optional or required | Description |
|-------------------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| spec.clusterRole | Optional | The cluster role (i.e. ClusterRole) to monitor. This defaults to cluster-admin if not specified. |
| spec.severity | Optional | Informs the user of the severity when the policy is non-compliant. Use the following parameter values: low , medium , high , or critical . |
| spec.remediationAction | Optional | Specifies the remediation of your policy. Enter inform . The IAM policy controller only supports the inform feature. |

| Field | Optional or required | Description |
|--|----------------------|--|
| spec.ignoreClusterRoleBindings | Optional | A list of regular expression (regex) values that indicate which cluster role binding names to ignore. These regular expression values must follow Go regexp syntax . By default, all cluster role bindings that have a name that starts with system: are ignored. It is recommended to set this to a stricter value. To not ignore any cluster role binding names, set the list to a single value of <code>.^</code> or some other regular expression that never matches. |
| spec.maxClusterRoleBindingUsers | Required | Maximum number of IAM role bindings that are available before a policy is considered non-compliant. |

2.3.3.3. IAM policy sample

See [policy-limitclusteradmin.yaml](#) to view the IAM policy sample. See [Managing security policies](#) for more information. Refer to [Policy controllers](#) for more topics.

2.3.4. Policy set controller

The policy set controller aggregates the policy status scoped to policies that are defined in the same namespace. Create a policy set (**PolicySet**) to group policies that are in the same namespace. All policies in the **PolicySet** are placed together in a selected cluster by creating a **PlacementBinding** to bind the **PolicySet** and **Placement**. The policy set is deployed to the hub cluster.

Additionally, when a policy is a part of multiple policy sets, existing and new **Placement** resources remain in the policy. When a user removes a policy from the policy set, the policy is not applied to the cluster that is selected in the policy set, but the placements remain. The policy set controller only checks for violations in clusters that include the policy set placement.

Note: The Red Hat Advanced Cluster Management hardening sample policy set uses cluster placement. If you use cluster placement, bind the namespace containing the policy to the managed cluster set. See [Deploying policies to your cluster](#) for more details on using cluster placement.

Learn more details about the policy set structure in the following sections:

- [Policy set controller YAML structure](#)
- [Policy set controller YAML table](#)
- [Policy set sample](#)

2.3.4.1. Policy set YAML structure

Your policy set might resemble the following YAML file:

```

apiVersion: policy.open-cluster-management.io/v1beta1
kind: PolicySet
metadata:
  name: demo-policyset
spec:
  policies:
  - policy-demo

---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: demo-policyset-pb
placementRef:
  apiGroup: apps.open-cluster-management.io
  kind: PlacementRule
  name: demo-policyset-pr
subjects:
- apiGroup: policy.open-cluster-management.io
  kind: PolicySet
  name: demo-policyset

---
apiVersion: apps.open-cluster-management.io
kind: PlacementRule
metadata:
  name: demo-policyset-pr
spec:
  clusterConditions:pagewidth:
  - status: "True"
    type: ManagedCLusterConditionAvailable
  clusterSelectors:
  matchExpressions:
  - key: name
    operator: In
    values:
    - local-cluster

```

2.3.4.2. Policy set table

View the following parameter table for descriptions:

Table 2.5. Parameter table

| Field | Optional or required | Description |
|-------------------|----------------------|---|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1beta1 . |
| kind | Required | Set the value to PolicySet to indicate the type of policy. |

| Field | Optional or required | Description |
|----------------------|----------------------|---|
| metadata.name | Required | The name for identifying the policy resource. |
| spec | Required | Add configuration details for your policy. |
| spec.policies | Optional | The list of policies that you want to group together in the policy set. |

2.3.4.3. Policy set sample

```

apiVersion: policy.open-cluster-management.io/v1beta1
kind: PolicySet
metadata:
  name: pci
  namespace: default
spec:
  description: Policies for PCI compliance
  policies:
    - policy-pod
    - policy-namespace
status:
  compliant: NonCompliant
  placement:
    - placementBinding: binding1
      placementRule: placement1
    policySet: policyset-ps

```

See the *Creating policy sets* section in the [Managing security policies](#) topic. Also view the stable **PolicySets**, which require the Policy Generator for deployment, [PolicySets-- Stable](#). See the [Policy Generator](#) documentation.

2.4. POLICY CONTROLLER ADVANCED CONFIGURATION

You can customize policy controller configurations on your managed clusters using the **ManagedClusterAddOn** custom resources. The following **ManagedClusterAddOns** configure the policy framework, **governance-policy-framework**, **config-policy-controller**, **cert-policy-controller**, and **iam-policy-controller**.

Required access: Cluster administrator

2.4.1. Configure the concurrency

You can configure the concurrency of the configuration policy controller for each managed cluster to change how many configuration policies it can evaluate at the same time. To change the default value of **2**, set the **policy-evaluation-concurrency** annotation with a non-zero integer within quotes. You can set the value on the **ManagedClusterAddOn** object called **config-policy-controller** in the managed cluster namespace of the hub cluster.

Note: Higher concurrency values increase CPU and memory utilization on the **config-policy-controller** pod, Kubernetes API server, and OpenShift API server.

In the following YAML example, concurrency is set to **5** on the managed cluster called **cluster1**:

```
apiVersion: addon.open-cluster-management.io/v1alpha1
kind: ManagedClusterAddOn
metadata:
  name: config-policy-controller
  namespace: cluster1
  annotations:
    policy-evaluation-concurrency: "5"
spec:
  installNamespace: open-cluster-management-agent-addon
```

2.4.2. Configure debug log

Configure debug logs for each policy controller. To receive the second level of debugging information for the Kubernetes configuration controller, add the **log-level** annotation with the value of **2** to the **ManagedClusterAddOn** custom resource. By default, the **log-level** is set to **0**, which means you receive informative messages. View the following example:

```
apiVersion: addon.open-cluster-management.io/v1alpha1
kind: ManagedClusterAddOn
metadata:
  name: config-policy-controller
  namespace: cluster1
  annotations:
    log-level: "2"
spec:
  installNamespace: open-cluster-management-agent-addon
```

2.4.3. Governance metric

The policy framework exposes metrics that show policy distribution and compliance. Use the **policy_governance_info** metric on the hub cluster to view trends and analyze any policy failures. See the following topics for an overview of metrics:

2.4.3.1. Metric: `policy_governance_info`

The **policy_governance_info** is collected by OpenShift Container Platform monitoring, and some aggregate data is collected by Red Hat Advanced Cluster Management observability, if it is enabled.

Note: If observability is enabled, you can enter a query for the metric from the Grafana *Explore* page.

When you create a policy, you are creating a *root* policy. The framework watches for root policies as well as **PlacementRules** and **PlacementBindings**, to determine where to create *propagated* policies in order to distribute the policy to managed clusters. For both root and propagated policies, a metric of **0** is recorded if the policy is compliant, and **1** if it is non-compliant.

The **policy_governance_info** metric uses the following labels:

- **type:** The label values are **root** or **propagated**.

- **policy**: The name of the associated root policy.
- **policy_namespace**: The namespace on the hub cluster where the root policy was defined.
- **cluster_namespace**: The namespace for the cluster where the policy is distributed.

These labels and values enable queries that can show us many things happening in the cluster that might be difficult to track.

Note: If the metrics are not needed, and there are any concerns about performance or security, this feature can be disabled. Set the **DISABLE_REPORT_METRICS** environment variable to **true** in the propagator deployment. You can also add **policy_governance_info** metric to the observability allowlist as a custom metric. See [Adding custom metrics](#) for more details.

2.4.3.2. Metric: *config_policies_evaluation_duration_seconds*

The **config_policies_evaluation_duration_seconds** histogram tracks the number of seconds it takes to process all configuration policies that are ready to be evaluated on the cluster. Use the following metrics to query the histogram:

- **config_policies_evaluation_duration_seconds_bucket**: The buckets are cumulative and represent seconds with the following possible entries: 1, 3, 9, 10.5, 15, 30, 60, 90, 120, 180, 300, 450, 600, and greater.
- **config_policies_evaluation_duration_seconds_count**: The count of all events.
- **config_policies_evaluation_duration_seconds_sum**: The sum of all values.

Use the **config_policies_evaluation_duration_seconds** metric to determine if the **ConfigurationPolicy evaluationInterval** setting needs to be changed for resource intensive policies that do not need frequent evaluation. You can also increase the concurrency at the cost of higher resource utilization on the Kubernetes API server. See *Configure the concurrency* section for more details.

To receive information about the time used to evaluate configuration policies, perform a Prometheus query that resembles the following expression:

```
rate(config_policies_evaluation_duration_seconds_sum[10m])/rate
(config_policies_evaluation_duration_seconds_count[10m])
```

The **config-policy-controller** pod running on managed clusters in the **open-cluster-management-agent-addon** namespace calculates the metric. The **config-policy-controller** does not send the metric to observability by default.

2.4.4. Verify configuration changes

When the new configuration is applied by the controller, the **ManifestApplied** parameter is updated in the **ManagedClusterAddOn**. That condition timestamp can be used to verify the configuration correctly. For example, this command can verify when the **cert-policy-controller** on the **local-cluster** was updated:

```
oc get -n local-cluster managedclusteraddon cert-policy-controller | grep -B4 'type: ManifestApplied'
```

You might receive the following output:

```
- lastTransitionTime: "2023-01-26T15:42:22Z"
```

```

message: manifests of addon are applied successfully
reason: AddonManifestApplied
status: "True"
type: ManifestApplied

```

Return to the [Governance](#) page for more topics.

2.5. SUPPORTED POLICIES

View the supported policies to learn how to define rules, processes, and controls on the hub cluster when you create and manage policies in Red Hat Advanced Cluster Management for Kubernetes.

2.5.1. Table of sample configuration policies

View the following sample configuration policies:

Table 2.6. Table list of configuration policies

| Policy sample | Description |
|--|--|
| Namespace policy | Ensure consistent environment isolation and naming with Namespaces. See the Kubernetes Namespace documentation . |
| Pod policy | Ensure cluster workload configuration. See the Kubernetes Pod documentation . |
| Memory usage policy | Limit workload resource usage by using Limit Ranges. See the Limit Range documentation . |
| Pod security policy (Deprecated) | Ensure consistent workload security. See the Kubernetes Pod security policy documentation . |
| Role policy Role binding policy | Manage role permissions and bindings by using Roles and Role Bindings. See the Kubernetes RBAC documentation . |
| Security content constraints (SCC) policy | Manage workload permissions with Security Context Constraints. See the Openshift Security Context Constraints documentation . |
| ETCD encryption policy | Ensure data security with etcd encryption. See the Openshift etcd encryption documentation . |
| Compliance operator policy | Deploy the Compliance Operator to scan and enforce the compliance state of clusters leveraging OpenSCAP. See the Openshift Compliance Operator documentation . |

| Policy sample | Description |
|--|---|
| Compliance operator E8 scan | After applying the Compliance operator policy, deploy an Essential 8 (E8) scan to check for compliance with E8 security profiles. See the Openshift Compliance Operator documentation . |
| Compliance operator CIS scan | After applying the Compliance operator policy, deploy a Center for Internet Security (CIS) scan to check for compliance with CIS security profiles. See the Openshift Compliance Operator documentation . |
| Image vulnerability policy | Deploy the Container Security Operator and detect known image vulnerabilities in pods running on the cluster. See the Container Security Operator GitHub . |
| Gatekeeper operator deployment | Gatekeeper is an admission webhook that enforces custom resource definition-based policies that are run by the Open Policy Agent (OPA) policy engine. See the Gatekeeper documentation. |
| Gatekeeper compliance policy | After deploying Gatekeeper to the clusters, deploy this sample Gatekeeper policy that ensures namespaces that are created on the cluster are labeled as specified. |

2.5.2. Support matrix for out-of-box policies

Table 2.7. Support matrix

| Policy | Red Hat OpenShift Container Platform 3.11 | Red Hat OpenShift Container Platform 4 |
|---|---|--|
| Memory usage policy | x | x |
| Namespace policy | x | x |
| Image vulnerability policy | x | x |
| Pod policy | x | x |
| Pod security policy (deprecated) | | |
| Role policy | x | x |
| Role binding policy | x | x |
| Security Context Constraints policy (SCC) | x | x |

| Policy | Red Hat OpenShift Container Platform 3.11 | Red Hat OpenShift Container Platform 4 |
|----------------------------|---|--|
| ETCD encryption policy | | x |
| Gatekeeper policy | | x |
| Compliance operator policy | | x |
| E8 scan policy | | x |
| OpenShift CIS scan policy | | x |
| Policy set | | x |

View the following policy samples to view how specific policies are applied:

- [Image vulnerability policy](#)
- [Memory usage policy](#)
- [Namespace policy](#)
- [Pod policy](#)
- [Pod security policy](#)
- [Role policy](#)
- [Role binding policy](#)
- [Security context constraints policy](#)
- [ETCD encryption policy](#)
- [Compliance operator policy](#)
- [E8 scan policy](#)
- [OpenShift CIS scan policy](#)
- [Policy set controller](#)

Refer to [Governance](#) for more topics.

2.5.3. Memory usage policy

The Kubernetes configuration policy controller monitors the status of the memory usage policy. Use the memory usage policy to limit or restrict your memory and compute usage. For more information, see *Limit Ranges* in the [Kubernetes documentation](#).

Learn more details about the memory usage policy structure in the following sections:

- [Memory usage policy YAML structure](#)
- [Memory usage policy table](#)
- [Memory usage policy sample](#)

2.5.3.1. Memory usage policy YAML structure

Your memory usage policy might resemble the following YAML file:

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        namespaceSelector:
          exclude:
          include:
          matchLabels:
          matchExpressions:
        object-templates:
          - complianceType: mustonlyhave
            objectDefinition:
              apiVersion: v1
              kind: LimitRange
              metadata:
                name:
              spec:
                limits:
                  - default:
                      memory:
                    defaultRequest:
                      memory:
                type:
  ...

```

2.5.3.2. Memory usage policy table

Table 2.8. Parameter table

| Field | Optional or required | Description |
|---|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because the value overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.3.3. Memory usage policy sample

See the [policy-limitmemory.yaml](#) to view a sample of the policy. See [Managing security policies](#) for more details. Refer to the [Policy overview](#) documentation, and to [Kubernetes configuration policy controller](#) to view other configuration policies that are monitored by the controller.

2.5.4. Namespace policy

The Kubernetes configuration policy controller monitors the status of your namespace policy. Apply the namespace policy to define specific rules for your namespace.

Learn more details about the namespace policy structure in the following sections:

- [Namespace policy YAML structure](#)
- [Namespace policy table](#)
- [Namespace policy sample](#)

2.5.4.1. Namespace policy YAML structure

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        object-templates:
          - complianceType:
            objectDefinition:
              kind: Namespace
              apiVersion: v1
              metadata:
                name:
            ...

```

2.5.4.2. Namespace policy YAML table

| Field | Optional or required | Description |
|-------------------------------|----------------------|---|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because it overrides any values provided in spec.policy-templates . |

| Field | Optional or required | Description |
|---|----------------------|--|
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.4.3. Namespace policy sample

See [policy-namespace.yaml](#) to view the policy sample.

See [Managing security policies](#) for more details. Refer to [Policy overview](#) documentation, and to the [Kubernetes configuration policy controller](#) to learn about other configuration policies.

2.5.5. Image vulnerability policy

Apply the image vulnerability policy to detect if container images have vulnerabilities by leveraging the Container Security Operator. The policy installs the Container Security Operator on your managed cluster if it is not installed.

The image vulnerability policy is checked by the Kubernetes configuration policy controller. For more information about the Security Operator, see the *Container Security Operator* from the [Quay repository](#).

Notes:

- Image vulnerability policy is not functional during a disconnected installation.
- The [Image vulnerability policy](#) is not supported on the IBM Power and IBM Z architectures. It relies on the [Quay Container Security Operator](#). There are no **ppc64le** or **s390x** images in the [container-security-operator registry](#).

View the following sections to learn more:

- [Image vulnerability policy YAML structure](#)
- [Image vulnerability policy sample](#)

2.5.5.1. Image vulnerability policy YAML structure

When you create the container security operator policy, it involves the following policies:

- A policy that creates the subscription (**container-security-operator**) to reference the name and channel. This configuration policy must have **spec.remediationAction** set to **enforce** to create the resources. The subscription pulls the profile, as a container, that the subscription supports. View the following example:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-imagemanifestvuln-example-sub
spec:
  remediationAction: enforce # will be overridden by remediationAction in parent policy
  severity: high
  object-templates:
    - complianceType: musthave
      objectDefinition:
        apiVersion: operators.coreos.com/v1alpha1
        kind: Subscription
        metadata:
          name: container-security-operator
          namespace: openshift-operators
        spec:
          # channel: quay-v3.3 # specify a specific channel if desired
          installPlanApproval: Automatic
          name: container-security-operator
          source: redhat-operators
          sourceNamespace: openshift-marketplace

```

- An **inform** configuration policy to audit the **ClusterServiceVersion** to ensure that the container security operator installation succeeded. View the following example:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-imagemanifestvuln-status
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: high
  object-templates:
    - complianceType: musthave
      objectDefinition:
        apiVersion: operators.coreos.com/v1alpha1
        kind: ClusterServiceVersion
        metadata:
          namespace: openshift-operators
        spec:
          displayName: Red Hat Quay Container Security Operator
        status:
          phase: Succeeded # check the CSV status to determine if operator is running or not

```

- An **inform** configuration policy to audit whether any **ImageManifestVuln** objects were created by the image vulnerability scans. View the following example:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-imagemanifestvuln-example-imv
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: high
  namespaceSelector:

```

```

exclude: ["kube-*"]
include: ["*"]
object-templates:
- complianceType: mustnothave # mustnothave any ImageManifestVuln object
  objectDefinition:
    apiVersion: secscan.quay.redhat.com/v1alpha1
    kind: ImageManifestVuln # checking for a Kind

```

2.5.5.2. Image vulnerability policy sample

See [policy-imagemanifestvuln.yaml](#). See [Managing security policies](#) for more information. Refer to [Kubernetes configuration policy controller](#) to view other configuration policies that are monitored by the configuration controller.

2.5.6. Pod policy

The Kubernetes configuration policy controller monitors the status of your pod policies. Apply the pod policy to define the container rules for your pods. A pod must exist in your cluster to use this information.

Learn more details about the pod policy structure in the following sections:

- [Pod policy YAML structure](#)
- [Pod policy table](#)
- [Pod policy sample](#)

2.5.6.1. Pod policy YAML structure

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
  - objectDefinition:
    apiVersion: policy.open-cluster-management.io/v1
    kind: ConfigurationPolicy
    metadata:
      name:
    spec:
      remediationAction:
      severity:
      namespaceSelector:
        exclude:
        include:
        matchLabels:

```

```

matchExpressions:
object-templates:
- complianceType:
objectDefinition:
  apiVersion: v1
  kind: Pod
  metadata:
    name:
  spec:
    containers:
    - image:
      name:
  ...

```

2.5.6.2. Pod policy table

Table 2.9. Parameter table

| Field | Optional or required | Description |
|---|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because the value overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.6.3. Pod policy sample

See [policy-pod.yaml](#) to view the policy sample.

Refer to [Kubernetes configuration policy controller](#) to view other configuration policies that are monitored by the configuration controller, and see the [Policy overview documentation](#) to see a full description of the policy YAML structure and additional fields. Return to [Managing configuration policies](#) documentation to manage other policies.

2.5.7. Pod security policy (Deprecated)

The Kubernetes configuration policy controller monitors the status of the pod security policy. Apply a pod security policy to secure pods and containers.

Learn more details about the pod security policy structure in the following sections:

- [Pod security policy YAML structure](#)
- [Pod security policy table](#)
- [Pod security policy sample](#)

2.5.7.1. Pod security policy YAML structure

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
        apiVersion: policy.open-cluster-management.io/v1
        kind: ConfigurationPolicy
        metadata:
          name:
        spec:
          remediationAction:
          severity:
          namespaceSelector:
            exclude:
            include:
          matchLabels:
          matchExpressions:
        object-templates:
          - complianceType:
              objectDefinition:
                apiVersion: policy/v1beta1
                kind: PodSecurityPolicy
                metadata:
                  name:
                  annotations:
                    seccomp.security.alpha.kubernetes.io/allowedProfileNames:

```

```

spec:
  privileged:
  allowPrivilegeEscalation:
  allowedCapabilities:
  volumes:
  hostNetwork:
  hostPorts:
  hostIPC:
  hostPID:
  runAsUser:
  seLinux:
  supplementalGroups:
  fsGroup:
  ...

```

2.5.7.2. Pod security policy table

Table 2.10. Parameter table

| Field | Optional or required | Description |
|---|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because the value overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.7.3. Pod security policy sample

The support of pod security policies is removed from OpenShift Container Platform 4.12 and later, and from Kubernetes v1.25 and later. If you apply a **PodSecurityPolicy** resource, you might receive the following non-compliant message:

```
violation - couldn't find mapping resource with kind PodSecurityPolicy, please check if you have CRD
deployed
```

- For more information including the deprecation notice, see *Pod Security Policies* in the [Kubernetes documentation](#).
- See [policy-psp.yaml](#) to view the sample policy. View [Managing configuration policies](#) for more information.
- Refer to the [Policy overview](#) documentation for a full description of the policy YAML structure, and [Kubernetes configuration policy controller](#) to view other configuration policies that are monitored by the controller.

2.5.8. Role policy

The Kubernetes configuration policy controller monitors the status of role policies. Define roles in the **object-template** to set rules and permissions for specific roles in your cluster.

Learn more details about the role policy structure in the following sections:

- [Role policy YAML structure](#)
- [Role policy table](#)
- [Role policy sample](#)

2.5.8.1. Role policy YAML structure

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        namespaceSelector:
          exclude:
```

```

include:
  matchLabels:
  matchExpressions:
object-templates:
- complianceType:
  objectDefinition:
    apiVersion: rbac.authorization.k8s.io/v1
    kind: Role
    metadata:
      name:
    rules:
      - apiGroups:
        resources:
        verbs:
      ...
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-policy-role
  namespace:
placementRef:
  name: placement-policy-role
  kind: PlacementRule
  apiGroup: apps.open-cluster-management.io
subjects:
- name: policy-role
  kind: Policy
  apiGroup: policy.open-cluster-management.io
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-policy-role
  namespace:
spec:
  clusterConditions:
    - type: ManagedClusterConditionAvailable
      status: "True"
  clusterSelector:
    matchExpressions:
      []
  ...

```

2.5.8.2. Role policy table

Table 2.11. Parameter table

| Field | Optional or required | Description |
|-------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |

| Field | Optional or required | Description |
|---|----------------------|--|
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because the value overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.8.3. Role policy sample

Apply a role policy to set rules and permissions for specific roles in your cluster. For more information on roles, see [Role-based access control](#). View a sample of a role policy, see [policy-role.yaml](#).

To learn how to manage role policies, refer to [Managing configuration policies](#) for more information. See the [Kubernetes configuration policy controller](#) to view other configuration policies that are monitored the controller.

2.5.9. Role binding policy

The Kubernetes configuration policy controller monitors the status of your role binding policy. Apply a role binding policy to bind a policy to a namespace in your managed cluster.

Learn more details about the namespace policy structure in the following sections:

- [Role binding policy YAML structure](#)
- [Role binding policy table](#)
- [Role binding policy sample](#)

2.5.9.1. Role binding policy YAML structure

```
apiVersion: policy.open-cluster-management.io/v1
```

```

kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        namespaceSelector:
          exclude:
          include:
          matchLabels:
          matchExpressions:
        object-templates:
          - complianceType:
            objectDefinition:
              kind: RoleBinding # role binding must exist
              apiVersion: rbac.authorization.k8s.io/v1
              metadata:
                name:
              subjects:
                - kind:
                  name:
                  apiGroup:
              roleRef:
                kind:
                name:
                apiGroup:
            ...

```

2.5.9.2. Role binding policy table

| Field | Optional or required | Description |
|-------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |

| Field | Optional or required | Description |
|---|----------------------|---|
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional since it overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.9.3. Role binding policy sample

See [policy-rolebinding.yaml](#) to view the policy sample. For a full description of the policy YAML structure and additional fields, see the [Policy overview documentation](#). Refer to [Kubernetes configuration policy controller](#) documentation to learn about other configuration policies.

2.5.10. Security Context Constraints policy

The Kubernetes configuration policy controller monitors the status of your Security Context Constraints (SCC) policy. Apply an Security Context Constraints (SCC) policy to control permissions for pods by defining conditions in the policy.

Learn more details about SCC policies in the following sections:

- [SCC policy YAML structure](#)
- [SCC policy table](#)
- [SCC policy sample](#)

2.5.10.1. SCC policy YAML structure

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:

```

```

policy.open-cluster-management.io/standards:
policy.open-cluster-management.io/categories:
policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        namespaceSelector:
          exclude:
          include:
          matchLabels:
          matchExpressions:
        object-templates:
        - complianceType:
            objectDefinition:
              apiVersion: security.openshift.io/v1
              kind: SecurityContextConstraints
              metadata:
                name:
              allowHostDirVolumePlugin:
              allowHostIPC:
              allowHostNetwork:
              allowHostPID:
              allowHostPorts:
              allowPrivilegeEscalation:
              allowPrivilegedContainer:
              fsGroup:
              readOnlyRootFilesystem:
              requiredDropCapabilities:
              runAsUser:
              seLinuxContext:
              supplementalGroups:
              users:
              volumes:
            ...

```

2.5.10.2. SCC policy table

| Field | Optional or required | Description |
|-------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |

| Field | Optional or required | Description |
|---|----------------------|---|
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional since it overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

For explanations on the contents of a SCC policy, see [Managing Security Context Constraints](#) from the OpenShift Container Platform documentation.

2.5.10.3. SCC policy sample

Apply a Security context constraints (SCC) policy to control permissions for pods by defining conditions in the policy. For more information see, [Managing Security Context Constraints \(SCC\)](#).

See [policy-scc.yaml](#) to view the policy sample. For a full description of the policy YAML structure and additional fields, see the [Policy overview](#) documentation. Refer to [Kubernetes configuration policy controller](#) documentation to learn about other configuration policies.

2.5.11. ETCD encryption policy

Apply the **etcd-encryption** policy to detect, or enable encryption of sensitive data in the ETCD data-store. The Kubernetes configuration policy controller monitors the status of the **etcd-encryption** policy. For more information, see [Encrypting etcd data](#) in the OpenShift Container Platform documentation.

Note: The ETCD encryption policy only supports Red Hat OpenShift Container Platform 4 and later.

Learn more details about the **etcd-encryption** policy structure in the following sections:

- [ETCD encryption policy YAML structure](#)
- [ETCD encryption policy table](#)
- [ETCD encryption policy sample](#)

2.5.11.1. ETCD encryption policy YAML structure

Your **etcd-encryption** policy might resemble the following YAML file:

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name:
  namespace:
  annotations:
    policy.open-cluster-management.io/standards:
    policy.open-cluster-management.io/categories:
    policy.open-cluster-management.io/controls:
spec:
  remediationAction:
  disabled:
  policy-templates:
    - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name:
      spec:
        remediationAction:
        severity:
        object-templates:
          - complianceType:
            objectDefinition:
              apiVersion: config.openshift.io/v1
              kind: APIServer
              metadata:
                name:
              spec:
                encryption:
                ...

```

2.5.11.2. ETCD encryption policy table

Table 2.12. Parameter table

| Field | Optional or required | Description |
|---------------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to Policy to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| metadata.namespace | Required | The namespace of the policy. |

| Field | Optional or required | Description |
|---|----------------------|---|
| spec.remediationAction | Optional | Specifies the remediation of your policy. The parameter values are enforce and inform . This value is optional because it overrides any values provided in spec.policy-templates . |
| spec.disabled | Required | Set the value to true or false . The disabled parameter provides the ability to enable and disable your policies. |
| spec.policy-templates[].objectDefinition | Required | Used to list configuration policies containing Kubernetes objects that must be evaluated or applied to the managed clusters. |

2.5.11.3. ETCD encryption policy sample

See [policy-etcdencryption.yaml](#) for the policy sample. See the [Policy overview](#) documentation and the [Kubernetes configuration policy controller](#) to view additional details on policy and configuration policy fields.

2.5.12. Compliance Operator policy

You can use the Compliance Operator to automate the inspection of numerous technical implementations and compare those against certain aspects of industry standards, benchmarks, and baselines. The Compliance Operator is not an auditor. To be compliant or certified with these various standards, you need to engage an authorized auditor such as a Qualified Security Assessor (QSA), Joint Authorization Board (JAB), or other industry recognized regulatory authority to assess your environment.

Recommendations that are generated from the Compliance Operator are based on generally available information and practices regarding such standards, and might assist you with remediations, but actual compliance is your responsibility. Work with an authorized auditor to achieve compliance with a standard.

For the latest updates, see the [Compliance Operator release notes](#).

2.5.12.1. Compliance Operator policy overview

You can install the Compliance Operator on your managed cluster by using the Compliance Operator policy. The Compliance Operator policy is created as a Kubernetes configuration policy in Red Hat Advanced Cluster Management. OpenShift Container Platform supports the Compliance Operator policy.

Note: The [Compliance operator policy](#) relies on the OpenShift Container Platform Compliance Operator, which is not supported on the IBM Power or IBM Z architectures. See [Understanding the Compliance Operator](#) in the OpenShift Container Platform documentation for more information about the Compliance Operator.

2.5.12.2. Compliance operator resources

When you create a Compliance Operator policy, the following resources are created:

- A Compliance Operator namespace (**openshift-compliance**) for the operator installation:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: comp-operator-ns
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: high
  object-templates:
  - complianceType: musthave
    objectDefinition:
      apiVersion: v1
      kind: Namespace
      metadata:
        name: openshift-compliance
```

- An operator group (**compliance-operator**) to specify the target namespace:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: comp-operator-operator-group
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: high
  object-templates:
  - complianceType: musthave
    objectDefinition:
      apiVersion: operators.coreos.com/v1
      kind: OperatorGroup
      metadata:
        name: compliance-operator
        namespace: openshift-compliance
      spec:
        targetNamespaces:
        - openshift-compliance
```

- A subscription (**comp-operator-subscription**) to reference the name and channel. The subscription pulls the profile, as a container, that it supports:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: comp-operator-subscription
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: high
  object-templates:
  - complianceType: musthave
    objectDefinition:
```

```

apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: compliance-operator
  namespace: openshift-compliance
spec:
  channel: "4.7"
  installPlanApproval: Automatic
  name: compliance-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace

```

2.5.12.3. Additional resources

- For more information, see [Managing the Compliance Operator](#) in the OpenShift Container Platform documentation for more details.
- After you install the Compliance Operator policy, the following pods are created: **compliance-operator**, **ocp4**, and **rhcos4**. See a sample of the [policy-compliance-operator-install.yaml](#).
- You can also create and apply the E8 scan policy and OpenShift CIS scan policy, after you have installed the Compliance Operator. For more information, see [E8 scan policy](#) and [OpenShift CIS scan policy](#).
- To learn about managing Compliance Operator policies, see [Managing security policies](#) for more details. Refer to [Kubernetes configuration policy controller](#) for more topics about configuration policies.

2.5.13. E8 scan policy

An Essential 8 (E8) scan policy deploys a scan that checks the master and worker nodes for compliance with the E8 security profiles. You must install the compliance operator to apply the E8 scan policy.

The E8 scan policy is created as a Kubernetes configuration policy in Red Hat Advanced Cluster Management. OpenShift Container Platform 4.7 and 4.6, support the E8 scan policy. For more information, see *Understanding the Compliance Operator* in the [OpenShift Container Platform documentation](#) for more details.

2.5.13.1. E8 scan policy resources

When you create an E8 scan policy the following resources are created:

- A **ScanSettingBinding** resource (**e8**) to identify which profiles to scan:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-suite-e8
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: musthave # this template checks if scan has completed by checking the
      status field
    objectDefinition:

```

```

apiVersion: compliance.openshift.io/v1alpha1
kind: ScanSettingBinding
metadata:
  name: e8
  namespace: openshift-compliance
profiles:
- apiGroup: compliance.openshift.io/v1alpha1
  kind: Profile
  name: ocp4-e8
- apiGroup: compliance.openshift.io/v1alpha1
  kind: Profile
  name: rhcos4-e8
settingsRef:
  apiGroup: compliance.openshift.io/v1alpha1
  kind: ScanSetting
  name: default

```

- A **ComplianceSuite** resource (**compliance-suite-e8**) to verify if the scan is complete by checking the **status** field:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-suite-e8
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: musthave # this template checks if scan has completed by checking the
      status field
    objectDefinition:
      apiVersion: compliance.openshift.io/v1alpha1
      kind: ComplianceSuite
      metadata:
        name: e8
        namespace: openshift-compliance
      status:
        phase: DONE

```

- A **ComplianceCheckResult** resource (**compliance-suite-e8-results**) which reports the results of the scan suite by checking the **ComplianceCheckResult** custom resources (CR):

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-suite-e8-results
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: mustnothave # this template reports the results for scan suite: e8 by
      looking at ComplianceCheckResult CRs
    objectDefinition:
      apiVersion: compliance.openshift.io/v1alpha1
      kind: ComplianceCheckResult

```

```

metadata:
  namespace: openshift-compliance
labels:
  compliance.openshift.io/check-status: FAIL
  compliance.openshift.io/suite: e8

```

Note: Automatic remediation is supported. Set the remediation action to **enforce** to create **ScanSettingBinding** resource.

See a sample of the [policy-compliance-operator-e8-scan.yaml](#). See [Managing security policies](#) for more information. **Note:** After your E8 policy is deleted, it is removed from your target cluster or clusters.

2.5.14. OpenShift CIS scan policy

An OpenShift CIS scan policy deploys a scan that checks the master and worker nodes for compliance with the OpenShift CIS security benchmark. You must install the compliance operator to apply the OpenShift CIS policy.

The OpenShift CIS scan policy is created as a Kubernetes configuration policy in Red Hat Advanced Cluster Management. OpenShift Container Platform 4.9, 4.7, and 4.6, support the OpenShift CIS scan policy. For more information, see [Understanding the Compliance Operator](#) in the OpenShift Container Platform documentation for more details.

2.5.14.1. OpenShift CIS resources

When you create an OpenShift CIS scan policy the following resources are created:

- A **ScanSettingBinding** resource (**cis**) to identify which profiles to scan:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-cis-scan
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: musthave # this template creates ScanSettingBinding:cis
      objectDefinition:
        apiVersion: compliance.openshift.io/v1alpha1
        kind: ScanSettingBinding
        metadata:
          name: cis
          namespace: openshift-compliance
        profiles:
          - apiGroup: compliance.openshift.io/v1alpha1
            kind: Profile
            name: ocp4-cis
          - apiGroup: compliance.openshift.io/v1alpha1
            kind: Profile
            name: ocp4-cis-node
        settingsRef:

```

```

apiGroup: compliance.openshift.io/v1alpha1
kind: ScanSetting
name: default

```

- A **ComplianceSuite** resource (**compliance-suite-cis**) to verify if the scan is complete by checking the **status** field:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-suite-cis
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: musthave # this template checks if scan has completed by checking the
      status field
    objectDefinition:
      apiVersion: compliance.openshift.io/v1alpha1
      kind: ComplianceSuite
      metadata:
        name: cis
        namespace: openshift-compliance
      status:
        phase: DONE

```

- A **ComplianceCheckResult** resource (**compliance-suite-cis-results**) which reports the results of the scan suite by checking the **ComplianceCheckResult** custom resources (CR):

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: compliance-suite-cis-results
spec:
  remediationAction: inform
  severity: high
  object-templates:
    - complianceType: mustnothave # this template reports the results for scan suite: cis by
      looking at ComplianceCheckResult CRs
    objectDefinition:
      apiVersion: compliance.openshift.io/v1alpha1
      kind: ComplianceCheckResult
      metadata:
        namespace: openshift-compliance
        labels:
          compliance.openshift.io/check-status: FAIL
          compliance.openshift.io/suite: cis

```

See a sample of the [policy-compliance-operator-cis-scan.yaml](#) file. For more information on creating policies, see [Managing security policies](#).

2.6. MANAGE SECURITY POLICIES

Use the *Governance* dashboard to create, view, and manage your security policies and policy violations. You can create YAML files for your policies from the CLI and console.

2.6.1. Governance page

The following tabs are displayed on the *Governance* page:

- **Overview**
View the following summary cards from the *Overview* tab: *Policy set violations*, *Policy violations*, *Clusters*, *Categories*, *Controls*, and *Standards*.
- **Policy sets**
Create and manage hub cluster policy sets.
- **Policies**
Create and manage security policies. The table of policies lists the following details of a policy: *Name*, *Namespace*, *Status*, *Remediation*, *Policy set*, *Cluster violations*, *Source*, *Automation* and *Created*.

You can edit, enable or disable, set remediation to inform or enforce, or remove a policy by selecting the **Actions** icon. You can view the categories and standards of a specific policy by selecting the drop-down arrow to expand the row.

Complete bulk actions by selecting multiple policies and clicking the **Actions** button. You can also customize your policy table by clicking the **Filter** button.

When you select a policy in the table list, the following tabs of information are displayed from the console:

- *Details*: Select the *Details* tab to view policy details and placement details. In the *Placement* table, the *Compliance* column provides links to view the compliance of the clusters that are displayed.
- *Results*: Select the *Results* tab to view a table list of all clusters that are associated to the policy. From the *Message* column, click the **View details** link to view the template details, template YAML, and related resources. You can also view related resources. Click the **View history** link to view the violation message and a time of the last report.

2.6.2. Governance automation configuration

If there is a configured automation for a specific policy, you can select the automation to view more details. View the following descriptions of the schedule frequency options for your automation:

- *Manual run*: Manually set this automation to run once. After the automation runs, it is set to **disabled**. **Note**: You can only select *Manual run* mode when the schedule frequency is disabled.
- *Run once mode*: When a policy is violated, the automation runs one time. After the automation runs, it is set to **disabled**. After the automation is set to **disabled**, you must continue to run the automation manually. When you run *once mode*, the extra variable of **target_clusters** is automatically supplied with the list of clusters that violated the policy. The Ansible Automation Platform Job template must have **PROMPT ON LAUNCH** enabled for the **EXTRA VARIABLES** section (also known as **extra_vars**).
- *Run everyEvent mode*: When a policy is violated, the automation runs every time for each unique policy violation per managed cluster. Use the **DelayAfterRunSeconds** parameter to set the minimum seconds before an automation can be restarted on the same cluster. If the policy is

violated multiple times during the delay period and kept in the violated state, the automation runs one time after the delay period. The default is 0 seconds and is only applicable for the **everyEvent** mode. When you run **everyEvent** mode, the extra variable of **target_clusters** and Ansible Automation Platform Job template is the same as *once mode*.

- *Disable automation*: When the scheduled automation is set to **disabled**, the automation does not run until the setting is updated.

The following variables are automatically provided in the **extra_vars** of the Ansible Automation Platform Job:

- **policy_name**: The name of the non-compliant root policy that initiates the Ansible Automation Platform job on the hub cluster.
- **policy_namespace**: The namespace of the root policy.
- **hub_cluster**: The name of the hub cluster, which is determined by the value in the **clusters DNS** object.
- **policy_sets**: This parameter contains all associated policy set names of the root policy. If the policy is not within a policy set, the **policy_set** parameter is empty.
- **policy_violations**: This parameter contains a list of non-compliant cluster names, and the value is the policy **status** field for each non-compliant cluster.

Review the following topics to learn more about creating and updating your security policies:

- [Managing security policies](#)
- [Managing configuration policies](#)
- [Managing gatekeeper policies](#)
- [Configuring Ansible Automation Platform for governance](#)
- [Governance](#)

2.6.3. Configuring Ansible Automation Platform for governance

Red Hat Advanced Cluster Management for Kubernetes governance can be integrated with Red Hat Ansible Automation Platform to create policy violation automations. You can configure the automation from the Red Hat Advanced Cluster Management console.

- [Prerequisites](#)
- [Creating a policy violation automation from the console](#)
- [Creating a policy violation automation from the CLI](#)

2.6.3.1. Prerequisites

- Red Hat OpenShift Container Platform 4.5 or later
- You must have Ansible Automation Platform version 3.7.3 or a later version installed. It is best practice to install the latest supported version of Ansible Automation Platform. See [Red Hat Ansible Automation Platform documentation](#) for more details.

- Install the Ansible Automation Platform Resource Operator from the Operator Lifecycle Manager. In the *Update Channel* section, select **stable-2.x-cluster-scoped**. Select the **All namespaces on the cluster (default)** installation mode.
Note: Ensure that the Ansible Automation Platform job template is idempotent when you run it. If you do not have Ansible Automation Platform Resource Operator, you can find it from the Red Hat OpenShift Container Platform *OperatorHub* page.

For more information about installing and configuring Red Hat Ansible Automation Platform, see [Setting up Ansible tasks](#).

2.6.3.2. Creating a policy violation automation from the console

After you log into your Red Hat Advanced Cluster Management hub cluster, select **Governance** from the navigation menu, and then click on the *Policies* tab to view the policy tables.

Configure an automation for a specific policy by clicking **Configure** in the *Automation* column. You can create automation when the policy automation panel appears. From the *Ansible credential* section, click the drop-down menu to select an Ansible credential. If you need to add a credential, see [Managing credentials overview](#).

Note: This credential is copied to the same namespace as the policy. The credential is used by the **AnsibleJob** resource that is created to initiate the automation. Changes to the Ansible credential in the *Credentials* section of the console is automatically updated.

After a credential is selected, click the Ansible job drop-down list to select a job template. In the *Extra variables* section, add the parameter values from the **extra_vars** section of the **PolicyAutomation**. Select the frequency of the automation. You can select *Run once mode*, *Run everyEvent mode*, or *Disable automation*.

Save your policy violation automation by selecting **Submit**. When you select the *View Job* link from the Ansible job details side panel, the link directs you to the job template on the *Search* page. After you successfully create the automation, it is displayed in the *Automation* column.

Note: When you delete a policy that has an associated policy automation, the policy automation is automatically deleted as part of clean up.

Your policy violation automation is created from the console.

2.6.3.3. Creating a policy violation automation from the CLI

Complete the following steps to configure a policy violation automation from the CLI:

1. From your terminal, log in to your Red Hat Advanced Cluster Management hub cluster using the **oc login** command.
2. Find or create a policy that you want to add an automation to. Note the policy name and namespace.
3. Create a **PolicyAutomation** resource using the following sample as a guide:

```
apiVersion: policy.open-cluster-management.io/v1beta1
kind: PolicyAutomation
metadata:
  name: policynamespace-policy-automation
spec:
  automationDef:
```

```

extra_vars:
  your_var: your_value
name: Policy Compliance Template
secret: ansible-tower
type: AnsibleJob
mode: disabled
policyRef: policyname

```

4. The Automation template name in the previous sample is **Policy Compliance Template**. Change that value to match your job template name.
5. In the **extra_vars** section, add any parameters you need to pass to the Automation template.
6. Set the mode to either **once**, **everyEvent**, or **disabled**.
7. Set the **policyRef** to the name of your policy.
8. Create a secret in the same namespace as this **PolicyAutomation** resource that contains the Ansible Automation Platform credential. In the previous example, the secret name is **ansible-tower**. Use the [sample from application lifecycle](#) to see how to create the secret.
9. Create the **PolicyAutomation** resource.

Notes:

- An immediate run of the policy automation can be initiated by adding the following annotation to the **PolicyAutomation** resource:

```

metadata:
  annotations:
    policy.open-cluster-management.io/rerun: "true"

```

- When the policy is in **once** mode, the automation runs when the policy is non-compliant. The **extra_vars** variable, named **target_clusters** is added and the value is an array of each managed cluster name where the policy is non-compliant.
- When the policy is in **everyEvent** mode and the **DelayAfterRunSeconds** exceeds the defined time value, the policy is non-compliant and the automation runs for every policy violation.

2.6.4. Deploying policies using GitOps

You can deploy a set of policies across a fleet of managed clusters with the governance framework. You can add to the open source community, [policy-collection](#) by contributing to and using the policies in the repository. For more information, see [Contributing a custom policy](#). Policies in each of the **stable** and **community** folders from the open source community are further organized according to [NIST Special Publication 800-53](#).

Continue reading to learn best practices to use GitOps to automate and track policy updates and creation through a Git repository.

Prerequisite: Before you begin, be sure to fork the **policy-collection** repository.

- [Customizing your local repository](#)
- [Committing to your local repository](#)

- [Deploying policies to your cluster](#)
- [Verifying GitOps policy deployments from the console](#)
- [Verifying GitOps policy deployments from the CLI](#)

2.6.4.1. Customizing your local repository

Customize your local repository by consolidating the **stable** and **community** policies into a single folder. Remove the policies you do not want to use. Complete the following steps to customize your local repository:

1. Create a new directory in the repository to hold the policies that you want to deploy. Be sure that you are in your local **policy-collection** repository on your main default branch for GitOps. Run the following command:

```
mkdir my-policies
```

2. Copy all of the **stable** and **community** policies into your **my-policies** directory. Start with the **community** policies first, in case the **stable** folder contains duplicates of what is available in the community. Run the following commands:

```
cp -R community/* my-policies/
```

```
cp -R stable/* my-policies/
```

Now that you have all of the policies in a single parent directory structure, you can edit the policies in your fork.

Tips:

- It is best practice to remove the policies you are not planning to use.
- Learn about policies and the definition of the policies from the following list:
 - Purpose: Understand what the policy does.
 - Remediation Action: Does the policy only inform you of compliance, or enforce the policy and make changes? See the **spec.remediationAction** parameter. If changes are enforced, make sure you understand the functional expectation. Remember to check which policies support enforcement. For more information, view the *Validate* section. **Note:** The **spec.remediationAction** set for the policy overrides any remediation action that is set in the individual **spec.policy-templates**.
 - Placement: What clusters is the policy deployed to? By default, most policies target the clusters with the **environment: dev** label. Some policies may target OpenShift Container Platform clusters or another label. You can update or add additional labels to include other clusters. When there is no specific value, the policy is applied to all of your clusters. You can also create multiple copies of a policy and customize each one if you want to use a policy that is configured one way for one set of clusters and configured another way for another set of clusters.

2.6.4.2. Committing to your local repository

After you are satisfied with the changes you have made to your directory, commit and push your changes to Git so that they can be accessed by your cluster.

Note: This example is used to show the basics of how to use policies with GitOps, so you might have a different workflow to get changes to your branch.

Complete the following steps:

1. From your terminal, run **git status** to view your recent changes in your directory that you previously created. Add your new directory to the list of changes to be committed with the following command:

```
git add my-policies/
```

2. Commit the changes and customize your message. Run the following command:

```
git commit -m "Policies to deploy to the hub cluster"
```

3. Push the changes to the branch of your forked repository that is used for GitOps. Run the following command:

```
git push origin <your_default_branch>master
```

Your changes are committed.

2.6.4.3. Deploying policies to your cluster

After you push your changes, you can deploy the policies to your Red Hat Advanced Cluster Management for Kubernetes installation. Post deployment, your hub cluster is connected to your Git repository. Any further changes to your chosen branch of the Git repository is reflected in your cluster.

Note: By default, policies deployed with GitOps use the **merge** reconcile option. If you want to use the **replace** reconcile option instead, add the **apps.open-cluster-management.io/reconcile-option: replace** annotation to the **Subscription** resource. See [Application Lifecycle](#) for more details.

The **deploy.sh** script creates **Channel** and **Subscription** resources in your hub cluster. The channel connects to the Git repository, and the subscription specifies the data to bring to the cluster through the channel. As a result, all policies defined in the specified subdirectory are created on your hub. After the policies are created by the subscription, Red Hat Advanced Cluster Management analyzes the policies and creates additional policy resources in the namespace associated with each managed cluster that the policy is applied to, based on the defined placement rule.

The policy is then copied to the managed cluster from its respective managed cluster namespace on the hub cluster. As a result, the policies in your Git repository are pushed to all managed clusters that have labels that match the **clusterSelector** that are defined in the placement rule of your policy.

Complete the following steps:

1. From the **policy-collection** folder, run the following command to change the directory:

```
cd deploy
```

2. Make sure that your command line interface (CLI) is configured to create resources on the correct cluster with the following command:

```
oc cluster-info
```

The output of the command displays the API server details for the cluster, where Red Hat Advanced Cluster Management is installed. If the correct URL is not displayed, configure your CLI to point to the correct cluster. See [Using the OpenShift CLI](#) for more information.

3. Create a namespace where your policies are created to control access and to organize the policies. Run the following command:

```
oc create namespace policy-namespace
```

4. Run the following command to deploy the policies to your cluster:

```
./deploy.sh -u https://github.com/<your-repository>/policy-collection -p my-policies -n policy-namespace
```

Replace **your-repository** with your Git user name or repository name.

Note: For reference, the full list of arguments for the **deploy.sh** script uses the following syntax:

```
./deploy.sh [-u <url>] [-b <branch>] [-p <path/to/dir>] [-n <namespace>] [-a|--name <resource-name>]
```

View the following explanations for each argument:

- URL: The URL to the repository that you forked from the main **policy-collection** repository. The default URL is <https://github.com/stolostron/policy-collection.git>.
- Branch: Branch of the Git repository to point to. The default branch is **main**.
- Subdirectory Path: The subdirectory path you created to contain the policies you want to use. In the previous sample, we used the **my-policies** subdirectory, but you can also specify which folder you want start with. For example, you can use **my-policies/AC-Access-Control**. The default folder is **stable**.
- Namespace: The namespace where the resources and policies are created on the hub cluster. These instructions use the **policy-namespace** namespace. The default namespace is **policies**.
- Name Prefix: Prefix for the **Channel** and **Subscription** resources. The default is **demo-stable-policies**.

After you run the **deploy.sh** script, any user with access to the repository can commit changes to the branch, which pushes changes to existing policies on your clusters.

Note: To deploy policies with subscriptions, complete the following steps:

1. Bind the **open-cluster-management:subscription-admin** ClusterRole to the user creating the subscription.
2. If you are using an allow list in the subscription, include the following API entries:

```
- apiVersion: policy.open-cluster-management.io/v1
  kinds:
    - "*"
- apiVersion: policy.open-cluster-management.io/v1beta1
  kinds:
```

```

- apiVersion: apps.open-cluster-management.io/v1
  kinds:
  - PlacementRule
- apiVersion: cluster.open-cluster-management.io/v1beta1
  kinds:
  - Placement

```

2.6.4.4. Verifying GitOps policy deployments from the console

Verify that your changes were applied to your policies from the console. You can also make more changes to your policy from the console, however the changes are reverted when the **Subscription** is reconciled with the Git repository. Complete the following steps:

1. Log in to your Red Hat Advanced Cluster Management cluster.
2. From the navigation menu, select **Governance**.
3. Locate the policies that you deployed in the table. Policies that are deployed using GitOps have a *Git* label in the *Source* column. Click the label to view the details for the Git repository.

2.6.4.4.1. Verifying GitOps policy deployments from the CLI

Complete the following steps:

1. Check for the following policy details:
 - Why is a specific policy compliant or non-compliant on the clusters that it was distributed to?
 - Are the policies applied to the correct clusters?
 - If this policy is not distributed to any clusters, why?
2. Identify the GitOps deployed policies that you created or modified. The GitOps deployed policies can be identified by the annotation that is applied automatically. Annotations for the GitOps deployed policies resemble the following paths:

```
apps.open-cluster-management.io/hosting-deployable: policies/deploy-stable-policies-Policy-policy-role9
```

```
apps.open-cluster-management.io/hosting-subscription: policies/demo-policies
```

```
apps.open-cluster-management.io/sync-source: subgbk8s-policies/demo-policies
```

GitOps annotations are valuable to see which subscription created the policy. You can also add your own labels to your policies so that you can write runtime queries that select policies based on labels.

For example, you can add a label to a policy with the following command:

```
oc label policies.policy.open-cluster-management.io <policy-name> -n <policy-namespace> <key>=<value>
```

Then, you can query policies that have labels with the following command:

```
oc get policies.policy.open-cluster-management.io -n <policy-namespace> -l <key>=<value>
```

Your policies are deployed using GitOps.

2.6.5. Support for templates in configuration policies

Configuration policies support the inclusion of Golang text templates in the object definitions. These templates are resolved at runtime either on the hub cluster or the target managed cluster using configurations related to that cluster. This gives you the ability to define configuration policies with dynamic content, and inform or enforce Kubernetes resources that are customized to the target cluster.

- [Prerequisite](#)
- [Template functions](#)
- [Support for hub cluster templates in configuration policies](#)
 - [Template processing](#)
 - [Special annotation for reprocessing](#)
 - [Object template processing](#)
 - [Bypass template processing](#)
 - [Comparison of hub cluster and managed cluster templates](#)

2.6.5.1. Prerequisite

- The template syntax must be conformed to the Golang template language specification, and the resource definition generated from the resolved template must be a valid YAML. See the Golang documentation about [Package templates](#) for more information. Any errors in template validation are recognized as policy violations. When you use a custom template function, the values are replaced at runtime.

2.6.5.2. Template functions

Template functions, such as resource-specific and generic **lookup** template functions, are available for referencing Kubernetes resources on the hub cluster (using the `{{hub ... hub}}` delimiters), or managed cluster (using the `{{ ... }}` delimiters). See [Support for hub cluster templates in configuration policies](#) for more details. The resource-specific functions are used for convenience and makes content of the resources more accessible. If you use the generic function, **lookup**, which is more advanced, it is best to be familiar with the YAML structure of the resource that is being looked up. In addition to these functions, utility functions like **base64enc**, **base64dec**, **indent**, **autoindent**, **toInt**, **toBool**, and more are also available.

To conform templates with YAML syntax, templates must be set in the policy resource as strings using quotes or a block character (`|` or `>`). This causes the resolved template value to also be a string. To override this, consider using **toInt** or **toBool** as the final function in the template to initiate further processing that forces the value to be interpreted as an integer or boolean respectively.

Continue reading to view descriptions and examples for some of the custom template functions that are supported:

- [fromSecret function](#)

- [fromConfigMap](#) function
- [fromClusterClaim](#) function
- [lookup](#) function
- [base64enc](#) function
- [base64dec](#) function
- [indent](#) function
- [autoindent](#) function
- [toInt](#) function
- [toBool](#) function
- [protect](#) function
- [toLiteral](#) function
- [Open source community functions](#)

2.6.5.2.1. [fromSecret](#) function

The **fromSecret** function returns the value of the given data key in the secret. View the following syntax for the function:

```
func fromSecret (ns string, secretName string, datakey string) (dataValue string, err error)
```

When you use this function, enter the namespace, name, and data key of a Kubernetes **Secret** resource. You must use the same namespace that is used for the policy when using the function in a hub cluster template. See, [Support for hub cluster templates in configuration policies](#) for more details.

Note: When you use this function with hub cluster templates, the output is automatically encrypted using the [protect](#) function.

You receive a policy violation if the Kubernetes **Secret** resource does not exist on the target cluster. If the data key does not exist on the target cluster, the value becomes an empty string. View the following configuration policy that enforces a **Secret** resource on the target cluster. The value for the **PASSWORD** data key is a template that references the secret on the target cluster:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromsecret
  namespace: test
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
```

```

objectDefinition:
  apiVersion: v1
  data:
    USER_NAME: YWRtaW4=
    PASSWORD: '{{ fromSecret "default" "localsecret" "PASSWORD" }}'
  kind: Secret
  metadata:
    name: demosecret
    namespace: test
  type: Opaque
remediationAction: enforce
severity: low

```

2.6.5.2.2. *fromConfigMap* function

The **fromConfigMap** function returns the value of the given data key in the ConfigMap. View the following syntax for the function:

```
func fromConfigMap (ns string, configmapName string, datakey string) (dataValue string, err Error)
```

When you use this function, enter the namespace, name, and data key of a Kubernetes **ConfigMap** resource. You must use the same namespace that is used for the policy using the function in a hub cluster template. See, [Support for hub cluster templates in configuration policies](#) for more details. You receive a policy violation if the Kubernetes **ConfigMap** resource does not exist on the target cluster. If the data key does not exist on the target cluster, the value becomes an empty string. View the following configuration policy that enforces a Kubernetes resource on the target managed cluster. The value for the **log-file** data key is a template that retrieves the value of the **log-file** from the ConfigMap, **logs-config** from the **default** namespace, and the **log-level** is set to the data key **log-level**.

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromcm-lookup
  namespace: test-templates
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        kind: ConfigMap
        apiVersion: v1
        metadata:
          name: demo-app-config
          namespace: test
        data:
          app-name: sampleApp
          app-description: "this is a sample app"
          log-file: '{{ fromConfigMap "default" "logs-config" "log-file" }}'
          log-level: '{{ fromConfigMap "default" "logs-config" "log-level" }}'
        remediationAction: enforce
        severity: low

```

2.6.5.2.3. *fromClusterClaim* function

The **fromClusterClaim** function returns the value of the **Spec.Value** in the **ClusterClaim** resource. View the following syntax for the function:

```
func fromClusterClaim (clusterclaimName string) (dataValue string, err Error)
```

When you use this function, enter the name of a Kubernetes **ClusterClaim** resource. You receive a policy violation if the **ClusterClaim** resource does not exist. View the following example of the configuration policy that enforces a Kubernetes resource on the target managed cluster. The value for the **platform** data key is a template that retrieves the value of the **platform.open-cluster-management.io** cluster claim. Similarly, it retrieves values for **product** and **version** from the **ClusterClaim**:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-clusterclaims
  namespace: default
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        kind: ConfigMap
        apiVersion: v1
        metadata:
          name: sample-app-config
          namespace: default
        data:
          # Configuration values can be set as key-value properties
          platform: '{{ fromClusterClaim "platform.open-cluster-management.io" }}'
          product: '{{ fromClusterClaim "product.open-cluster-management.io" }}'
          version: '{{ fromClusterClaim "version.openshift.io" }}'
      remediationAction: enforce
      severity: low
```

2.6.5.2.4. *lookup* function

The **lookup** function returns the Kubernetes resource as a JSON compatible map. If the requested resource does not exist, an empty map is returned. If the resource does not exist and the value is provided to another template function, you might get the following error: **invalid value; expected string**.

Note: Use the **default** template function, so the correct type is provided to later template functions. See the *Open source community functions* section.

View the following syntax for the function:

```
func lookup (apiversion string, kind string, namespace string, name string) (value string, err Error)
```

When you use this function, enter the API version, kind, namespace, and name of the Kubernetes resource. You must use the same namespace that is used for the policy within the hub cluster template. See, [Support for hub cluster templates in configuration policies](#) for more details. View the following example of the configuration policy that enforces a Kubernetes resource on the target managed cluster. The value for the **metrics-url** data key is a template that retrieves the **v1/Service** Kubernetes resource **metrics** from the **default** namespace, and is set to the value of the **Spec.ClusterIP** in the queried resource:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-lookup
  namespace: test-templates
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        kind: ConfigMap
        apiVersion: v1
        metadata:
          name: demo-app-config
          namespace: test
        data:
          # Configuration values can be set as key-value properties
          app-name: sampleApp
          app-description: "this is a sample app"
          metrics-url: |
            http://{{ (lookup "v1" "Service" "default" "metrics").spec.clusterIP }}:8080
      remediationAction: enforce
      severity: low
```

2.6.5.2.5. *base64enc* function

The **base64enc** function returns a **base64** encoded value of the input **data string**. View the following syntax for the function:

```
func base64enc (data string) (enc-data string)
```

When you use this function, enter a string value. View the following example of the configuration policy that uses the **base64enc** function:

```
apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromsecret
  namespace: test
spec:
  namespaceSelector:
    exclude:
      - kube-*
```

```

include:
- default
object-templates:
- complianceType: musthave
objectDefinition:
...
data:
  USER_NAME: '{{ fromConfigMap "default" "myconfigmap" "admin-user" | base64enc }}'

```

2.6.5.2.6. *base64dec* function

The **base64dec** function returns a **base64** decoded value of the input **enc-data string**. View the following syntax for the function:

```
func base64dec (enc-data string) (data string)
```

When you use this function, enter a string value. View the following example of the configuration policy that uses the **base64dec** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromsecret
  namespace: test
spec:
  namespaceSelector:
    exclude:
    - kube-*
    include:
    - default
  object-templates:
  - complianceType: musthave
  objectDefinition:
  ...
  data:
    app-name: |
      '{{ ( lookup "v1" "Secret" "testns" "mytestsecret" ) .data.appname ) | base64dec }}'

```

2.6.5.2.7. *indent* function

The **indent** function returns the padded **data string**. View the following syntax for the function:

```
func indent (spaces int, data string) (padded-data string)
```

When you use this function, enter a data string with the specific number of spaces. View the following example of the configuration policy that uses the **indent** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromsecret
  namespace: test
spec:
  namespaceSelector:

```

```

exclude:
- kube-*
include:
- default
object-templates:
- complianceType: musthave
objectDefinition:
...
data:
  Ca-cert: |
    {{ ( index ( lookup "v1" "Secret" "default" "mycert-tls" ).data "ca.pem" ) | base64dec | indent 4
}}

```

2.6.5.2.8. *autoindent* function

The **autoindent** function acts like the **indent** function that automatically determines the number of leading spaces based on the number of spaces before the template. View the following example of the configuration policy that uses the **autoindent** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-fromsecret
  namespace: test
spec:
  namespaceSelector:
    exclude:
    - kube-*
    include:
    - default
  object-templates:
  - complianceType: musthave
    objectDefinition:
    ...
    data:
      Ca-cert: |
        {{ ( index ( lookup "v1" "Secret" "default" "mycert-tls" ).data "ca.pem" ) | base64dec |
autoindent }}

```

2.6.5.2.9. *toInt* function

The **toInt** function casts and returns the integer value of the input value. Also, when this is the last function in the template, there is further processing of the source content. This is to ensure that the value is interpreted as an integer by the YAML. View the following syntax for the function:

```
func toInt (input interface{}) (output int)
```

When you use this function, enter the data that needs to be casted as an integer. View the following example of the configuration policy that uses the **toInt** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-template-function

```

```

namespace: test
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        ...
        spec:
          vlanid: |
            {{ (fromConfigMap "site-config" "site1" "vlan") | toInt }}

```

2.6.5.2.10. *toBool* function

The **toBool** function converts the input string into a boolean, and returns the boolean. Also, when this is the last function in the template, there is further processing of the source content. This is to ensure that the value is interpreted as a boolean by the YAML. View the following syntax for the function:

```
func toBool (input string) (output bool)
```

When you use this function, enter the string data that needs to be converted to a boolean. View the following example of the configuration policy that uses the **toBool** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: demo-template-function
  namespace: test
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        ...
        spec:
          enabled: |
            {{ (fromConfigMap "site-config" "site1" "enabled") | toBool }}

```

2.6.5.2.11. *protect* function

The **protect** function enables you to encrypt a string in a hub cluster policy template. It is automatically decrypted on the managed cluster when the policy is evaluated. View the following example of the configuration policy that uses the **protect** function:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:

```

```

name: demo-template-function
namespace: test
spec:
  namespaceSelector:
    exclude:
      - kube-*
    include:
      - default
  object-templates:
    - complianceType: musthave
      objectDefinition:
        ...
      spec:
        enabled: |
          {{hub (lookup "v1" "Secret" "default" "my-hub-secret").data.message | protect hub}}

```

In the previous YAML example, there is an existing hub cluster policy template that is defined to use the **lookup** function. On the replicated policy in the managed cluster namespace, the value might resemble the following syntax: **\$ocm_encrypted:okrrBqt72ol+3WT/0vxel3vGa+wpLD7Z0ZxFMLvL204=**

Each encryption algorithm used is AES-CBC using 256-bit keys. Each encryption key is unique per managed cluster and is automatically rotated every 30 days.

This ensures that your decrypted value is to never be stored in the policy on the managed cluster.

To force an immediate rotation, delete the **policy.open-cluster-management.io/last-rotated** annotation on the **policy-encryption-key** Secret in the managed cluster namespace on the hub cluster. Policies are then reprocessed to use the new encryption key.

2.6.5.2.12. *toLiteral* function

The **toLiteral** function removes any quotation marks around the template string after it is processed. You can use this function to convert a JSON string from a ConfigMap field to a JSON value in the manifest. Run the following function to remove quotation marks from the **key** parameter value:

```
key: '{{ "[\"10.10.10.10\", \"1.1.1.1\"]" | toLiteral }}'
```

After using the **toLiteral** function, the following update is displayed:

```
key: ["10.10.10.10", "1.1.1.1"]
```

2.6.5.2.13. Open source community functions

Additionally, Red Hat Advanced Cluster Management supports the following template functions that are included from the **sprig** open source project:

- **cat**
- **contains**
- **default**
- **empty**
- **fromJson**

- **hasPrefix**
- **hasSuffix**
- **join**
- **list**
- **lower**
- **mustFromJson**
- **quote**
- **replace**
- **semver**
- **semverCompare**
- **split**
- **splitn**
- **ternary**
- **trim**
- **until**
- **untilStep**
- **upper**

See the [Sprig Function Documentation](#) for more details.

2.6.5.3. Support for hub cluster templates in configuration policies

In addition to managed cluster templates that are dynamically customized to the target cluster, Red Hat Advanced Cluster Management also supports hub cluster templates to define configuration policies using values from the hub cluster. This combination reduces the need to create separate policies for each target cluster or hardcode configuration values in the policy definitions.

Hub cluster templates are based on Golang text template specifications, and the `{{hub ... hub}}` delimiter indicates a hub cluster template in a configuration policy.

For security, both resource-specific and the generic lookup functions in hub cluster templates are restricted to the namespace of the policy on the hub cluster. View the [Comparison of hub and managed cluster templates](#) for additional details.

Important: If you use hub cluster templates to propagate secrets or other sensitive data, the sensitive data exists in the managed cluster namespace on the hub cluster and on the managed clusters where that policy is distributed. The template content is expanded in the policy, and policies are not encrypted by the OpenShift Container Platform ETCD encryption support. To address this, use **fromSecret**, which automatically encrypts the values from the Secret, or **protect** to encrypt other values.

2.6.5.3.1. Template processing

A configuration policy definition can contain both hub cluster and managed cluster templates. Hub cluster templates are processed first on the hub cluster, then the policy definition with resolved hub cluster templates is propagated to the target clusters. On the managed cluster, the **ConfigurationPolicyController** processes any managed cluster templates in the policy definition and then enforces or verifies the fully resolved object definition.

2.6.5.3.2. Special annotation for reprocessing

Hub cluster templates are resolved to the data in the referenced resources during policy creation, or when the referenced resources are updated.

If you need to manually initiate an update, use the special annotation, **policy.open-cluster-management.io/trigger-update**, to indicate changes for the data referenced by the templates. Any change to the special annotation value automatically initiates template processing. Additionally, the latest contents of the referenced resource are read and updated in the policy definition that is propagated for processing on managed clusters. A way to use this annotation is to increment the value by one each time.

2.6.5.4. Object template processing

Set object templates with a YAML string representation. The **object-template-raw** parameter is an optional parameter that supports advanced templating use-cases, such as if-else and the **range** function. The following example is defined to add the **species-category: mammal** label to any ConfigMap in the **default** namespace that has a **name** key equal to **Sea Otter**:

```
object-templates-raw: |
  {{- range (lookup "v1" "ConfigMap" "default" "").items }}
  {{- if eq .data.name "Sea Otter" }}
  - complianceType: musthave
    objectDefinition:
      kind: ConfigMap
      apiVersion: v1
      metadata:
        name: {{ .metadata.name }}
        namespace: {{ .metadata.namespace }}
        labels:
          species-category: mammal
  {{- end }}
  {{- end }}
```

Note: While **spec.object-templates** and **spec.object-templates-raw** are optional, one of the two parameter fields must be set.

2.6.5.4.1. Bypass template processing

You might create a policy that contains a template that is not intended to be processed by Red Hat Advanced Cluster Management. By default, Red Hat Advanced Cluster Management processes all templates.

To bypass template processing for your hub cluster, you must change **{{ template content }}** to **{{ `{{ template content }}` }}**.

Alternatively, you can add the following annotation in the **ConfigurationPolicy** section of your **Policy**: **policy.open-cluster-management.io/disable-templates: "true"**. When this annotation is included, the previous workaround is not necessary. Template processing is bypassed for the **ConfigurationPolicy**.

See the following table for a comparison of hub cluster and managed cluster templates:

2.6.5.4.2. Comparison of hub cluster and managed cluster templates

Table 2.13. Comparison table

| Templates | Hub cluster | Managed cluster |
|----------------|---|---|
| Syntax | Golang text template specification | Golang text template specification |
| Delimiter | <code>{{hub ... hub}}</code> | <code>{{ ... }}</code> |
| Context | A .ManagedClusterName variable is available, which at runtime, resolves to the name of the target cluster where the policy is propagated. | No context variables |
| Access control | You can only reference namespaced Kubernetes objects that are in the same namespace as the Policy resource. | You can reference any resource on the cluster. |
| Functions | <p>A set of template functions that support dynamic access to Kubernetes resources and string manipulation. See Template functions for more information. See the Access control row for lookup restrictions.</p> <p>The fromSecret template function on the hub cluster stores the resulting value as an encrypted string on the replicated policy, in the managed cluster namespace.</p> <p>The equivalent call might use the following syntax: <code>{{hub "(lookup "v1" "Secret" "default" "my-hub-secret").data.message protect hub}}</code></p> | <p>A set of template functions support dynamic access to Kubernetes resources and string manipulation. See Template functions for more information.</p> |

| Templates | Hub cluster | Managed cluster |
|-------------------------|---|--|
| Function output storage | The output of template functions are stored in Policy resource objects in each applicable managed cluster namespace on the hub cluster, before it is synced to the managed cluster. This means that any sensitive results from template functions are readable by anyone with read access to the Policy resource objects on the hub cluster, and read access with ConfigurationPolicy resource objects on the managed clusters. Additionally, if etcd encryption is enabled, the Policy and ConfigurationPolicy resource objects are not encrypted. It is best to carefully consider this when using template functions that return sensitive output (e.g. from a secret). | The output of template functions are not stored in policy related resource objects. |
| Processing | Processing occurs at runtime on the hub cluster during propagation of replicated policies to clusters. Policies and the hub cluster templates within the policies are processed on the hub cluster only when templates are created or updated. | Processing occurs in the ConfigurationPolicyController on the managed cluster. Policies are processed periodically, which automatically updates the resolved object definition with data in the referenced resources. |
| Processing errors | Errors from the hub cluster templates are displayed as violations on the managed clusters the policy applies to. | Errors from the managed cluster templates are displayed as violations on the specific target cluster where the violation occurred. |

2.6.6. Managing security policies

Create a security policy to report and validate your cluster compliance based on your specified security standards, categories, and controls.

View the following sections:

- [Creating a security policy](#)
 - [Creating a security policy from the command line interface](#)
 - [Viewing your security policy from the CLI](#)

- [Creating a cluster security policy from the console](#)
- [Viewing your security policy from the console](#)
- [Creating policy sets from the CLI](#)
- [Creating policy sets from the console](#)
- [Updating security policies](#)
 - [Disabling security policies](#)
- [Deleting a security policy](#)
 - [Deleting policy sets from the console](#)
- [Cleaning up resources that are created by policies](#)

2.6.6.1. Creating a security policy

You can create a security policy from the command line interface (CLI) or from the console.

Required access: Cluster administrator

Important: You must define a placement rule and placement binding to apply your policy to a specific cluster. Enter a valid value for the *Cluster selector* field to define a **PlacementRule** and **PlacementBinding**. See [Resources that support support set-based requirements](#) in the Kubernetes documentation for a valid expression. View the definitions of the objects that are required for your Red Hat Advanced Cluster Management policy:

- *PlacementRule*: Defines a *cluster selector* where the policy must be deployed.
- *PlacementBinding*: Binds the placement to a placement rule.

View more descriptions of the policy YAML files in the [Policy overview](#).

2.6.6.1.1. Creating a security policy from the command line interface

Complete the following steps to create a policy from the command line interface (CLI):

1. Create a policy by running the following command:

```
kubectl create -f policy.yaml -n <policy-namespace>
```

2. Define the template that the policy uses. Edit your **.yaml** file by adding a **policy-templates** field to define a template. Your policy might resemble the following YAML file:

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy1
spec:
  remediationAction: "enforce" # or inform
  disabled: false # or true
  namespaceSelector:
    include:
```

```

- "default"
- "my-namespace"
policy-templates:
- objectDefinition:
  apiVersion: policy.open-cluster-management.io/v1
  kind: ConfigurationPolicy
  metadata:
    name: operator
    # namespace: # will be supplied by the controller via the namespaceSelector
  spec:
    remediationAction: "inform"
    object-templates:
      - complianceType: "musthave" # at this level, it means the role must exist and must
        have the following rules
        apiVersion: rbac.authorization.k8s.io/v1
        kind: Role
        metadata:
          name: example
        objectDefinition:
          rules:
            - complianceType: "musthave" # at this level, it means if the role exists the rule is a
              musthave
            apiGroups: ["extensions", "apps"]
            resources: ["deployments"]
            verbs: ["get", "list", "watch", "create", "delete", "patch"]

```

3. Define a **PlacementRule**. Be sure to change the **PlacementRule** to specify the clusters where the policies need to be applied by adjusting the **clusterSelector**. View [Placement rule samples overview](#)

Your **PlacementRule** might resemble the following content:

```

apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement1
spec:
  clusterConditions:
    - type: ManagedClusterConditionAvailable
      status: "True"
  clusterNames:
    - "cluster1"
    - "cluster2"
  - clusterSelector
    matchLabels:
      cloud: IBM

```

4. Define a **PlacementBinding** to bind your policy to your **PlacementRule**. Your **PlacementBinding** might resemble the following YAML sample:

```

apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding1
placementRef:
  name: placement1

```

```

apiGroup: apps.open-cluster-management.io
kind: PlacementRule
subjects:
- name: policy1
  apiGroup: policy.open-cluster-management.io
  kind: Policy

```

2.6.6.1.1.1. Viewing your security policy from the CLI

Complete the following steps to view your security policy from the CLI:

1. View details for a specific security policy by running the following command:

```

kubectl get policies.policy.open-cluster-management.io <policy-name> -n <policy-namespace> -o yaml

```

2. View a description of your security policy by running the following command:

```

kubectl describe policies.policy.open-cluster-management.io <policy-name> -n <policy-namespace>

```

2.6.6.1.1.2. Creating a cluster security policy from the console

After you log into your Red Hat Advanced Cluster Management, navigate to the *Governance* page and click **Create policy**.

As you create your new policy from the console, a YAML file is also created in the YAML editor. To view the YAML editor, select the toggle at the beginning of the *Create policy* form to enable it.

Complete the *Create policy* form, then select the **Submit** button.

Your YAML file might resemble the following policy:

```

apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy-pod
  annotations:
    policy.open-cluster-management.io/categories:
'SystemAndCommunicationsProtections,SystemAndInformationIntegrity'
    policy.open-cluster-management.io/controls: 'control example'
    policy.open-cluster-management.io/standards: 'NIST,HIPAA'
spec:
  complianceType: musthave
  namespaces:
    exclude: ["kube*"]
    include: ["default"]
    pruneObjectBehavior: None
  object-templates:
  - complianceType: musthave
    objectDefinition:
      apiVersion: v1
      kind: Pod
      metadata:

```

```

    name: pod1
  spec:
    containers:
    - name: pod-name
      image: 'pod-image'
      ports:
      - containerPort: 80
  remediationAction: enforce
  disabled: false
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-pod
placementRef:
  name: placement-pod
  kind: PlacementRule
  apiGroup: apps.open-cluster-management.io
subjects:
- name: policy-pod
  kind: Policy
  apiGroup: policy.open-cluster-management.io
---
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-pod
spec:
  clusterConditions: []
  clusterSelector:
    matchLabels:
      cloud: "IBM"

```

Click **Create Policy**. A security policy is created from the console.

2.6.6.1.2.1. Viewing your security policy from the console

View any security policy and its status from the console. Navigate to the *Governance* page to view a table list of your policies. **Note:** You can filter the table list of your policies by selecting the *Policies* tab or *Cluster violations* tab.

Select one of your policies to view more details. The *Details*, *Clusters*, and *Templates* tabs are displayed. When the cluster or policy status cannot be determined, the following message is displayed: **No status**.

2.6.6.1.3. Creating policy sets from the CLI

By default, the policy set is created with no policies or placements. You must create a placement for the policy set and have at least one policy that exists on your cluster. When you create a policy set, you can add numerous policies. Run the following command to create a policy set from the CLI:

```
kubectl apply -f <policyset-filename>
```

2.6.6.1.4. Creating policy sets from the console

From the navigation menu, select **Governance**. Then select the *Policy sets* tab. Select the **Create policy set** button and complete the form. After you add the details for your policy set, select the **Submit** button.

View the stable **Policysets**, which require the Policy Generator for deployment, [PolicySets-- Stable](#).

2.6.6.2. Updating security policies

Learn to update security policies by viewing the following section.

2.6.6.2.1. Adding a policy to a policy set from the CLI

Run the following command to edit your policy set: **kubectl edit policysets your-policyset-name**

Add the policy name to the list in the **policies** section of the policy set. Apply your added policy in the placement section of your policy set with the following command: **kubectl apply -f your-added-policy.yaml**. A **PlacementBinding** and **PlacementRule** are created. **Note:** If you delete the placement binding, the policy is still placed by the policy set.

2.6.6.2.2. Adding a policy to a policy set from the console

Add a policy to the policy set by selecting the *Policy sets* tab. Select the Actions icon and select **Edit**. The *Edit policy set* form appears.

Navigate to the *Policies* section of the form to select a policy to add to the policy set.

2.6.6.2.3. Disabling security policies

Your policy is enabled by default. Disable your policy from the console.

After you log into your Red Hat Advanced Cluster Management for Kubernetes console, navigate to the *Governance* page to view a table list of your policies.

Select the **Actions** icon > **Disable policy**. The *Disable Policy* dialog box appears.

Click **Disable policy**. Your policy is disabled.

2.6.6.3. Deleting a security policy

Delete a security policy from the CLI or the console.

- Delete a security policy from the CLI:
 - a. Delete a security policy by running the following command:

```
kubectl delete policies.policy.open-cluster-management.io <policy-name> -n <policy-namespace>
```

After your policy is deleted, it is removed from your target cluster or clusters. Verify that your policy is removed by running the following command: **kubectl get policies.policy.open-cluster-management.io <policy-name> -n <policy-namespace>**

- Delete a security policy from the console:

From the navigation menu, click **Governance** to view a table list of your policies. Click the **Actions** icon for the policy you want to delete in the policy violation table.

Click **Remove**. From the *Remove policy* dialog box, click **Remove policy**

2.6.6.3.1. Deleting policy sets from the console

From the *Policy sets* tab, select the **Actions** icon for the policy set. When you click **Delete**, the *Permanently delete Policyset?* dialogue box appears.

Click the **Delete** button.

To manage other policies, see [Managing security policies](#) for more information. Refer to [Governance](#) for more topics about policies.

2.6.6.4. Cleaning up resources that are created by policies

Use the **pruneObjectBehavior** parameter in a configuration policy to clean up resources that are created by the policy. When **pruneObjectBehavior** is set, the related objects are only cleaned up after the configuration policy (or parent policy) associated with them is deleted. View the following descriptions of the values that can be used for the parameter:

- **DeleteIfCreated**: Cleans up any resources created by the policy.
- **DeleteAll**: Cleans up all resources managed by the policy.
- **None**: This is the default value and maintains the same behavior from previous releases, where no related resources are deleted.

You can set the value directly in the YAML as you create a policy from the CLI. From the console, you can select the value in the *Prune Object Behavior* section of the *Policy templates* step.

Notes:

- If a policy that installs an operator has the **pruneObjectBehavior** parameter defined, then additional clean up is needed to complete the operator uninstall. You might need to delete the operator **ClusterServiceVersion** object as part of this cleanup.
- As you disable the **config-policy-addon** resource on the managed cluster, the **pruneObjectBehavior** is ignored. To automatically clean up the related resources on the policies, you must remove the policies from the managed cluster before the add-on is disabled.

2.6.7. Managing configuration policies

Learn to create, apply, view, and update your configuration policies.

Required access: Administrator or cluster administrator

- [Creating a configuration policy](#)
 - [Creating a configuration policy from the CLI](#)
 - [Viewing your configuration policy from the CLI](#)
 - [Creating a configuration policy from the console](#)
 - [Viewing your configuration policy from the console](#)
- [Updating configuration policies](#)

- [Disabling configuration policies](#)
- [Deleting a configuration policy](#)

2.6.7.1. Creating a configuration policy

You can create a YAML file for your configuration policy from the command line interface (CLI) or from the console.

If you have an existing Kubernetes manifest, consider using the Policy Generator to automatically include the manifests in a policy. See the [Policy Generator](#) documentation. View the following sections to create a configuration policy:

2.6.7.1.1. Creating a configuration policy from the CLI

Complete the following steps to create a configuration policy from the (CLI):

1. Create a YAML file for your configuration policy. Run the following command:

```
oc create -f configpolicy-1.yaml
```

Your configuration policy might resemble the following policy:

```
apiVersion: policy.open-cluster-management.io/v1
kind: Policy
metadata:
  name: policy-1
  namespace: my-policies
policy-templates:
- apiVersion: policy.open-cluster-management.io/v1
  kind: ConfigurationPolicy
  metadata:
    name: mustonlyhave-configuration
  spec:
    namespaceSelector:
      include: ["default"]
      exclude: ["kube-system"]
    remediationAction: inform
    disabled: false
    complianceType: mustonlyhave
    object-templates:
```

2. Apply the policy by running the following command:

```
oc apply -f <policy-file-name> --namespace=<namespace>
```

3. Verify and list the policies by running the following command:

```
oc get policies.policy.open-cluster-management.io --namespace=<namespace>
```

Your configuration policy is created.

2.6.7.1.2. Viewing your configuration policy from the CLI

Complete the following steps to view your configuration policy from the CLI:

1. View details for a specific configuration policy by running the following command:

```
oc get policies.policy.open-cluster-management.io <policy-name> -n <namespace> -o yaml
```

2. View a description of your configuration policy by running the following command:

```
oc describe policies.policy.open-cluster-management.io <name> -n <namespace>
```

2.6.7.1.3. Creating a configuration policy from the console

As you create a configuration policy from the console, a YAML file is also created in the YAML editor.

1. Log in to your cluster from the console, and select **Governance** from the navigation menu.
2. Click **Create policy**. Specify the policy you want to create by selecting one of the configuration policies for the specification parameter.
3. Continue with configuration policy creation by completing the policy form. Enter or select the appropriate values for the following fields:
 - Name
 - Specifications
 - Cluster selector
 - Remediation action
 - Standards
 - Categories
 - Controls
4. Click **Create**. Your configuration policy is created.

2.6.7.1.4. Viewing your configuration policy from the console

View any configuration policy and its status from the console.

After you log into your cluster from the console, select **Governance** to view a table list of your policies.

Note: You can filter the table list of your policies by selecting the *All policies* tab or *Cluster violations* tab.

Select one of your policies to view more details. The *Details*, *Clusters*, and *Templates* tabs are displayed.

2.6.7.2. Updating configuration policies

Learn to update configuration policies by viewing the following section.

2.6.7.2.1. Disabling configuration policies

Disable your configuration policy. Similar to the instructions mentioned earlier, log in and navigate to the *Governance* page to complete the tasks.

1. Select the **Actions** icon for a configuration policy from the table list, then click **Disable**. The *Disable Policy* dialog box appears.
2. Click **Disable policy**.

The policy is disabled, but not deleted.

2.6.7.3. Deleting a configuration policy

Delete a configuration policy from the CLI or the console.

- Delete a configuration policy from the CLI with the following procedure:
 1. Run the following command to delete the policy from your target cluster or clusters:


```
oc delete policies.policy.open-cluster-management.io <policy-name> -n <namespace>
```
 2. Verify that your policy is removed by running the following command:


```
oc get policies.policy.open-cluster-management.io <policy-name> -n <namespace>
```
- Delete a configuration policy from the console with the following procedure:
 1. From the navigation menu, click **Governance** to view a table list of your policies.
 2. Click the **Actions** icon for the policy you want to delete in the policy violation table, then click **Remove**.
 3. From the *Remove policy* dialog box, click **Remove policy**.

Your policy is deleted.

See configuration policy samples that are supported by Red Hat Advanced Cluster Management from the [CM-Configuration-Management folder](#).

Alternatively, you can refer to the [Table of sample configuration policies](#) to view other configuration policies that are monitored by the controller. For details to manage other policies, refer to [Managing security policies](#).

2.6.8. Managing gatekeeper operator policies

Use the gatekeeper operator policy to install the gatekeeper operator and gatekeeper on a managed cluster. Learn to create, view, and update your gatekeeper operator policies in the following sections.

Required access: Cluster administrator

- [Installing gatekeeper using a gatekeeper operator policy](#)
- [Creating a gatekeeper policy from the console](#)
 - [Gatekeeper operator custom resource](#)
- [Upgrading gatekeeper and the gatekeeper operator](#)
- [Updating gatekeeper operator policy](#)

- [Viewing gatekeeper operator policy from the console](#)
- [Disabling gatekeeper operator policy](#)
- [Deleting gatekeeper operator policy](#)
- [Uninstalling gatekeeper policy, gatekeeper, and gatekeeper operator policy](#)

2.6.8.1. Installing gatekeeper using a gatekeeper operator policy

Use the governance framework to install the gatekeeper operator. Gatekeeper operator is available in the OpenShift Container Platform catalog. See *Adding Operators to a cluster* in the [OpenShift Container Platform documentation](#) for more information.

Use the configuration policy controller to install the gatekeeper operator policy. During the install, the operator group and subscription pull the gatekeeper operator to install it in your managed cluster. Then, the gatekeeper operator creates a gatekeeper CR to configure gatekeeper. View the [Gatekeeper operator CR](#) sample.

Gatekeeper operator policy is monitored by the Red Hat Advanced Cluster Management configuration policy controller, where **enforce** remediation action is supported. Gatekeeper operator policies are created automatically by the controller when set to **enforce**.

2.6.8.2. Creating a gatekeeper policy from the console

Install the gatekeeper policy by creating a gatekeeper policy from the console. Alternatively, you can view a sample YAML to deploy [policy-gatekeeper-operator.yaml](#).

After you log into your cluster, navigate to the *Governance* page.

Select **Create policy**. As you complete the form, select **Gatekeeper Operator** from the *Specifications* field. The parameter values for your policy are automatically populated and the policy is set to **inform** by default. Set your remediation action to **enforce** to install gatekeeper.

Note: Default values are generated by the operator. See [Gatekeeper Helm Chart](#) for an explanation of the optional parameters that can be used for the gatekeeper operator policy.

2.6.8.2.1. Gatekeeper operator CR

```
apiVersion: operator.gatekeeper.sh/v1alpha1
kind: Gatekeeper
metadata:
  name: gatekeeper
spec:
  audit:
    replicas: 1
    logLevel: DEBUG
    auditInterval: 10s
    constraintViolationLimit: 55
    auditFromCache: Enabled
    auditChunkSize: 66
    emitAuditEvents: Enabled
  resources:
    limits:
      cpu: 500m
      memory: 150Mi
```

```

requests:
  cpu: 500m
  memory: 130Mi
validatingWebhook: Enabled
webhook:
  replicas: 2
  logLevel: ERROR
  emitAdmissionEvents: Enabled
  failurePolicy: Fail
resources:
  limits:
    cpu: 480m
    memory: 140Mi
  requests:
    cpu: 400m
    memory: 120Mi
nodeSelector:
  region: "EMEA"
affinity:
  podAffinity:
    requiredDuringSchedulingIgnoredDuringExecution:
      - labelSelector:
          matchLabels:
            auditKey: "auditValue"
        topologyKey: topology.kubernetes.io/zone
tolerations:
  - key: "Example"
    operator: "Exists"
    effect: "NoSchedule"
podAnnotations:
  some-annotation: "this is a test"
  other-annotation: "another test"

```

2.6.8.3. Upgrading gatekeeper and the gatekeeper operator

You can upgrade the versions for gatekeeper and the gatekeeper operator. When you install the gatekeeper operator with the gatekeeper operator policy, notice the value for **installPlanApproval**. The operator upgrades automatically when **installPlanApproval** is set to **Automatic**.

You must approve the upgrade of the gatekeeper operator manually, for each cluster, when **installPlanApproval** is set to **Manual**.

2.6.8.4. Updating gatekeeper operator policy

Learn to update the gatekeeper operator policy by viewing the following section.

2.6.8.4.1. Viewing gatekeeper operator policy from the console

View your gatekeeper operator policy and its status from the console.

After you log into your cluster from the console, click **Governance** to view a table list of your policies.

Note: You can filter the table list of your policies by selecting the *Policies* tab or *Cluster violations* tab.

Select the **policy-gatekeeper-operator** policy to view more details. View the policy violations by selecting the *Clusters* tab.

2.6.8.4.2. Disabling gatekeeper operator policy

Disable your gatekeeper operator policy.

After you log into your Red Hat Advanced Cluster Management for Kubernetes console, navigate to the *Governance* page to view a table list of your policies.

Select the **Actions** icon for the **policy-gatekeeper-operator** policy, then click **Disable**. The *Disable Policy* dialog box appears.

Click **Disable policy**. Your **policy-gatekeeper-operator** policy is disabled.

2.6.8.5. Deleting gatekeeper operator policy

Delete the gatekeeper operator policy from the CLI or the console.

- Delete gatekeeper operator policy from the CLI:
 - a. Delete gatekeeper operator policy by running the following command:

```
kubectl delete policies.policy.open-cluster-management.io <policy-gatekeeper-operator-name> -n <namespace>
```

After your policy is deleted, it is removed from your target cluster or clusters.

- b. Verify that your policy is removed by running the following command:

```
kubectl get policies.policy.open-cluster-management.io <policy-gatekeeper-operator-name> -n <namespace>
```

- Delete gatekeeper operator policy from the console:

Navigate to the *Governance* page to view a table list of your policies.

Similar to the previous console instructions, click the **Actions** icon for the **policy-gatekeeper-operator** policy. Click **Remove** to delete the policy. From the *Remove policy* dialog box, click **Remove policy**.

Your gatekeeper operator policy is deleted.

2.6.8.6. Uninstalling gatekeeper policy, gatekeeper, and gatekeeper operator policy

Complete the following steps to uninstall gatekeeper policy, gatekeeper, and gatekeeper operator policy:

1. Remove the gatekeeper **Constraint** and **ConstraintTemplate** that is applied on your managed cluster:
 - a. Edit your gatekeeper operator policy. Locate the **ConfigurationPolicy** template that you used to create the gatekeeper **Constraint** and **ConstraintTemplate**.
 - b. Change the value for **complianceType** of the **ConfigurationPolicy** template to **mustnothave**.
 - c. Save and apply the policy.
2. Remove gatekeeper instance from your managed cluster:

- a. Edit your gatekeeper operator policy. Locate the **ConfigurationPolicy** template that you used to create the Gatekeeper custom resource (CR).
 - b. Change the value for **complianceType** of the **ConfigurationPolicy** template to **mustnothave**.
3. Remove the gatekeeper operator that is on your managed cluster:
 - a. Edit your gatekeeper operator policy. Locate the **ConfigurationPolicy** template that you used to create the Subscription CR.
 - b. Change the value for **complianceType** of the **ConfigurationPolicy** template to **mustnothave**.

Gatekeeper policy, gatekeeper, and gatekeeper operator policy are uninstalled.

See [Integrating gatekeeper constraints and constraint templates](#) for details about gatekeeper. For a list of topics to integrate third-party policies with the product, see [Integrate third-party policy controllers](#).

2.6.9. Managing operator policies in disconnected environments

You might need to deploy Red Hat Advanced Cluster Management for Kubernetes policies on Red Hat OpenShift Container Platform clusters that are not connected to the internet (disconnected). If the policies you deploy are used to deploy policies that install an Operator Lifecycle Manager operator, you must follow the procedure for [Mirroring an Operator catalog](#).

Complete the following steps to validate access to the operator images:

1. See [Verify required packages are available](#) to validate that packages you require to use with policies are available. You must validate availability for each image registry used by any managed cluster that the following policies are deployed to:
 - **container-security-operator**
 - **Deprecated: gatekeeper-operator-product**
 - **compliance-operator**
2. See [Configure image content source policies](#) to validate that the sources are available. The image content source policies must exist on each of the disconnected managed clusters and can be deployed using a policy to simplify the process. See the following table of image source locations:

| Governance policy type | Image source location |
|------------------------|--------------------------------------|
| Container security | registry.redhat.io/quay |
| Compliance | registry.redhat.io/compliance |
| Gatekeeper | registry.redhat.io/rhacm2 |

2.7. POLICY DEPENDENCIES

Dependencies can be used to activate a policy or policy template when the dependency criteria are satisfied. The following fields are checked on the managed cluster, **dependencies** and

extraDependencies. When a dependency is not met, the template status of the replicated policy template displays more details.

Required access: Policy administrator

View the following policy dependency example, where the **ScanSettingBinding** is only created if the **upstream-compliance-operator** policy is already compliant on the managed cluster:

```

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: moderate-compliance-scan
  namespace: default
spec:
  dependencies:
  - apiVersion: policy.open-cluster-management.io/v1
    compliance: Compliant
    kind: Policy
    name: upstream-compliance-operator
    namespace: default
  disabled: false
  policy-templates:
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name: moderate-compliance-scan
      spec:
        object-templates:
        - complianceType: musthave
          objectDefinition:
            apiVersion: compliance.openshift.io/v1alpha1
            kind: ScanSettingBinding
            metadata:
              name: moderate
              namespace: openshift-compliance
            profiles:
            - apiGroup: compliance.openshift.io/v1alpha1
              kind: Profile
              name: ocp4-moderate
            - apiGroup: compliance.openshift.io/v1alpha1
              kind: Profile
              name: ocp4-moderate-node
            settingsRef:
              apiGroup: compliance.openshift.io/v1alpha1
              kind: ScanSetting
              name: default
            remediationAction: enforce
            severity: low

```

Note: A dependency cannot be used to apply a policy on one cluster based on the status of a policy in another cluster.

2.8. SECURE THE HUB CLUSTER

Secure your Red Hat Advanced Cluster Management for Kubernetes installation by enhancing the hub cluster security. Complete the following steps:

1. Secure Red Hat OpenShift Container Platform. For more information, see [OpenShift Container Platform security and compliance](#).
2. Setup role-based access control (RBAC). For more information, see [Role-based access control](#).
3. Customize certificates, see [Certificates](#).
4. Define your cluster credentials, see [Managing credentials overview](#)
5. Review the policies that are available to help you harden your cluster security. See [Supported policies](#)

2.9. INTEGRATE THIRD-PARTY POLICY CONTROLLERS

Integrate third-party policies to create custom annotations within the policy templates to specify one or more compliance standards, control categories, and controls.

You can also use the third-party party policies from the [policy-collection/community](#).

Learn to integrate the following third-party policies:

- [Integrating gatekeeper constraints and constraint templates](#)
- [Policy Generator](#)
 - [Generating a policy to install an Operator](#)

2.9.1. Integrating gatekeeper constraints and constraint templates

Gatekeeper is a validating webhook that enforces custom resource definition (CRD) based policies that are run with the Open Policy Agent (OPA). You can install gatekeeper on your cluster by using the gatekeeper operator policy. Gatekeeper policy can be used to evaluate Kubernetes resource compliance. You can leverage a OPA as the policy engine, and use Rego as the policy language.

Prerequisite: You must have a Red Hat Advanced Cluster Management for Kubernetes or Red Hat OpenShift Container Platform Plus subscription to install Gatekeeper and apply Gatekeeper policies to your cluster. Gatekeeper is supported only on OpenShift Container Platform versions, except version 3.11, that are supported by the latest version of Red Hat Advanced Cluster Management.

The gatekeeper policy is created as a Kubernetes configuration policy in Red Hat Advanced Cluster Management. Gatekeeper policies include constraint templates (**ConstraintTemplates**) and **Constraints**, audit templates, and admission templates. For more information, see the [Gatekeeper upstream repository](#).

Red Hat Advanced Cluster Management supports version 3.3.0 for Gatekeeper and applies the following constraint templates in your Red Hat Advanced Cluster Management gatekeeper policy:

- **ConstraintTemplates** and constraints: Use the **policy-gatekeeper-k8srequiredlabels** policy to create a gatekeeper constraint template on the managed cluster.

apiVersion: policy.open-cluster-management.io/v1

```

kind: ConfigurationPolicy
metadata:
  name: policy-gatekeeper-k8srequiredlabels
spec:
  remediationAction: enforce # will be overridden by remediationAction in parent policy
  severity: low
  object-templates:
    - complianceType: musthave
      objectDefinition:
        apiVersion: templates.gatekeeper.sh/v1beta1
        kind: ConstraintTemplate
        metadata:
          name: k8srequiredlabels
        spec:
          crd:
            spec:
              names:
                kind: K8sRequiredLabels
              validation:
                # Schema for the `parameters` field
                openAPIV3Schema:
                  properties:
                    labels:
                      type: array
                      items: string
          targets:
            - target: admission.k8s.gatekeeper.sh
              rego: |
                package k8srequiredlabels
                violation[{"msg": msg, "details": {"missing_labels": missing}}] {
                  provided := {label | input.review.object.metadata.labels[label]}
                  required := {label | label := input.parameters.labels[_]}
                  missing := required - provided
                  count(missing) > 0
                  msg := sprintf("you must provide labels: %v", [missing])
                }
            - complianceType: musthave
              objectDefinition:
                apiVersion: constraints.gatekeeper.sh/v1beta1
                kind: K8sRequiredLabels
                metadata:
                  name: ns-must-have-gk
                spec:
                  match:
                    kinds:
                      - apiGroups: [""]
                        kinds: ["Namespace"]
                    namespaces:
                      - e2etestsuccess
                      - e2etestfail
                  parameters:
                    labels: ["gatekeeper"]

```

- audit template: Use the **policy-gatekeeper-audit** to periodically check and evaluate existing resources against the gatekeeper policies that are enforced to detect existing misconfigurations.

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-gatekeeper-audit
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: low
  object-templates:
    - complianceType: musthave
      objectDefinition:
        apiVersion: constraints.gatekeeper.sh/v1beta1
        kind: K8sRequiredLabels
        metadata:
          name: ns-must-have-gk
        status:
          totalViolations: 0

```

- admission template: Use the **policy-gatekeeper-admission** to check for misconfigurations that are created by the gatekeeper admission webhook:

```

apiVersion: policy.open-cluster-management.io/v1
kind: ConfigurationPolicy
metadata:
  name: policy-gatekeeper-admission
spec:
  remediationAction: inform # will be overridden by remediationAction in parent policy
  severity: low
  object-templates:
    - complianceType: mustnothave
      objectDefinition:
        apiVersion: v1
        kind: Event
        metadata:
          namespace: openshift-gatekeeper-system # set it to the actual namespace where gatekeeper is running if different
        annotations:
          constraint_action: deny
          constraint_kind: K8sRequiredLabels
          constraint_name: ns-must-have-gk
          event_type: violation

```

- For more details, see [policy-gatekeeper-sample.yaml](#).
- For more details, see [What is OPA Gatekeeper](#).
- See [Managing configuration policies](#) for more information about managing other policies.

Refer to [Governance](#) for more topics on the security framework.

2.9.2. Policy Generator

The Policy Generator is a part of the Red Hat Advanced Cluster Management for Kubernetes application lifecycle subscription GitOps workflow that generates Red Hat Advanced Cluster Management policies using Kustomize. The Policy Generator builds Red Hat Advanced Cluster Management policies from Kubernetes manifest YAML files, which are provided through a

PolicyGenerator manifest YAML file that is used to configure it. The Policy Generator is implemented as a Kustomize generator plug-in. For more information on Kustomize, see the [Kustomize documentation](#).

The Policy Generator version bundled in this version of Red Hat Advanced Cluster Management is v1.10.0. View the following topics to for more information:

- [Policy Generator capabilities](#)
- [Policy Generator configuration structure](#)
- [Policy Generator configuration reference table](#)
- [Related resources](#)

2.9.2.1. Policy Generator capabilities

The Policy Generator and its integration with the Red Hat Advanced Cluster Management application lifecycle subscription GitOps workflow simplifies the distribution of Kubernetes resource objects to managed OpenShift Container Platform clusters, and Kubernetes clusters through Red Hat Advanced Cluster Management policies. In particular, use the Policy Generator to complete the following actions:

- Convert any Kubernetes manifest files to Red Hat Advanced Cluster Management configuration policies, including manifests created from a Kustomize directory.
- Patch the input Kubernetes manifests before they are inserted into a generated Red Hat Advanced Cluster Management policy.
- Generate additional configuration policies to be able to report on Gatekeeper policy violations through Red Hat Advanced Cluster Management for Kubernetes.
- Generate policy sets on the hub cluster.

2.9.2.2. Policy Generator configuration structure

The Policy Generator is a Kustomize generator plug-in that is configured with a manifest of the **PolicyGenerator** kind and **policy.open-cluster-management.io/v1** API version.

To use the plug-in, start by adding a **generators** section in a **kustomization.yaml** file. View the following example:

```
generators:
- policy-generator-config.yaml
```

The **policy-generator-config.yaml** file referenced in the previous example is a YAML file with the instructions of the policies to generate. A simple **PolicyGenerator** configuration file might resemble the following example:

```
apiVersion: policy.open-cluster-management.io/v1
kind: PolicyGenerator
metadata:
  name: config-data-policies
policyDefaults:
  namespace: policies
  policySets: []
policies:
```

- name: config-data
- manifests:
 - path: configmap.yaml

The **configmap.yaml** represents a Kubernetes manifest YAML file to be included in the policy. Alternatively, you can set the path to a Kustomize directory, or a directory with multiple Kubernetes manifest YAML files. View the following example:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: my-config
  namespace: default
data:
  key1: value1
  key2: value2
```

The generated **Policy**, along with the generated **PlacementRule** and **PlacementBinding** might resemble the following example:

```
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-config-data
  namespace: policies
spec:
  clusterConditions:
  - status: "True"
    type: ManagedClusterConditionAvailable
  clusterSelector:
    matchExpressions: []
---
apiVersion: policy.open-cluster-management.io/v1
kind: PlacementBinding
metadata:
  name: binding-config-data
  namespace: policies
placementRef:
  apiGroup: apps.open-cluster-management.io
  kind: PlacementRule
  name: placement-config-data
subjects:
- apiGroup: policy.open-cluster-management.io
  kind: Policy
  name: config-data
---
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: config-data
  namespace: policies
```

```

spec:
  disabled: false
  policy-templates:
  - objectDefinition:
      apiVersion: policy.open-cluster-management.io/v1
      kind: ConfigurationPolicy
      metadata:
        name: config-data
      spec:
        object-templates:
        - complianceType: musthave
          objectDefinition:
            apiVersion: v1
            data:
              key1: value1
              key2: value2
            kind: ConfigMap
            metadata:
              name: my-config
              namespace: default
            remediationAction: inform
            severity: low

```

2.9.2.3. Policy Generator configuration reference table

Note that all the fields in the **policyDefaults** section except for **namespace** can be overridden for each policy.

Table 2.14. Parameter table

| Field | Optional or required | Description |
|--------------------------------------|----------------------|--|
| apiVersion | Required | Set the value to policy.open-cluster-management.io/v1 . |
| kind | Required | Set the value to PolicyGenerator to indicate the type of policy. |
| metadata.name | Required | The name for identifying the policy resource. |
| placementBindingDefaults.name | Optional | If multiple policies use the same placement, this name is used to generate a unique name for the resulting PlacementBinding , binding the placement with the array of policies that reference it. |
| policyDefaults | Required | Any default value listed here is overridden by an entry in the policies array except for namespace . |

| Field | Optional or required | Description |
|--|----------------------|--|
| policyDefaults.namespace | Required | The namespace of all the policies. |
| policyDefaults.complianceType | Optional | Determines the policy controller behavior when comparing the manifest to objects on the cluster. The values that you can use are musthave , mustonlyhave , or mustnothave . The default value is musthave . |
| policyDefaults.metadataComplianceType | Optional | Overrides complianceType when comparing the manifest metadata section to objects on the cluster. The values that you can use are musthave , and mustonlyhave . The default value is empty (<code>{}</code>) to avoid overriding the complianceType for metadata. |
| policyDefaults.categories | Optional | Array of categories to be used in the policy.open-cluster-management.io/categories annotation. The default value is CM Configuration Management . |
| policyDefaults.controls | Optional | Array of controls to be used in the policy.open-cluster-management.io/controls annotation. The default value is CM-2 Baseline Configuration . |
| policyDefaults.standards | Optional | An array of standards to be used in the policy.open-cluster-management.io/standards annotation. The default value is NIST SP 800-53 . |
| policyDefaults.policyAnnotations | Optional | Annotations that the policy includes in the metadata.annotations section. It is applied for all policies unless specified in the policy. The default value is empty (<code>{}</code>). |

| Field | Optional or required | Description |
|--|--|--|
| policyDefaults.configurationPolicyAnnotations | Optional | Key-value pairs of annotations to set on generated configuration policies. For example, you can disable policy templates by defining the following parameter: {"policy.open-cluster-management.io/disable-templates": "true"} . The default value is empty ({}). |
| policyDefaults.severity | Optional | The severity of the policy violation. The default value is low . |
| policyDefaults.disabled | Optional | Whether the policy is disabled, meaning it is not propagated and no status as a result. The default value is false to enable the policy. |
| policyDefaults.remediationAction | Optional | The remediation mechanism of your policy. The parameter values are enforce and inform . The default value is inform . |
| policyDefaults.namespaceSelector | Required for namespaced objects that do not have a namespace specified | Determines namespaces in the managed cluster that the object is applied to. The include and exclude parameters accept file path expressions to include and exclude namespaces by name. The matchExpressions and matchLabels parameters specify namespaces to include by label. See the Kubernetes labels and selectors documentation. The resulting list is compiled by using the intersection of results from all parameters. |

| Field | Optional or required | Description |
|---|----------------------|---|
| policyDefaults.evaluationInterval | Optional | Use the parameters compliant and noncompliant to specify the frequency for a policy to be evaluated when in a particular compliance state. When managed clusters have low CPU resources, the evaluation interval can be increased to reduce CPU usage on the Kubernetes API. These are in the format of durations. For example, " 1h25m3s " represents 1 hour, 25 minutes, and 3 seconds. These can also be set to "never" to avoid evaluating the policy after it has become a particular compliance state. |
| policyDefaults.dependencies | Optional | A list of objects that must be in specific compliance states before this policy is applied. Cannot be specified when policyDefaults.orderPolicies is set to true . |
| policyDefaults.dependencies[].name | Required | The name of the object being depended on. |
| policyDefaults.dependencies[].namespace | Optional | The namespace of the object being depended on. The default is the namespace of policies set for the Policy Generator. |
| policyDefaults.dependencies[].compliance | Optional | The compliance state the object needs to be in. The default value is Compliant . |
| policyDefaults.dependencies[].kind | Optional | The kind of the object. By default, the kind is set to Policy , but can also be other kinds that have compliance state, such as ConfigurationPolicy . |
| policyDefaults.dependencies[].apiVersion | Optional | The API version of the object. The default value is policy.open-cluster-management.io/v1 . |

| Field | Optional or required | Description |
|--|----------------------|---|
| policyDefaults.extraDependencies | Optional | A list of objects that must be in specific compliance states before this policy is applied. The dependencies that you define are added to each policy template (for example, ConfigurationPolicy) separately from the dependencies list. Cannot be specified when policyDefaults.orderManifests is set to true . |
| policyDefaults.extraDependencies[].name | Required | The name of the object being depended on. |
| policyDefaults.extraDependencies[].namespace | Optional | The namespace of the object being depended on. By default, the value is set to the namespace of policies set for the Policy Generator. |
| policyDefaults.extraDependencies[].compliance | Optional | The compliance state the object needs to be in. The default value is Compliant . |
| policyDefaults.extraDependencies[].kind | Optional | The kind of the object. The default value is to Policy , but can also be other kinds that have a compliance state, such as ConfigurationPolicy . |
| policyDefaults.extraDependencies[].apiVersion | Optional | The API version of the object. The default value is policy.open-cluster-management.io/v1 . |
| policyDefaults.ignorePending | Optional | Bypass compliance status checks when the Policy Generator is waiting for its dependencies to reach their desired states. The default value is false . |

| Field | Optional or required | Description |
|--|----------------------|--|
| policyDefaults.orderPolicies | Optional | Automatically generate dependencies on the policies so they are applied in the order you defined in the policies list. By default, the value is set to false . Cannot be specified at the same time as policyDefaults.dependencies . |
| policyDefaults.orderManifests | Optional | Automatically generate extraDependencies on policy templates so that they are applied in the order you defined in the manifests list for that policy. Cannot be specified when policyDefaults consolidateManifests is set to true . Cannot be specified at the same time as policyDefaults.extraDependencies . |
| policyDefaults consolidateManifests | Optional | This determines if a single configuration policy is generated for all the manifests being wrapped in the policy. If set to false , a configuration policy per manifest is generated. The default value is true . |
| policyDefaults.informGatekeeperPolicies | Optional | When the policy references a violated gatekeeper policy manifest, this determines if an additional configuration policy is generated in order to receive policy violations in Red Hat Advanced Cluster Management. The default value is true . |
| policyDefaults.informKyvernoPolicies | Optional | When the policy references a Kyverno policy manifest, this determines if an additional configuration policy is generated to receive policy violations in Red Hat Advanced Cluster Management, when the Kyverno policy is violated. The default value is true . |

| Field | Optional or required | Description |
|--|----------------------|--|
| policyDefaults.policySets | Optional | Array of policy sets that the policy joins. Policy set details can be defined in the policySets section. When a policy is part of a policy set, a placement binding is not generated for the policy since one is generated for the set. Set policies[].generatePlacementWhenInSet or policyDefaults.generatePlacementWhenInSet to override policyDefaults.policySets . |
| policyDefaults.generatePlacementWhenInSet | Optional | When a policy is part of a policy set, by default, the generator does not generate the placement for this policy since a placement is generated for the policy set. Set generatePlacementWhenInSet to true to deploy the policy with both policy placement and policy set placement. The default value is false . |
| policyDefaults.placement | Optional | The placement configuration for the policies. This defaults to a placement configuration that matches all clusters. |
| policyDefaults.placement.name | Optional | Specifying a name to consolidate placement rules that contain the same cluster selectors. |
| policyDefaults.placement.placementName | Optional | Define this parameter to use a placement that already exists on the cluster. A Placement is not created, but a PlacementBinding binds the policy to this Placement . |
| policyDefaults.placement.placementPath | Optional | To reuse an existing placement, specify the path here relative to the kustomization.yaml file. If provided, this placement rule is used by all policies by default. See clusterSelectors to generate a new Placement . |

| Field | Optional or required | Description |
|---|----------------------|---|
| policyDefaults.placement.clusterSelectors | Optional | Specify a placement by defining a cluster selector in the following format, key:value . See placementPath to specify an existing file. |
| policyDefaults.placement.placementRuleName | Optional | To use a placement rule that already exists on the cluster, specify its name here. A PlacementRule is not created, but a PlacementBinding binds the policy to this PlacementRule . |
| policyDefaults.placement.placementRulePath | Optional | To reuse an existing placement rule, specify the path here relative to the kustomization.yaml file. If provided, this placement rule is used by all policies by default. See labelSelector to generate a new PlacementRule . |
| policyDefaults.placement.labelSelector | Optional | Specify a placement rule by defining a cluster selector in the following format, key:value . See placementRulePath to specify an existing file. |
| policies | Required. | The list of policies to create along with overrides to either the default values, or the values that are set in policyDefaults . See policyDefaults for additional fields and descriptions. |
| policies[].name | Required | The name of the policy to create. |

| Field | Optional or required | Description |
|---------------------------------------|----------------------|--|
| policies[].manifests | Required | The list of Kubernetes object manifests to include in the policy, along with overrides to either the default values, the values set in this policies item, or the values set in policyDefaults . See policyDefaults for additional fields and descriptions. When consolidateManifests is set to true , only complianceType and metadataComplianceType can be overridden at the policies[].manifests level. |
| policies[].manifests[].path | Required | Path to a single file, a flat directory of files, or a Kustomize directory relative to the kustomization.yaml file. If the directory is a Kustomize directory, the generator runs Kustomize against the directory before generating the policies. |
| policies[].manifests[].patches | Optional | A list of Kustomize patches to apply to the manifest at the path. If there are multiple manifests, the patch requires the apiVersion , kind , metadata.name , and metadata.namespace (if applicable) fields to be set so Kustomize can identify the manifest that the patch applies to. If there is a single manifest, the metadata.name and metadata.namespace fields can be patched. |
| policySets | Optional | The list of policy sets to create. To include a policy in a policy set, use policyDefaults.policySets , policies[].policySets , or policySets.policies . |
| policySets[].name | Required | The name of the policy set to create. |
| policySets[].description | Optional | The description of the policy set to create. |

| Field | Optional or required | Description |
|-------------------------------|----------------------|--|
| policySets[].policies | Optional | The list of policies to be included in the policy set. If policyDefaults.policySets or policies[].policySets is also specified, the lists are merged. |
| policySets[].placement | Optional | The placement configuration for the policy set. This defaults to a placement configuration that matches all clusters. See policyDefaults.placement for placement documentation, however policyDefaults.placement settings do not apply to policy sets. |

2.9.2.4. Related resources

- See [Generating a policy to install an Operator](#)
- See the [policy-generator-plugin](#) repository for more details.
- See [Policy set controller](#) for more details.
- Return to the [Integrate third-party policy controllers](#) documentation.
- Refer to the [Governance](#) documentation for more topics.
- See [Applying Kustomize](#) for more information.
- See an example of a [kustomization.yaml](#) file.
- See [Gatekeeper](#) for more details.

2.9.3. Generating a policy to install an Operator

A common use of Red Hat Advanced Cluster Management policies is to install an Operator on one or more managed Red Hat OpenShift Container Platform clusters. Continue reading to learn how to generate policies using the Policy Generator, and to install the OpenShift GitOps Operator with a generated policy:

- [Generating a policy that installs OpenShift GitOps](#)
- [Generating a policy that installs the Compliance Operator](#)
- [Using policy dependencies with OperatorGroups](#)
- [Related resources](#)

2.9.3.1. Generating a policy that installs OpenShift GitOps

You can generate a policy that installs OpenShift GitOps by using the Policy Generator. The OpenShift GitOps operator offers the *all namespaces* installation mode, which is used in the following example. Create a **Subscription** manifest file called **openshift-gitops-subscription.yaml**, similar to the following example:

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: openshift-gitops-operator
  namespace: openshift-operators
spec:
  channel: stable
  name: openshift-gitops-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

To pin to a specific version of the operator, add the following parameter and value: **spec.startingCSV: openshift-gitops-operator.v<version>**. Replace **<version>** with your preferred version.

A **PolicyGenerator** configuration file is required. Use the configuration file named **policy-generator-config.yaml** to generate a policy to install OpenShift GitOps on all OpenShift Container Platform managed clusters. See the following example:

```
apiVersion: policy.open-cluster-management.io/v1
kind: PolicyGenerator
metadata:
  name: install-openshift-gitops
policyDefaults:
  namespace: policies
  placement:
    clusterSelectors:
      vendor: "OpenShift"
  remediationAction: enforce
policies:
  - name: install-openshift-gitops
    manifests:
      - path: openshift-gitops-subscription.yaml
```

The last required file is **kustomization.yaml**, which requires the following configuration:

```
generators:
  - policy-generator-config.yaml
```

The generated policy might resemble the following file:

```
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-install-openshift-gitops
  namespace: policies
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-openshift-gitops
- namespace: policies

```
placementRef:
```

- apiGroup: apps.open-cluster-management.io
- kind: PlacementRule
- name: placement-install-openshift-gitops

```
subjects:
```

- apiGroup: policy.open-cluster-management.io
- kind: Policy
- name: install-openshift-gitops

```
---
```

```
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-openshift-gitops
```

```
namespace: policies
```

```
spec:
```

```
disabled: false
```

```
policy-templates:
```

- objectDefinition:

```
  apiVersion: policy.open-cluster-management.io/v1
```

```
  kind: ConfigurationPolicy
```

```
  metadata:
```

```
    name: install-openshift-gitops
```

```
  spec:
```

```
    object-templates:
```

- complianceType: musthave

```
    objectDefinition:
```

```
      apiVersion: operators.coreos.com/v1alpha1
```

```
      kind: Subscription
```

```
      metadata:
```

```
        name: openshift-gitops-operator
```

```
        namespace: openshift-operators
```

```
      spec:
```

```
        channel: stable
```

```
        name: openshift-gitops-operator
```

```
        source: redhat-operators
```

```
        sourceNamespace: openshift-marketplace
```

```
    remediationAction: enforce
```

```
    severity: low
```

Generated policies from manifests in the OpenShift Container Platform documentation is supported. Any configuration guidance from the OpenShift Container Platform documentation can be applied using the Policy Generator.

2.9.3.2. Generating a policy that installs the Compliance Operator

For an operator that uses the *namespaced* installation mode, such as the Compliance Operator, an **OperatorGroup** manifest is also required.

Create a YAML file with a **Namespace**, a **Subscription**, and an **OperatorGroup** manifest called **compliance-operator.yaml**. The following example installs these manifests in the **compliance-operator** namespace:

```
apiVersion: v1
kind: Namespace
metadata:
  name: openshift-compliance
---
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: compliance-operator
  namespace: openshift-compliance
spec:
  targetNamespaces:
    - compliance-operator
---
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: compliance-operator
  namespace: openshift-compliance
spec:
  channel: release-0.1
  name: compliance-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

A **PolicyGenerator** configuration file is required. View the following **PolicyGenerator** policy example that installs the Compliance Operator on all OpenShift Container Platform managed clusters:

```
apiVersion: policy.open-cluster-management.io/v1
kind: PolicyGenerator
metadata:
  name: install-compliance-operator
policyDefaults:
  namespace: policies
  placement:
    clusterSelectors:
      vendor: "OpenShift"
  remediationAction: enforce
policies:
  - name: install-compliance-operator
    manifests:
      - path: compliance-operator.yaml
```

The last required file is **kustomization.yaml**, which requires the following configuration:

```
generators:
  - policy-generator-config.yaml
```

As a result, the generated policy resembles the following file:

```
apiVersion: apps.open-cluster-management.io/v1
kind: PlacementRule
metadata:
  name: placement-install-compliance-operator
  namespace: policies
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-compliance-operator
  namespace: policies
placementRef:
  apiGroup: apps.open-cluster-management.io
  kind: PlacementRule
  name: placement-install-compliance-operator
subjects:
  - apiGroup: policy.open-cluster-management.io
    kind: Policy
    name: install-compliance-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-compliance-operator
  namespace: policies
spec:
  disabled: false
  policy-templates:
    - objectDefinition:
        apiVersion: policy.open-cluster-management.io/v1
        kind: ConfigurationPolicy
        metadata:
          name: install-compliance-operator
        spec:
```

```

object-templates:
- complianceType: musthave
  objectDefinition:
    apiVersion: v1
    kind: Namespace
    metadata:
      name: openshift-compliance
- complianceType: musthave
  objectDefinition:
    apiVersion: operators.coreos.com/v1alpha1
    kind: Subscription
    metadata:
      name: compliance-operator
      namespace: openshift-compliance
    spec:
      channel: release-0.1
      name: compliance-operator
      source: redhat-operators
      sourceNamespace: openshift-marketplace
- complianceType: musthave
  objectDefinition:
    apiVersion: operators.coreos.com/v1
    kind: OperatorGroup
    metadata:
      name: compliance-operator
      namespace: openshift-compliance
    spec:
      targetNamespaces:
        - compliance-operator
  remediationAction: enforce
  severity: low

```

2.9.3.3. Using policy dependencies with *OperatorGroups*

When you install an operator with an **OperatorGroup** manifest, the **OperatorGroup** must exist on the cluster before the **Subscription** is created. Use the policy dependency feature along with the Policy Generator to ensure that the **OperatorGroup** policy is compliant prior to enforcing the **Subscription** policy.

Set up policy dependencies by listing the manifests in the order that you want. For example, you might want to create the namespace policy first, create the **OperatorGroup** next, and create the **Subscription** last.

Enable the **policyDefaults.orderManifests** parameter and disable **policyDefaults consolidateManifests** in the Policy Generator configuration manifest to automatically set up dependencies between the manifests.

2.9.3.4. Related resources

- See [Deploying policies using GitOps](#) for more details.
- Return to the [Integrate third-party policy controllers](#) documentation.
- See [Understanding OpenShift GitOps](#) and the [Operator](#) documentation for more details.
- See [Adding Operators to a cluster - Installing from OperatorHub using the CLI](#)

- See the [Compliance Operator documentation](#) for more details.
- See [all namespaces installation mode](#).
- See [namespaced installation mode](#).
- See [Using Init Containers to perform tasks before a pod is deployed](#).
- See [ArgoCD](#).
- View the following examples of YAML input that is supported in the OpenShift Container Platform documentation:
 - [Post-installation cluster tasks](#)
 - [Configuring the audit log policy](#)
 - [About forwarding logs to third-party systems](#)

2.9.4. Managing policy definitions with OpenShift GitOps (ArgoCD)

OpenShift GitOps, based on ArgoCD, can also be used to manage policy definitions. To allow this workflow, OpenShift GitOps must be granted access to create policies on the Red Hat Advanced Cluster Management hub cluster. Create the following **ClusterRole** resource called **openshift-gitops-policy-admin**, with access to create, read, update, and delete policies and placements. Your **ClusterRole** might resemble the following example:

```
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: openshift-gitops-policy-admin
rules:
- verbs:
  - get
  - list
  - watch
  - create
  - update
  - patch
  - delete
  apiGroups:
  - policy.open-cluster-management.io
  resources:
  - policies
  - placementbindings
- verbs:
  - get
  - list
  - watch
  - create
  - update
  - patch
  - delete
  apiGroups:
  - apps.open-cluster-management.io
  resources:
```

```

- placementrules
- verbs:
  - get
  - list
  - watch
  - create
  - update
  - patch
  - delete
apiGroups:
  - cluster.open-cluster-management.io
resources:
  - placements
  - placements/status
  - placementdecisions
  - placementdecisions/status

```

Create a **ClusterRoleBinding** object to grant the OpenShift GitOps service account access to the **openshift-gitops-policy-admin ClusterRole** object. Your **ClusterRoleBinding** might resemble the following example:

```

kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: openshift-gitops-policy-admin
subjects:
  - kind: ServiceAccount
    name: openshift-gitops-argocd-application-controller
    namespace: openshift-gitops
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: openshift-gitops-policy-admin

```

When a Red Hat Advanced Cluster Management policy definition is deployed with OpenShift GitOps, a copy of the policy is created in each managed cluster namespace. These copies are called replicated policies. To prevent OpenShift GitOps from repeatedly deleting this replicated policy or show that the ArgoCD **Application** is out of sync, the **argocd.argoproj.io/compare-options: ignoreExtraneous** annotation is automatically set on each replicated policy by the Red Hat Advanced Cluster Management policy framework.

There are labels and annotations used by ArgoCD to track objects. For replicated policies to not show up at all in ArgoCD, you can set **spec.copyPolicyMetadata** to **false** on the Red Hat Advanced Cluster Management policy definition to prevent these ArgoCD tracking labels and annotations from being copied to the replicated policy.

2.9.4.1. Integrating the Policy Generator with OpenShift GitOps (ArgoCD)

OpenShift GitOps, based on ArgoCD, can also be used to generate policies by using the Policy Generator through GitOps. Since the Policy Generator does not come preinstalled in the OpenShift GitOps container image, some customization must take place. In order to follow along, you must have the OpenShift GitOps Operator installed on the Red Hat Advanced Cluster Management hub cluster and be sure to log into the hub cluster.

In order for OpenShift GitOps to have access to the Policy Generator when you run Kustomize, an Init

Container is required to copy the Policy Generator binary from the Red Hat Advanced Cluster Management Application Subscription container image to the OpenShift GitOps container. Additionally, OpenShift GitOps must be configured to provide the **--enable-alpha-plugins** flag when it runs Kustomize. Start editing the OpenShift GitOps **argocd** object with the following command:

```
oc -n openshift-gitops edit argocd openshift-gitops
```

Then modify the OpenShift GitOps **argocd** object to contain the following additional YAML content. When a new major version of Red Hat Advanced Cluster Management is released and you want to update the Policy Generator to a newer version, you need to update the **registry.redhat.io/rhacm2/multicluster-operators-subscription-rhel8** image used by the Init Container to a newer tag. View the following example and replace **<version>** with 2.7 or your desired Red Hat Advanced Cluster Management version:

```
apiVersion: argoproj.io/v1alpha1
kind: ArgoCD
metadata:
  name: openshift-gitops
  namespace: openshift-gitops
spec:
  kustomizeBuildOptions: --enable-alpha-plugins
  repo:
    env:
      - name: KUSTOMIZE_PLUGIN_HOME
        value: /etc/kustomize/plugin
    initContainers:
      - args:
          - -c
          - cp /etc/kustomize/plugin/policy.open-cluster-management.io/v1/policygenerator/PolicyGenerator
            /policy-generator/PolicyGenerator
        command:
          - /bin/bash
        image: registry.redhat.io/rhacm2/multicluster-operators-subscription-rhel8:v<version>
        name: policy-generator-install
        volumeMounts:
          - mountPath: /policy-generator
            name: policy-generator
        volumeMounts:
          - mountPath: /etc/kustomize/plugin/policy.open-cluster-management.io/v1/policygenerator
            name: policy-generator
        volumes:
          - emptyDir: {}
            name: policy-generator
```

Now that OpenShift GitOps can use the Policy Generator, OpenShift GitOps must be granted access to create policies on the Red Hat Advanced Cluster Management hub cluster. Create the **ClusterRole** resource called **openshift-gitops-policy-admin**, with access to create, read, update, and delete policies and placements. Refer to the earlier **ClusterRole** example.

Additionally, create a **ClusterRoleBinding** object to grant the OpenShift GitOps service account access to the **openshift-gitops-policy-admin ClusterRole**. Your **ClusterRoleBinding** might resemble the following resource:

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
```

```
metadata:  
  name: openshift-gitops-policy-admin  
subjects:  
  - kind: ServiceAccount  
    name: openshift-gitops-argocd-application-controller  
    namespace: openshift-gitops  
roleRef:  
  apiGroup: rbac.authorization.k8s.io  
  kind: ClusterRole  
  name: openshift-gitops-policy-admin
```

2.9.4.2. Related resource

- See [Using policy dependencies with *OperatorGroups*](#).
- Return to [Generating a policy to install an Operator](#) .