

Red Hat OpenShift Service on AWS cluster configuration guide

Detail

Table of contents

I. Introduction	2
A. Congratulations on creating your first cluster	2
B. Purpose of the guide	2
II. Getting started with your Red Hat OpenShift cluster	2
A. Cluster configuration	2
B. Accessing Red Hat Hybrid Cloud Console	2
C. Setting up an identity provider (IdP)	3
D. Deploying applications	3
E. Configuring continuous integration and continuous delivery (CI/CD) pipelines	3
III. Key features and best practices	3
A. Monitoring and logging	3
B. Autoscaling	4
C. Backup and disaster recovery	4
D. Upgrades and updates	4
IV. Support, documentation, and training	4
A. Red Hat Customer Portal	4
B. Red Hat OpenShift Service on AWS product documentation	5
C. Training courses, workshops, and certifications	5
V. Common pitfalls and troubleshooting	5
A. Connectivity issues	5
B. Permissions and access control	5
C. Application deployment errors	5
D. Unexpected costs	5
VI. Community and forums	5
A. Red Hat OpenShift Commons	5

B. Slack.....	6
C. AWS Communities	6
D. GitHub	6
VII. Next steps for long-term success.....	6
A. Performance optimization.....	6
B. Security and compliance.....	6
C. Cost optimization.....	6
D. Governance and continuous improvement.....	6
VIII. Conclusion.....	7
A. Encouragement and commitment to customer success	7
B. Contact information for questions and support.....	7

I. Introduction

Congratulations on creating your first cluster. You have taken the first step in unlocking the potential of Red Hat® OpenShift® Service on AWS. This service allows you to fully take advantage of the scalability, reliability, and agility of Amazon Web Services (AWS) while using the powerful container orchestration capabilities of Red Hat OpenShift. In the following sections, we will guide you through the key steps to start deploying applications and optimizing your Red Hat OpenShift environment.

Purpose of the guide

This guide is intended to help you understand how to effectively use OpenShift Service on AWS. We will walk you through getting started, introduce key features and best practices, and outline steps for long-term success. This guide will serve as a valuable resource while you are using OpenShift Service on AWS.

The [OpenShift Service on AWS product documentation](#) on the [Red Hat Customer Portal](#) will be an important resource as you progress through the steps we outline below, so be sure to keep it handy.

II. Getting started with your Red Hat OpenShift cluster

A. Cluster configuration

Once your cluster is created, you will need to configure it to suit your specific needs. This involves setting up user access controls, defining resource quotas, installing any necessary operators, and more. To get the most out of your Red Hat OpenShift cluster, it is crucial to understand the [configuration options](#).

B. Accessing Red Hat Hybrid Cloud Console and Red Hat OpenShift Cluster Manager

[Red Hat Hybrid Cloud Console](#) provides a visual interface to interact with your cluster called [Red Hat OpenShift Cluster Manager](#). To access it, navigate to the web console URL provided after your cluster is created. [Log in](#) with your Red Hat credentials. The console provides a graphical interface that allows you to interact with your cluster and manage its resources.

C. Setting up an identity provider

An [identity provider \(IdP\)](#) serves a pivotal role in the Red Hat OpenShift ecosystem by managing user authentication and permissions. It verifies users' identities and controls their access rights to your Red Hat OpenShift cluster. This function establishes an essential layer of security, helping to ensure that only authorized individuals can access and manipulate the cluster. During the setup of your identity provider, ensuring that the correct users are granted access to your Red Hat OpenShift cluster is crucial. Carefully consider which individuals or groups within your organization require access, taking into account their roles and responsibilities. This step forms the foundation of your access control strategy, promoting security-focused and efficient operations in your OpenShift on AWS environment.

D. Deploying applications

Applications on Red Hat OpenShift are deployed as containers, which can be created from existing images or source code. Red Hat OpenShift supports a variety of programming languages, databases, and other services. You can [deploy applications on OpenShift Service on AWS](#) in several ways, including using the web console, the OpenShift command line interface (the `oc` command), and Kubernetes manifests. The OpenShift Developer perspective in the web console provides an intuitive interface for deploying and managing applications.

An important aspect of deploying applications in Red Hat OpenShift involves [configuring resource limits and requests](#). Setting appropriate resource requests and limits for your application's pods is a crucial best practice to ensure efficient utilization of your cluster's resources and maintain the stability of your services. Resource requests help the Kubernetes scheduler make more intelligent decisions about where to place pods while limits prevent single pods from depleting all of the available resources on a node.

It is equally critical to create [pod disruption budgets \(PDBs\)](#) for your deployments. PDBs limit the number of pods of a replicated application that can be simultaneously down. By setting a PDB, you can establish operational resilience with your applications, ensuring that a certain minimum number of pods are always available—even during voluntary disruptions, such as maintenance operations or during node autoscaling events. This is crucial for maintaining high availability and resilience of your application.

E. Configuring CI/CD pipelines

With [Red Hat OpenShift Pipelines](#), you can create continuous integration and continuous delivery (CI/CD) pipelines to automate the build, test, and deployment stages of your applications. You can configure these pipelines through the OpenShift Pipelines operator, which integrates Tekton—a Kubernetes-native CI/CD solution. Pipelines are defined using standard Kubernetes custom resource definitions (CRDs), and they are portable across Kubernetes distributions.

III. Key features and best practices

A. Monitoring and alerting

OpenShift Service on AWS provides several tools for [monitoring](#) and [logging](#). Regularly monitoring your cluster and applications helps identify potential issues before they affect your workloads. Monitoring helps you track your cluster's performance and health while logging helps you

troubleshoot issues. Red Hat recommends that you use the [AWS CloudWatch](#) solution for your logging needs. You can forward logs to AWS CloudWatch in addition to, or instead of, the default log store.

B. Autoscaling

To optimize resource utilization and handle workload fluctuations, OpenShift Service on AWS supports both horizontal and vertical autoscaling. [Horizontal pod autoscaler](#) (HPA) adjusts the number of pod replicas, while [vertical pod autoscaler](#) (VPA) adjusts resource limits. [Autoscaling](#) allows your applications to dynamically scale up or down based on their resource utilization. This helps ensure that your applications always have the resources they need while also optimizing your AWS resource utilization.

C. Backup and disaster recovery

Implementing a rigorous [backup](#) and disaster recovery strategy is an integral part of maintaining the health and longevity of your applications on OpenShift Service on AWS. Services through Amazon like Amazon S3 provide reliable storage solutions for your backup needs, and Red Hat OpenShift has built-in backup tools to simplify the process.

In addition to these backup measures, it is highly recommended to have deployment pipelines in place. [Deployment pipelines](#) provide an automated and consistent way to deploy and roll back applications, contributing to more precise and resilient clusters. They allow you to quickly and reliably restore your applications to a known good state in the event of a failure.

A combination of using robust backup tools, reliable storage, and precise deployment pipelines can greatly enhance your disaster recovery strategy and ensure operational continuity in the face of potential disruptions.

D. Updates

[Upgrading on OpenShift Service on AWS is an automated process](#), helping you keep your clusters up to date with the latest features and security updates. Before initiating an upgrade, ensure that your applications and data are backed up. These automated upgrades can help reduce downtime and ensure that your cluster is always running the latest and most security-hardened version of Red Hat OpenShift.

IV. Support, documentation, and training

A. Red Hat Customer Portal

The [Red Hat Customer Portal](#) is your one-stop shop for support, product documentation, and other resources, including:

- ▶ [Red Hat Knowledgebase](#): This is a comprehensive library of articles, tutorials, and solutions to common issues.
- ▶ [Case management](#): This is a system for reporting issues and tracking their resolution.
- ▶ [Troubleshooting](#): This allows you to resolve a server error message.
- ▶ [Labs](#): These are applications to help you improve performance, troubleshoot issues, identify security problems, and optimize configuration.

B. OpenShift Service on AWS documentation

[OpenShift Service on AWS product documentation](#) provides comprehensive guidance on using OpenShift Service on AWS, including detailed explanations of its features and how-to guides.

C. Training courses, workshops, and certifications

Red Hat offers a variety of training options to help you and your team get the most out of OpenShift Service on AWS. This includes online and in-person courses, hands-on workshops, and certifications. Learn how to deploy, access, and perform basic customizations to an OpenShift Service on AWS cluster with this [training course](#).

V. Common pitfalls and troubleshooting

Even the best planning and implementation cannot guarantee against encountering challenges while you use OpenShift Service on AWS. Here are some common pitfalls and troubleshooting tips:

A. Connectivity issues

If you are having trouble connecting to your cluster, ensure that your network and firewall settings allow traffic to and from the [necessary ports](#). Remember that Red Hat OpenShift requires specific ports to be open for different types of communication.

B. Permissions and access control

If you find that you are unable to perform certain actions, it may be due to your [user permissions](#). Check the roles and permissions assigned to your user account in the Red Hat OpenShift and AWS Identity Access Management (IAM) consoles.

C. Application deployment errors

Errors during application deployment can stem from a variety of issues, such as [misconfigured settings](#) or insufficient resources. Examine the error logs and events in the Red Hat OpenShift console for clues as to what went wrong.

D. Unexpected costs

If your AWS bill is higher than expected, check the resource usage in your Red Hat OpenShift cluster. Ensure that autoscaling is configured properly and that unused resources are being scaled down or terminated.

Remember, the [Red Hat Customer Portal](#) and [OpenShift Service on AWS product documentation](#) are invaluable resources for troubleshooting.

VI. Community and forums

While this guide provides a solid starting point for using OpenShift Service on AWS, the broader community can offer a wealth of knowledge and experience. Here are some forums and communities where you can ask questions, share your experiences, and learn from others:

A. Red Hat OpenShift Commons

[Red Hat OpenShift Commons](#) is a community of Red Hat OpenShift users, partners, customers, and contributors and a great place to connect with others, learn about the latest developments, and share your own experiences.

B. Slack

The [OpenShift Service on AWS community on Slack](#) is a great place to ask your burning technical questions and get answers from Red Hat product experts and the community of OpenShift Service on AWS users. The channel also occasionally features special guests.

C. AWS Communities

The [AWS Developer Communities](#) are a hub for AWS users to ask questions, share tips, and discuss best practices. There are specific forums for different AWS services, including a forum for OpenShift Service on AWS.

D. GitHub

Many OpenShift-related projects, including operators and sample applications, are hosted on [GitHub](#). Exploring these projects can provide insights into how others are using Red Hat OpenShift and OpenShift Service on AWS.

VII. Next steps for long-term success

A. Performance optimization

Regular health and performance monitoring of your cluster is conducive to optimal operation. OpenShift on AWS integrates with [AWS CloudWatch](#) and Red Hat [monitoring solutions](#) to provide comprehensive insights.

In addition, OpenShift Service on AWS includes telemetry data capabilities via [remote health monitoring](#). This data can be accessed through the [OpenShift Cluster Manager](#) and [Red Hat Insights Advisor](#). The telemetry data provides invaluable insights into the performance and utilization of your cluster, allowing for proactive identification of potential issues and suggesting improvements.

When used effectively, these insights, derived from telemetry data, can drastically improve your operational efficiency, reduce downtime, and facilitate informed decision-making. Remote health monitoring not only helps in maintaining the health of your cluster but also provides data-driven insights for optimization.

B. Security and compliance

OpenShift Service on AWS provides robust security features, including role-based access control (RBAC), security context constraints (SCC), and integration with AWS IAM. Regularly review your [security configurations](#) and compliance with industry standards.

C. Cost optimization

To manage your AWS costs effectively, monitor your resource usage and adjust your resource allocations with [autoscaling](#). Also, take advantage of savings plans and reserved instances.

D. Governance and continuous improvement

Establish clear governance policies and continually refine your operations based on your evolving needs and the latest best practices. Make use of Red Hat OpenShift's built-in policy management features and AWS's governance tools.

VIII. Conclusion

A. Encouragement and commitment to customer success

We are excited for you to embark on this journey with OpenShift Service on AWS and are committed to your success. We will continually strive to provide the best tools and support to help you achieve your goals.

B. Contact information for questions and support




For any questions or support needs, reach out to us through the [Red Hat Customer Portal](#). Our dedicated support team is ready to assist you. Remember, you can also turn to the OpenShift and AWS communities for help and advice.

Congratulations once again on setting up your Red Hat OpenShift Service on AWS cluster. With this configuration guide and the resources provided, you are well on your way to success. Happy coding.



About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

 facebook.com/redhatinc
 @RedHat
 linkedin.com/company/red-hat

North America
1 888 REDHAT1

**Europe, Middle East,
and Africa**
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com