

# Dashboards with Grafana

## DATA 503: Fundamentals of Data Engineering

Lucas P. Cordova, Ph.D.

2026-04-06

This lecture connects the presentation layer of a data pipeline to PostgreSQL on Railway. We review why dashboards matter, how Grafana fits the stack, how views and crosstabs prepare data for visualization, then follow a full path: deploy PostgreSQL 18, restore a class dump with `pg_restore`, add Grafana with a persistent config database, and build panels from curated views. A LaTeX lab handout supports in-class practice.

### Table of contents

<b>1</b>	<b>Learning Objectives</b>	<b>2</b>
<b>2</b>	<b>Part 1: Dashboards in the Pipeline</b>	<b>2</b>
<b>3</b>	<b>Part 2: Grafana Overview</b>	<b>3</b>
<b>4</b>	<b>Part 3: Views for Dashboards</b>	<b>5</b>
<b>5</b>	<b>Part 4: Dashboard-Ready Views (Review)</b>	<b>5</b>
<b>6</b>	<b>Part 5: Crosstabs and Pivot-Style Reporting</b>	<b>10</b>
<b>7</b>	<b>Part 6: Follow-Along — PostgreSQL 18 and Restore</b>	<b>12</b>
<b>8</b>	<b>Part 7: Follow-Along — Grafana on Railway</b>	<b>13</b>
<b>9</b>	<b>Part 8: Demo — Build Panels Step by Step</b>	<b>15</b>
<b>10</b>	<b>Part 9: Your Turn</b>	<b>16</b>
<b>11</b>	<b>Summary</b>	<b>16</b>

# 1 Learning Objectives

## 1.1 What You Will Be Able to Do

1. Explain where dashboards sit in a pipeline and why Grafana is a reasonable choice for SQL-backed analytics
2. Read and reason about dashboard-oriented views (aggregations, joins, classifications)
3. Interpret `crosstab()` pivots from `tablefunc` and know when a table panel in Grafana is the right visualization
4. Provision PostgreSQL 18 on Railway and restore a custom-format dump with `pg_restore` using the public database URL
5. Deploy Grafana on Railway with a separate `grafana` database for persistence and connect it to pipeline data with a read-only role
6. Build several panel types (bar, horizontal bar, grouped bar, table) from named views and pivot queries

## 1.2 Course Connection


This ties ingestion and modeling work to stakeholder-facing delivery: live, shareable visuals over the same database your pipeline maintains.

# 2 Part 1: Dashboards in the Pipeline

## 2.1 The Presentation Layer

Dashboards are the **presentation layer**. They turn vetted SQL into charts and tables other people can read without learning `JOIN` syntax.

## 2.2 Where Grafana Sits



...

Grafana does not replace your ETL jobs. It **reads** from PostgreSQL (ideally through views and restricted users) and renders results on a schedule or on demand.

## 2.3 What Makes a Good Pipeline Dashboard?

- **Live connection** to the database, not a one-off CSV export
- **Refresh** so panels stay current as data lands
- **Shareable** links for reviews and demos
- **Layered layout:** headline metrics on top, comparisons in the middle, detail tables below

## 3 Part 2: Grafana Overview

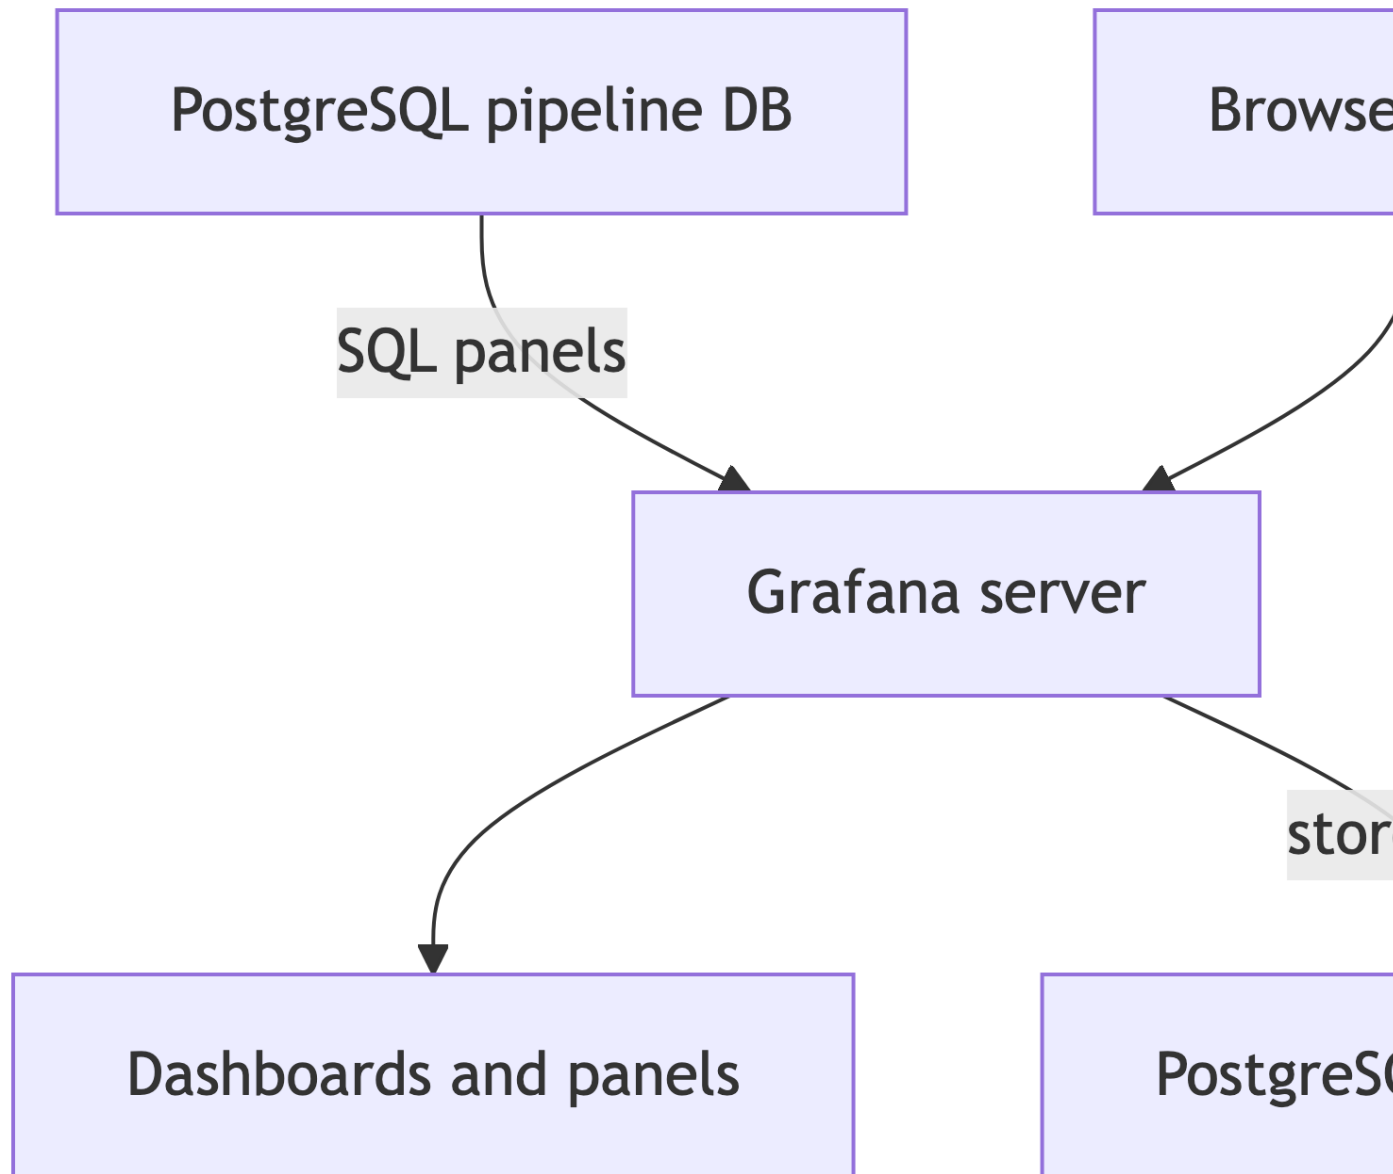
### 3.1 What Grafana Is

[Grafana](#) is an open-source observability and dashboarding platform. It started in metrics and logs, but its PostgreSQL data source is mature: you write SQL, Grafana maps columns to charts.

### 3.2 Why Grafana for This Course?

- **PostgreSQL-native** data source with a SQL editor
- **Many panel types:** time series, bar, table, stat, pie, heatmap, and more
- **Variables** for filters (dropdowns that rewrite queries)
- **Alerting** when a query crosses a threshold
- **Fits Railway:** can run next to your database; serverless mode reduces idle cost

### 3.3 Grafana Architecture (Two Databases)



1. **Analytics data** lives in your pipeline database (after restore, that is the class dataset).
2. **Grafana metadata** (dashboards, users, data source definitions) should live in a **separate** database so you can rebuild or clone pipeline data without losing UI work.

## 4 Part 3: Views for Dashboards

### 4.1 Why Views?

A **view** is a named `SELECT`. It behaves like a read-only table for clients.

- Hides long joins and aggregations behind one name
- Reuse the same logic from Grafana, APIs, and ad hoc queries
- Lets you grant `SELECT` on views to a read-only role without exposing base tables
- When logic changes, you update the view once, not every panel

...

Views do not magically cache results unless you use a **materialized** view. For class-scale data, ordinary views are fine.

### 4.2 Dataset After Restore

After you restore `railway_dump.dump`, you should see the census, business, temperature, and survey tables from the advanced-queries style curriculum, plus the views below (names may match exactly).

Table	Approximate role
<code>us_counties_pop_est_2019</code>	County population and components of change
<code>cbp_naics_72_establishments</code>	Hotels and accommodation establishments by area
<code>temperature_readings</code>	Daily readings for three stations
<code>ice_cream_survey</code>	Long-format survey responses

## 5 Part 4: Dashboard-Ready Views (Review)

### 5.1 View: `v_state_population_summary`

State-level totals and natural increase. Good for **sorted bar charts** or **tables**.

```
1 CREATE OR REPLACE VIEW v_state_population_summary AS
2 SELECT
3     state_name,
4     count(*) AS num_counties,
5     sum(pop_est_2019) AS total_pop,
6     round(avg(pop_est_2019), 0) AS avg_county_pop,
7     sum(births_2019) AS total_births,
```

```

8     sum(deaths_2019) AS total_deaths,
9     sum(births_2019) - sum(deaths_2019) AS natural_increase
10 FROM us_counties_pop_est_2019
11 GROUP BY state_name
12 ORDER BY total_pop DESC;

```

## 5.2 View: v\_state\_migration\_rates

Rates per thousand and a direction label. Good for **horizontal bar** charts and color by category.

```

1 CREATE OR REPLACE VIEW v_state_migration_rates AS
2 WITH state_migration AS (
3     SELECT
4         state_name,
5         sum(international_migr_2019 + domestic_migr_2019) AS net_migration,
6         sum(pop_est_2018) AS base_pop
7     FROM us_counties_pop_est_2019
8     GROUP BY state_name
9 )
10 SELECT
11     state_name,
12     base_pop,
13     net_migration,
14     round((net_migration / base_pop::numeric) * 1000, 2) AS migration_rate_per_thousand,
15     CASE
16         WHEN net_migration > 0 THEN 'Gaining'
17         WHEN net_migration < 0 THEN 'Losing'
18         ELSE 'Stable'
19     END AS migration_direction
20 FROM state_migration
21 ORDER BY migration_rate_per_thousand DESC;

```

## 5.3 View: v\_county\_growth\_summary

Aggregates counties by growth category. Natural **pie** or **stat** panels.

```

1 CREATE OR REPLACE VIEW v_county_growth_summary AS
2 WITH county_growth AS (
3     SELECT
4         county_name,
5         state_name,

```

```

6         pop_est_2019,
7         CASE
8             WHEN pop_est_2019 > pop_est_2018 THEN 'Growing'
9             WHEN pop_est_2019 < pop_est_2018 THEN 'Shrinking'
10            ELSE 'Stable'
11        END AS growth_category
12    FROM us_counties_pop_est_2019
13 )
14 SELECT
15     growth_category,
16     count(*) AS num_counties,
17     sum(pop_est_2019) AS total_pop,
18     round(avg(pop_est_2019), 0) AS avg_county_pop,
19     round(count(*)::numeric / (SELECT count(*) FROM us_counties_pop_est_2019) * 100, 1)
20     AS pct_of_counties
21 FROM county_growth
22 GROUP BY growth_category
23 ORDER BY num_counties DESC;

```

#### 5.4 View: v\_state\_estabs\_per\_capita

Establishments per thousand residents by state. Good **ranked bar** chart.

```

1 CREATE OR REPLACE VIEW v_state_estabs_per_capita AS
2 WITH estabs AS (
3     SELECT st, sum(establishments) AS establishment_count
4     FROM cbp_naics_72_establishments
5     GROUP BY st
6 ),
7 pop AS (
8     SELECT state_name, sum(pop_est_2018) AS total_pop
9     FROM us_counties_pop_est_2019
10    GROUP BY state_name
11 )
12 SELECT
13     pop.state_name,
14     pop.total_pop,
15     estabs.establishment_count,
16     round((estabs.establishment_count / pop.total_pop::numeric) * 1000, 1)
17     AS estabs_per_thousand
18 FROM estabs

```

```

19 JOIN pop ON estabs.st = pop.state_name
20 ORDER BY estabs_per_thousand DESC;

```

## 5.5 View: v\_seasonal\_temperatures

Season labels for each station. Ideal **grouped bar** chart.

```

1 CREATE OR REPLACE VIEW v_seasonal_temperatures AS
2 WITH labeled_readings AS (
3     SELECT
4         station_name,
5         CASE
6             WHEN date_part('month', observation_date) IN (12, 1, 2) THEN 'Winter'
7             WHEN date_part('month', observation_date) IN (3, 4, 5) THEN 'Spring'
8             WHEN date_part('month', observation_date) IN (6, 7, 8) THEN 'Summer'
9             WHEN date_part('month', observation_date) IN (9, 10, 11) THEN 'Fall'
10        END AS season,
11        max_temp,
12        min_temp
13    FROM temperature_readings
14 )
15 SELECT
16     station_name,
17     season,
18     round(avg(max_temp), 1) AS avg_max_temp,
19     round(avg(min_temp), 1) AS avg_min_temp,
20     count(*) AS num_readings
21 FROM labeled_readings
22 GROUP BY station_name, season
23 ORDER BY station_name,
24         CASE season
25             WHEN 'Winter' THEN 1
26             WHEN 'Spring' THEN 2
27             WHEN 'Summer' THEN 3
28             WHEN 'Fall' THEN 4
29        END;

```

## 5.6 View: v\_population\_tiers

County tiers by size. Strong **population share** story in one table.

```

1 CREATE OR REPLACE VIEW v_population_tiers AS
2 WITH tiers AS (
3     SELECT
4         county_name,
5         state_name,
6         pop_est_2019,
7         CASE
8             WHEN pop_est_2019 >= 500000 THEN 'Metro'
9             WHEN pop_est_2019 >= 100000 THEN 'Urban'
10            WHEN pop_est_2019 >= 50000 THEN 'Suburban'
11            ELSE 'Rural'
12        END AS tier
13    FROM us_counties_pop_est_2019
14 )
15 SELECT
16     tier,
17     count(*) AS num_counties,
18     sum(pop_est_2019) AS total_pop,
19     round(avg(pop_est_2019), 0) AS avg_county_pop,
20     round(sum(pop_est_2019)::numeric /
21         (SELECT sum(pop_est_2019) FROM us_counties_pop_est_2019) * 100, 1)
22     AS pct_of_total_pop
23 FROM tiers
24 GROUP BY tier
25 ORDER BY
26     CASE tier
27         WHEN 'Metro' THEN 1
28         WHEN 'Urban' THEN 2
29         WHEN 'Suburban' THEN 3
30         WHEN 'Rural' THEN 4
31     END;

```

## 5.7 Quick Check: Views vs. Inline SQL in Grafana

Why define views instead of pasting huge queries into every panel?

- A. Views always cache results for Grafana
- B. Views give a stable, reusable interface and cleaner permissions
- C. Grafana cannot run SQL
- D. Views refresh on a timer automatically

...

B. Non-materialized views still run fresh SQL. The win is naming, reuse, and security boundaries.

## 6 Part 5: Crosstabs and Pivot-Style Reporting

### 6.1 tablefunc and Long-to-Wide Data

Reporting often needs **wide** tables: one row per entity, many metric columns. Raw surveys and events are usually **long**: one row per observation.

...

PostgreSQL provides `crosstab()` through the extension:

```
1 CREATE EXTENSION IF NOT EXISTS tablefunc;
```

### 6.2 Ice Cream Survey Pivot

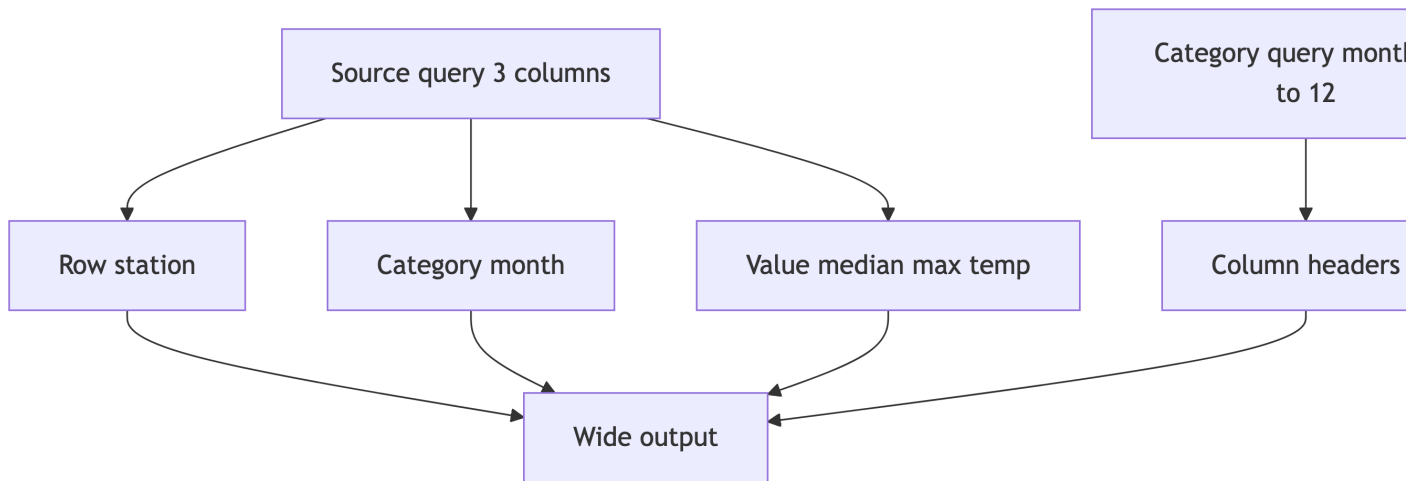
Long format: one row per person and flavor choice. Wide format: offices on rows, flavors as columns.

```
1 SELECT *
2 FROM crosstab(
3     'SELECT office, flavor, count(*)
4     FROM ice_cream_survey
5     GROUP BY office, flavor
6     ORDER BY office',
7     'SELECT flavor
8     FROM ice_cream_survey
9     GROUP BY flavor
10    ORDER BY flavor'
11 )
12 AS (office text,
13     chocolate bigint,
14     strawberry bigint,
15     vanilla bigint);
```

...

**Grafana note:** This returns a **table** shape. Use a **Table** panel, not a bar chart, unless you unpivot again in SQL.

### 6.3 Temperature Median Pivot by Month



```
1 SELECT *
2 FROM crosstab(
3     'SELECT station_name,
4         date_part('month', observation_date),
5         percentile_cont(.5)
6         WITHIN GROUP (ORDER BY max_temp)
7     FROM temperature_readings
8     GROUP BY station_name,
9         date_part('month', observation_date)
10    ORDER BY station_name',
11    'SELECT month FROM generate_series(1,12) month'
12 )
13 AS (station text,
14     jan numeric(3,0), feb numeric(3,0), mar numeric(3,0),
15     apr numeric(3,0), may numeric(3,0), jun numeric(3,0),
16     jul numeric(3,0), aug numeric(3,0), sep numeric(3,0),
17     oct numeric(3,0), nov numeric(3,0), dec numeric(3,0));
```

### 6.4 Quick Check: crosstab() Shape

The **source query** passed to `crosstab()` must return how many columns?

- A. 2
- B. 3
- C. 4

- **D.** It depends on the number of categories

...

**B.** Row id, category, value. The second query lists category labels for column order.

## 7 Part 6: Follow-Along — PostgreSQL 18 and Restore

### 7.1 Step 1: New Railway Project and Postgres 18

1. In [Railway](#), create a **new project** (empty is fine).
2. **Add PostgreSQL.** In the database service **Settings**, choose **PostgreSQL 18** if the platform exposes a version selector or image tag; otherwise use the image your instructor pins for this term.
3. Wait until the service is healthy. Open **Variables** (or **Connect**) and copy the **public** database URL Railway exposes for external access (often named `DATABASE_PUBLIC_URL`, or shown under **Public network / TCP Proxy** on the Connect tab). You will use this same URL style for Beekeeper, pgAdmin, `psql`, and `pg_restore` on your laptop.

...

**Important:** Anything you run **on your own computer** must use that **public** URL. Private hostnames such as `*.railway.internal` only resolve inside Railway's network and will fail from your laptop.

### 7.2 Step 2: Download the Dump

From Canvas, download `railway_dump.dump`. This file is a **custom-format** archive intended for PostgreSQL 18 and `pg_restore`.

...

If you cannot install client tools, use the optional plain `railway_dump.sql` from Canvas instead (run it in Beekeeper or pgAdmin, or with `psql -f`).

### 7.3 Step 3: Restore with `pg_restore`

In a terminal on the machine that has the file:

```
1 export DATABASE_PUBLIC_URL="postgresql://USER:PASSWORD@PUBLIC_HOST:PORT/railway"
2 pg_restore --verbose --no-owner --no-acl -d "$DATABASE_PUBLIC_URL" railway_dump.dump
```

...

- `--no-owner --no-acl` avoid errors about roles and permissions that do not exist on Railway.
- The database name in the URL is often `railway`; use whatever Railway shows.
- Client version: prefer `pg_restore 18` to match the server. Older clients sometimes fail against newer servers.

## 7.4 Step 4: Verify in a SQL Client

Connect with **Beekeeper Studio** or **pgAdmin** using the **public** URL (paste the full URL or fill host, port, user, password, and database from it). Confirm:

- Tables such as `us_counties_pop_est_2019`, `ice_cream_survey`, `temperature_readings`
- Views such as `v_state_population_summary`
- Extension `tablefunc` if the dump included it (needed for `crosstab()` examples)

## 7.5 Step 5: Plain SQL Fallback

If `pg_restore` is not available:

1. Download `railway_dump.sql` from Canvas.
2. Execute the entire script while connected to your Railway database (Beekeeper **Run**, pgAdmin **Query Tool**, or `psql "$DATABASE_PUBLIC_URL" -f railway_dump.sql`).

# 8 Part 7: Follow-Along — Grafana on Railway

## 8.1 Create the grafana Database

In Beekeeper or pgAdmin, connected to the **same** Postgres server:

```
1 CREATE DATABASE grafana;
```

...

Grafana will store dashboards here. Skipping this step usually means **lost dashboards** after redeploys.

## 8.2 Add Grafana (Template)

1. In the same Railway project, **Create** then **Template** and search **Grafana**.
2. Deploy. Note the **public URL** Grafana receives; set `GF_SERVER_ROOT_URL` to that full `https://...` value after the first deploy if the template asks for it.

## 8.3 Environment Variables (Checklist)

### Admin login

Key	Value
GF_SECURITY_ADMIN_USER	admin
GF_SECURITY_ADMIN_PASSWORD	strong password you keep

### Grafana config database

Key	Value
GF_DATABASE_TYPE	postgres
GF_DATABASE_HOST	<b>host:port</b> from the same <b>public</b> Postgres URL you use in Beekeeper (not a private internal hostname)
GF_DATABASE_NAME	grafana
GF_DATABASE_USER	postgres (unless Railway uses another superuser name)
GF_DATABASE_PASSWORD	from Postgres service variables
GF_DATABASE_SSL_MODE	usually <b>require</b> when using the public endpoint (match Railway's docs if the test fails)

...

Enable **Serverless** on Grafana if you want the service to sleep when idle.

## 8.4 Read-Only Role for Panel Queries

Still connected to the **pipeline** database (not **grafana**):

```
1 CREATE USER grafanareader WITH PASSWORD 'use_a_strong_password';
2
3 GRANT USAGE ON SCHEMA public TO grafanareader;
4 GRANT SELECT ON ALL TABLES IN SCHEMA public TO grafanareader;
5 ALTER DEFAULT PRIVILEGES IN SCHEMA public
6     GRANT SELECT TO grafanareader;
```

...

In Grafana, add the **PostgreSQL data source** for the pipeline database using the **public** host and port (same as Beekeeper), database name (often **railway**), user **grafanareader**, and **TLS/SSL mode require** unless Railway's Connect instructions say otherwise.

## 9 Part 8: Demo — Build Panels Step by Step

Assume you are logged into Grafana as admin and the PostgreSQL data source **Save & test** succeeded against the restored database.

### 9.1 Demo A: Top States by Population (Bar Chart)

1. **Dashboards** then **New** then **New dashboard**.
2. **Add visualization**. Pick your **PostgreSQL** data source.
3. Set query editor to **Code (SQL)**. Paste:

```
1 SELECT state_name, total_pop
2 FROM v_state_population_summary
3 ORDER BY total_pop DESC
4 LIMIT 15;
```

4. Visualization: **Bar chart**.
5. In **Panel options**, set category field to `state_name` and value field to `total_pop`.
6. Title the panel clearly, for example **Top 15 states by estimated 2019 population**.

### 9.2 Demo B: Migration Rates (Horizontal Bar)

1. **Add visualization** on the same dashboard.
2. SQL:

```
1 SELECT state_name, migration_rate_per_thousand, migration_direction
2 FROM v_state_migration_rates
3 ORDER BY migration_rate_per_thousand DESC
4 LIMIT 20;
```

3. Visualization: **Bar chart**, orientation **horizontal**.
4. Optional: under **Overrides**, color by field `migration_direction` (map Gaining, Losing, Stable to different colors).
5. Title: **Net migration rate per 1,000 residents (selected states)**.

### 9.3 Demo C: Seasonal Temperatures (Grouped Bar)

1. **Add visualization**.
2. SQL:

```
1 SELECT station_name, season, avg_max_temp
2 FROM v_seasonal_temperatures;
```

3. Visualization: **Bar chart**, mode **grouped** (exact control names vary by Grafana version; look for grouping by `station_name` with `season` as series or x-axis categories depending on the panel).
4. Title: **Average max temperature by station and season.**

## 9.4 Demo D (Optional Table): Ice Cream Crosstab

1. **Add visualization** then choose **Table**.
2. Paste the full `crosstab` query from Part 5 (ice cream example).
3. Run query. Confirm columns are office plus flavor counts.
4. Title: **Ice cream preferences by office (pivot).**

...

**Talking point:** Views feed most charts; pivots often stay as **tables** unless you reshape data further for chart types.

## 10 Part 9: Your Turn

### 10.1 During Class

Replicate the follow-along: **Postgres 18**, **restore**, **Grafana** with **grafana database**, **read-only user**, then build a dashboard with **at least four** panels. Mix views and at least one **table** (a view or a crosstab query).

### 10.2 Take-Home Reference

Download the lab steps as [grafana-lab-handout.pdf](#) or the LaTeX source [grafana-lab-handout.tex](#) if you want to edit or recompile locally. Your instructor may mirror the PDF on Canvas.

## 11 Summary

Topic	Takeaway
Pipeline role	Dashboards visualize vetted data; they are not the transform layer
Views	Named, reusable, permission-friendly interfaces for panels

---

Topic	Takeaway
Crosstabs	<code>crosstab()</code> builds wide tables; great for <b>Table</b> panels
Restore	<code>pg_restore --no-owner --no-acl -d "\$DATABASE_PUBLIC_URL" file.dump</code>
Grafana persistence	Separate <code>grafana</code> database for config
Security	Read-only DB user for Grafana queries

---

## 11.1 References

1. [Grafana documentation](#)
2. [Grafana PostgreSQL data source](#)
3. [Railway documentation](#)
4. [PostgreSQL CREATE VIEW](#)
5. [PostgreSQL tablefunc \(`crosstab`\)](#)
6. [PostgreSQL `pg\_restore`](#)