

How to keep your OpenShift clusters running like clockwork



Christian Hubinger

cloud native

openshift

2022-11-22

Developing and deploying applications on [OpenShift](#)? Great stuff. Losing out on time to market when your clusters stop running? Not so much.

With all their incredible benefits, clustered systems do come with challenges — one of them is maintaining zero down time. But there are some tricks and tools you can implement to work around them. In this insight, we cover how to keep your OpenShift clusters running so you can scale, update, and deploy your applications disruption-free.

Why well-functioning OpenShift clusters are vital for your business

When using a streamlined platform like OpenShift, you're already on track to [speeding up the delivery of new applications to your business](#). It's likely why your tech team sees value in it: OpenShift and container technology generally automate many steps in their workload and make it quicker, easier, and more efficient.

At the same time, containerized workloads often involve applications that consist of many individual components, which run separate to one another. This can make it challenging to monitor your containers during runtime and ensure compliance. Other factors — like application failures or cluster administration errors — can also result in downtime, which means plenty of sunk costs in the long run.

But when you improve an application's uptime, you end up with happier developers and users. Increasing your cluster efficiency can help prevent outages even if application instances or their infrastructure restart or issue problems. We've covered [why investing in OpenShift is a smart business](#)

[decision](#) and how it can help you stay ahead of the competition in the tech landscape — and some steps to help ensure that your clusters are performing at their best pace.

Best practices for keeping your OpenShift clusters running smoothly

OpenShift comes with sophisticated enterprise features that support compliance and multi-cluster management — so we've put together some best practices to get you started. Here's what we've tried and tested successfully at TRIGO to keep our OpenShift clusters running like clockwork:

- Deploy multiple replicas to prevent app unavailability: If you're running a single instance of your application and it gets deleted, you might find it completely (though temporarily) unavailable. To counter this, it's a good idea to deploy at least two running replicas of the application. This ensures you won't experience downtime when, for example, you've planned for an update and one instance goes down. It'll also come in handy when you experience unintended downtime as a result of an application crash.
- Run one process per container for better efficiency: When you run each process in a separate container, you can isolate them better. This helps you avoid issues with signal routing or any need for reaping zombie processes.
- Use Rolling Update for important applications: [The Rolling Strategy](#) is a deployment strategy for changing or upgrading application users experiencing disruptions. It's your best bet if you want to avoid downtime when your application is updating. The Rolling Strategy ensures that at least one pod is running while replacing pods from previous versions of applications to new versions.
- Use a Pod Disruption Budget (PDB) for improved maintenance: [a PDB](#) is an API object that limits how many pod replicas can be taken down by a cluster for maintenance. This comes in handy when running critical applications and ensuring that at least a minimum number of pods are available during maintenance or even during disruptions.
- Use Pod Topology Spread Constraints for a better-distributed workload: [Pod Topology Spread Constraints](#) are another OpenShift feature that ensures pods are distributed across domains. This can help prevent downtime, which might occur if a domain (like a node or availability zone) grows unhealthy.

- Include resource requests and limits in your pod definitions: The OpenShift Container Platform comes equipped with [internal pod scheduling processes](#). With these, you can specify resource requests and limits to prevent your applications from running out of memory or CPU usage since your cluster can schedule them more efficiently. This ensures that your application can continue running.
- Prioritize app health and resilience: In any software system, issues like connectivity or configuration errors can lead to containers stopping functioning properly. OpenShift Container Platform applications include features like [health check probes](#) that can detect and handle unhealthy containers periodically. You can specify which pods (and containers) you want to run a health check on to allow the cluster to restart your application or avoid routine traffic if it's not ready to handle requests yet.

[Other features](#) (like circuit breakers, timeouts, retries, and rate limiting) can also help your application perform better in the case of failures. These can prevent them from getting overloaded and improve performance during connectivity issues.

Along with these tips, we'd highly recommend using application monitoring and alerting tools like [Prometheus](#) and [Grafana](#). These tools help provide quick updates for the speed and health of applications in production. You could also give [OpenShift Service Mesh](#) a go, which lets you implement resilience measures without changing any of your application code.

If you don't have a dedicated platform team to maintain your OpenShift cluster, check out [TRIGO's OpenShift Consulting and Operating services](#), which we cover a little further below. Till then, here's a product that helps you manage your clusters with minimal stress: [Red Hat Hybrid Cloud Console Services](#), including its [OpenShift Cluster Manager](#).

Managing your clusters with Red Hat Hybrid Cloud Console services

[OpenShift Cluster Manager](#) is a managed service on the Red Hat Hybrid Cloud Console. You can use it to install, manage, and upgrade Red Hat OpenShift Container Platform clusters and create clusters on OpenShift Dedicated and OpenShift Service on AWS (ROSA). Using this also helps you:

- Centralize your OpenShift Container Platform and cloud services clusters to a single dashboard

- Get an overview of and create new clusters
- Monitor performance and troubleshoot stale clusters
- Get the support you need from Red Hat for managing your clusters

The Red Hat Hybrid Cloud Console also provides integrated services like the [Red Hat Insights Advisor](#). This helps you monitor the health and performance of your OpenShift Container Platform clusters and avoid downtime by highlighting whether your apps and services are available. The Insights Advisor can also help you identify security risks, affected clusters, and steps you can take.

And here's a final option for managing your OpenShift clusters and keeping them running smoothly: partnering with us at [TRIGO](#). We're a [custom software development service](#) with years of experience planning, implementing, and managing OpenShift technologies for our clients — when you need a team that brings both a professional and human perspective and is dedicated to your growth as a business.

How we're leveraging the power of OpenShift at TRIGO

At TRIGO, we've seen first-hand the incredible benefits a modernized, scalable OpenShift tech stack has brought to our major brand clients — like [Santander](#) and [3-S-IT](#). Because of this, we've doubled down on our services for addressing your needs when you're struggling to keep your OpenShift cluster running smoothly.

Partnering with us means our clients get the full benefits of our [OpenShift Consulting and Operations team](#). We provide you with the expertise you need to plan, implement, and manage your OpenShift clusters smoothly — with and without our support. Along with this, our team also supports our clients in the following:

- Maintaining your OpenShift clusters to keep them running smoothly
- Incident support to keep your systems and clusters running
- Accessing licensed OpenShift technology without having to go through Red Hat
- Gaining insights regarding what's working (and what isn't) for your tech
- Upgrading and/or migrating your tech across systems
- Getting the latest updates and insights for deploying your applications even faster

If you're interested in learning more about how our OpenShift consulting services work in practice, we'd love to hear from you. Feel free to book a first free

consultation with [Christian, our Head of Cloud Native Services](#), today — and let's get the conversation started.