

# Kubernetes log management

Whether you're a cloud native company that's used Kubernetes from the start, or you're adopting it for the first time, Kubernetes administration can be a monumental task. Here are our tips for adapting or extending your existing log management strategy to fit Kubernetes to gain critical visibility into your clusters.

## Issues with Kubernetes native logging

1

### Lack of centralization

Monitoring information is spread across multiple files on Kubernetes master and worker nodes.

2

### Log access

Kubernetes expects you to access application logs using the `kubectl` utility. Because of this, unlike on a conventional operating system, you can't use conventional text-manipulation tools (like `grep` and `awk`) to interact with Kubernetes log data unless you access the logs from outside the Kubernetes interface.

3

### Storage

By default, most Kubernetes distributions delete old data in application logs once the size of the log file exceeds 10 megabytes. If you haven't exported those logs before they're deleted, you will lose access to them.

4

### Log structure

Kubernetes doesn't automatically structure your log data, it just records whatever your containers dump to `stdout` or `stderr`. Whether that data is standardized and easy to work with depends on the way your containers are configured, not the way Kubernetes collects data from them.

+

## Tips for successful Kubernetes logging

### Don't use `kubectl` to manage logs

To analyze log data from all of the applications that you have running in Kubernetes, you'll need to connect log data to a third-party analytics and visualization tool.

### Don't settle for `stdout` and `stderr`

To ensure that you capture data from all of your applications, deploy an agent. The agent collects log data from the application in whichever form the application exposes it.

### Standardize logs

Kubernetes doesn't structure logs in a standard format, making them hard to query. Export your log data to a log manager that supports common logging formats, like Mezmo.

Using a log management solution like Mezmo, you can collect logs from all parts of your Kubernetes cluster, regardless of which Kubernetes distribution you use or whether your clusters run in the cloud, on-premises, or via a hybrid architecture. With Mezmo, three simple `kubectl` commands allow you to deploy a logging agent that will give you total visibility into your cluster. From there, view Kubernetes events in context with application logs, and easily set up alerts and views so teams understand the health of their applications.

To learn more about how Mezmo simplifies Kubernetes log management, [chat with a Mezmo solutions engineer](#) or start your fully-featured 14-day [free trial](#) today.