

Two pillars, one pane of glass

Grafana is doing double duty in this lab, and that's deliberate. It's the dashboard front-end for **metrics** (from InfluxDB) *and* — because Grafana can query multiple data sources — the front-end for **logs** (from Loki; see the next book).

```
Grafana (10.100.100.4)
|-- data source: InfluxDB -> metrics (how much, how fast, trending)
|-- data source: Loki     -> logs   (what happened, exact lines)
```

So when something looks off, you can pivot in one place: spot the CPU spike on a metric graph, then jump to the logs from that same window to see *why*. Metrics and logs are the first two "pillars of observability" (the third being traces), and having them behind one login is what makes investigating an incident feel like one workflow instead of three.

“ **Why we use this:** correlation is where observability pays off. A spike on a graph and a stack trace in the logs are far more useful together than apart. Pointing one Grafana at both your TSDB and your log store is a cheap way to get that correlation.

Revision #1

Created 2026-05-01 04:54:02 UTC by Eslam Shapsough

Updated 2026-05-01 05:03:02 UTC by Eslam Shapsough