Serverless vs containers: a cost-and-control comparison

26 October 2026 · 5 min read · blogs.iictp.org/serverless-vs-containers

Serverless and containers are two ways to run code without managing physical servers. They trade control for convenience in different amounts.

TL;DR The short version

- Serverless and containers both avoid managing servers, but trade control for convenience in different amounts.
- Serverless suits short, event-driven, spiky work and scales to zero; containers suit steady, long-running or specialised work.
- It is not either/or: many systems use both.

Read online, listen to it, or change the text size: https://blogs.iictp.org/serverless-vs-containers

Serverless functions

- You supply code; the provider runs it on demand and scales it, including to zero.
- You pay for use, which suits spiky or occasional load.
- Limits on run time, memory and start-up latency apply.

Containers

- You package the application with its dependencies and run it on a platform you configure.
- More control over runtime, networking and long-running work.
- More to operate, and you pay for capacity while it runs.

How to choose

1. Short, event-driven, spiky work leans towards serverless.
2. Steady, long-running or specialised work leans towards containers.
3. Check limits and start-up behaviour against your latency needs.
4. Estimate cost at your real volume before deciding.

It is not either/or

Many systems use both. Pick per workload, and keep the boundaries clean so you can move one later.

Related credential: Cloud Architecture Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. IICTP. Cloud Architecture Practitioner credential page. iictp.org/certifications/cap.

Published by IICTP, an independent technical credential body, on 26 October 2026. The current version of this post is at https://blogs.iictp.org/serverless-vs-containers. More guides: blogs.iictp.org.