




Data as Code: Exploring Data and Database Portability

New York Oracle User Group
April 27, 2023



Sean Scott

Oracle ACE Director
Managing Principal Consultant

-  @oraclesean
-  [linkedin.com/in/soscott/](https://www.linkedin.com/in/soscott/)
-  sean.scott@viscosityna.com



Sean Scott

Oracle ACE Director
Managing Principal Consultant

-  @oraclesean
-  [linkedin.com/in/soscott/](https://www.linkedin.com/in/soscott/)
-  sean.scott@viscosityna.com

www.viscosityna.com

DATABASE RELIABILITY ENGINEERING :: DEVOPS & AUTOMATION
HIGH AVAILABILITY :: BUSINESS CONTINUITY :: DISASTER RECOVERY

MODERNIZATION :: OBSERVABILITY :: ENGINEERED SYSTEMS

AutoUpgrade :: Zero Downtime Migrations :: Patching

Real Application Clusters :: Data Guard :: Sharding

Docker/Containers :: Terraform :: Ansible

Exadata :: Oracle Database Appliance

AHF :: TFA :: GIMR :: CHA

 @ViscosityNA

Oracle on Docker

Running Oracle Databases in Linux Containers

Oracle on Docker


Running Oracle Databases in
Linux Containers

—
Sean Scott


Apress®

Download a free sample chapter:
<https://oraclesean.com>

Databases are Special™



Phil Connors:
main




Feature Request

Customers are complaining that the Phil microservice is narcissistic and misogynistic.

Redesign Phil to be a decent human being.

Phil Connors:

- main
 - └ day-2
 - └ bank-robber
 - └ piano-student
 - └ redemption



Phil Connors:


main

└ day-2

└ bank-robber

└ piano-student

└ redemption



6:00 AM

Phil wakes up

7:00 AM

Encounters Ned

7:20 AM

Tapes segment

9:00 AM

Blizzard!

Phil Connors:


main

└ day-2

└ bank-robber

└ piano-student

└ redemption



6:00 AM

Phil wakes up

7:00 AM

Encounters Ned

7:20 AM

Tapes segment

9:00 AM

Blizzard!

Phil Connors:


main

└ day-2

└ bank-robber

└ piano-student

└ redemption



Alarm doesn't go off!

Ned is POTUS!

Multiple groundhogs?

Miami Beach?

6:00 AM

7:00 AM

7:20 AM

9:00 AM

Problems

- Every test further pollutes data.
- Restoring databases is operationally difficult, requiring:
 - Downtime
 - Specialized skills/DBA intervention
 - Every database is a production database
- Making databases idempotent is technically difficult

Common Solutions

- Build exceptions/elasticity into CI/CD pipelines
 - Technically difficult
 - Requires upkeep
 - Not absolute
- Assume certain errors can be ignored
- Technology solution (\$\$\$\$\$)


But... Data is Code

Data *is* Code

Every record in every database is created in one of three ways:

- **INSERT**
- **UPDATE**
- **DELETE**


Heterogenous replication = synchronous SQL



Optional
Transformation


Databases are Special™

"[F]or a long time, DBAs were in the business of crafting silos and snowflakes. Their tools were different, their hardware was different, and their languages were different."



Laine Campbell, Charity Majors
Database Reliability Engineering
O'Reilly


"The days in which this model can prove itself to be effective and sustainable are numbered."



Laine Campbell, Charity Majors
Database Reliability Engineering
O'Reilly


"Our systems are no more or less important than any other components serving the needs of the business.

[T]he components of database clusters are not sacred."

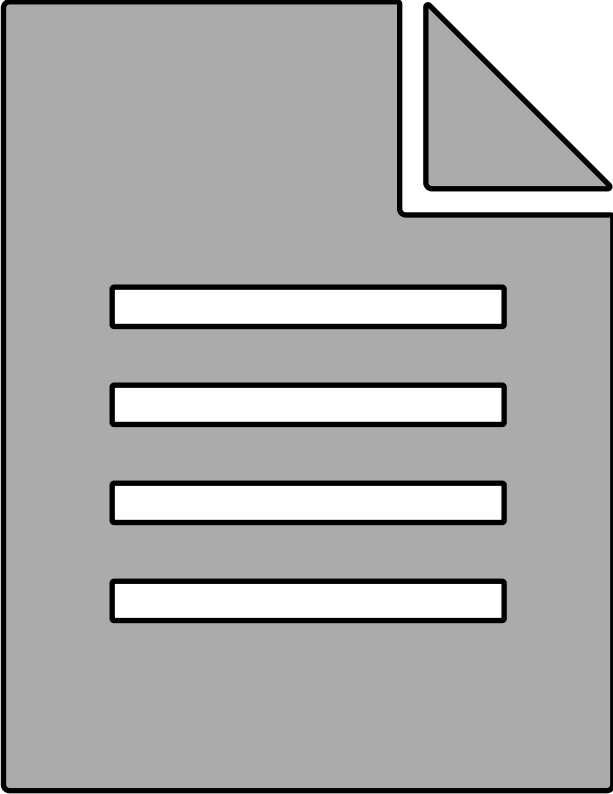


Laine Campbell, Charity Majors
Database Reliability Engineering
O'Reilly

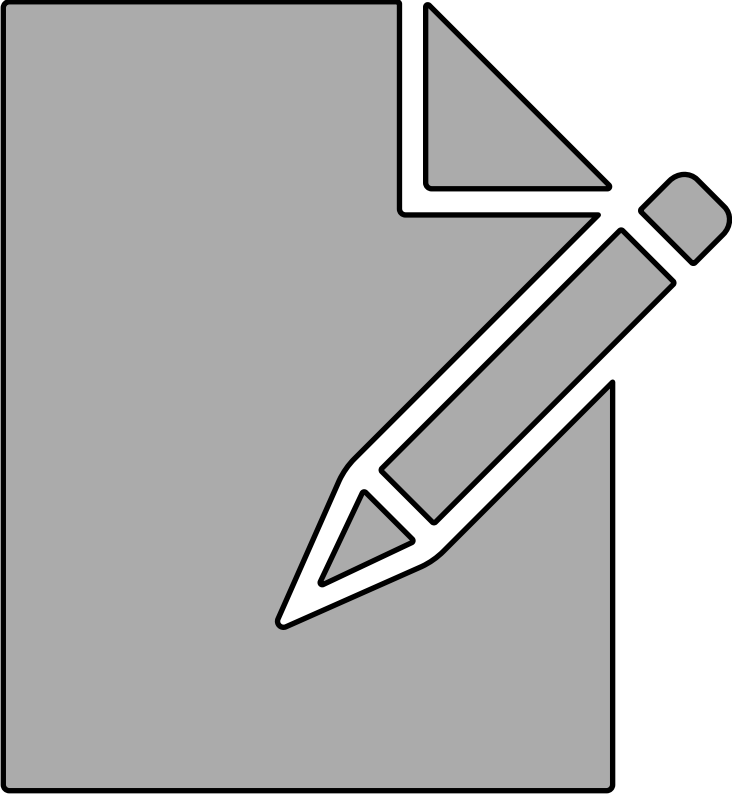
Databases are Special™



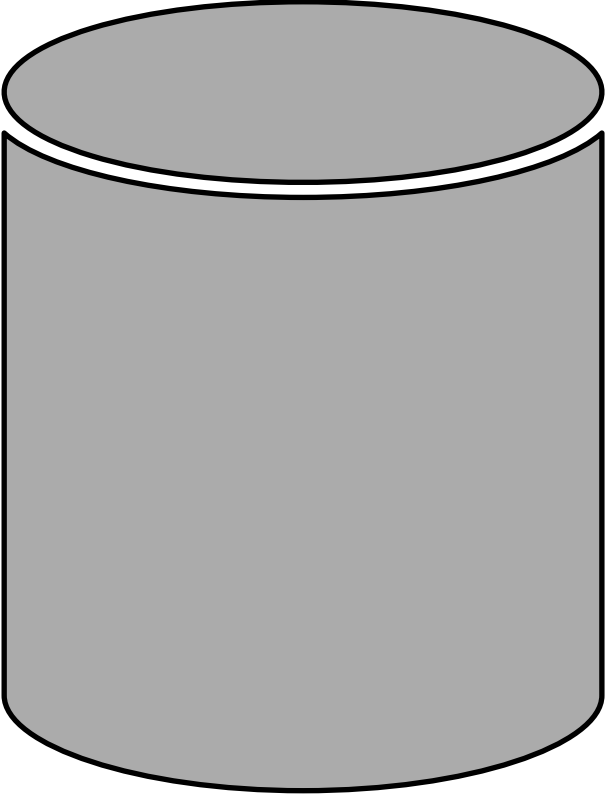
Databases are Special™



Software



Configuration



Data

What is Cloud Native?

Electricity is like a water hose


Voltage

Volts (V)



Current

Amps (A or I)





Resistance

Ohms (R or Ω)



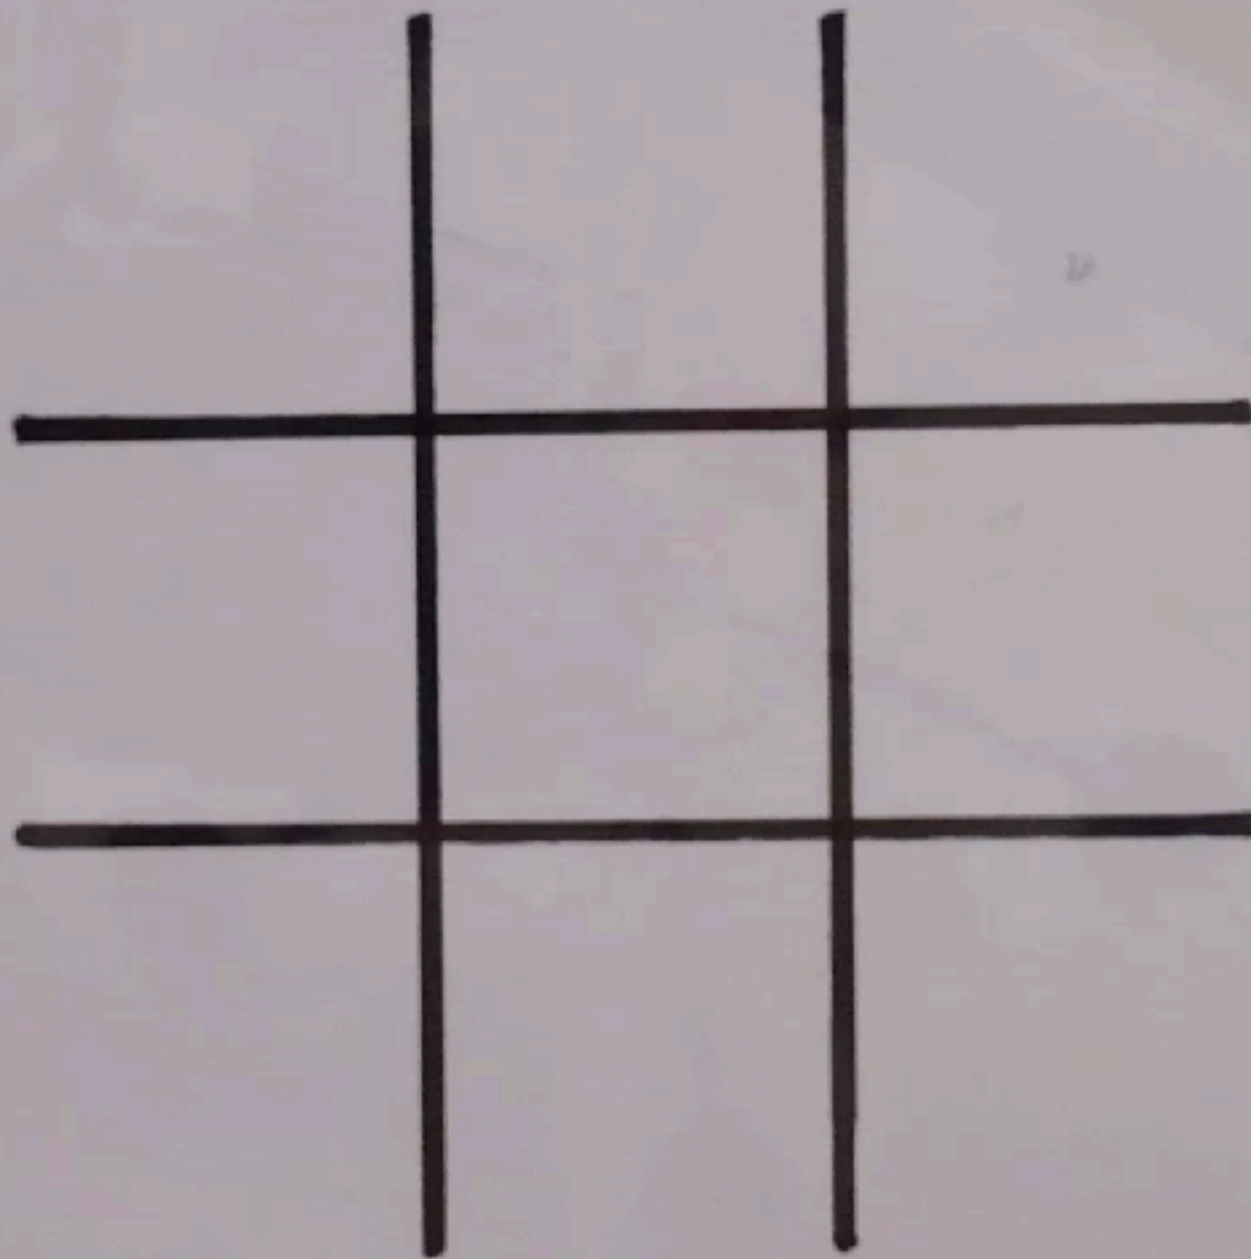
FREEING
ENERGY






An **Image** is the game board and rules;
Containers host and persist games;
The superimposed **Layer** saves moves;
Deleting a **Container** removes its **Layer**.

Stalemate!





X

O

X

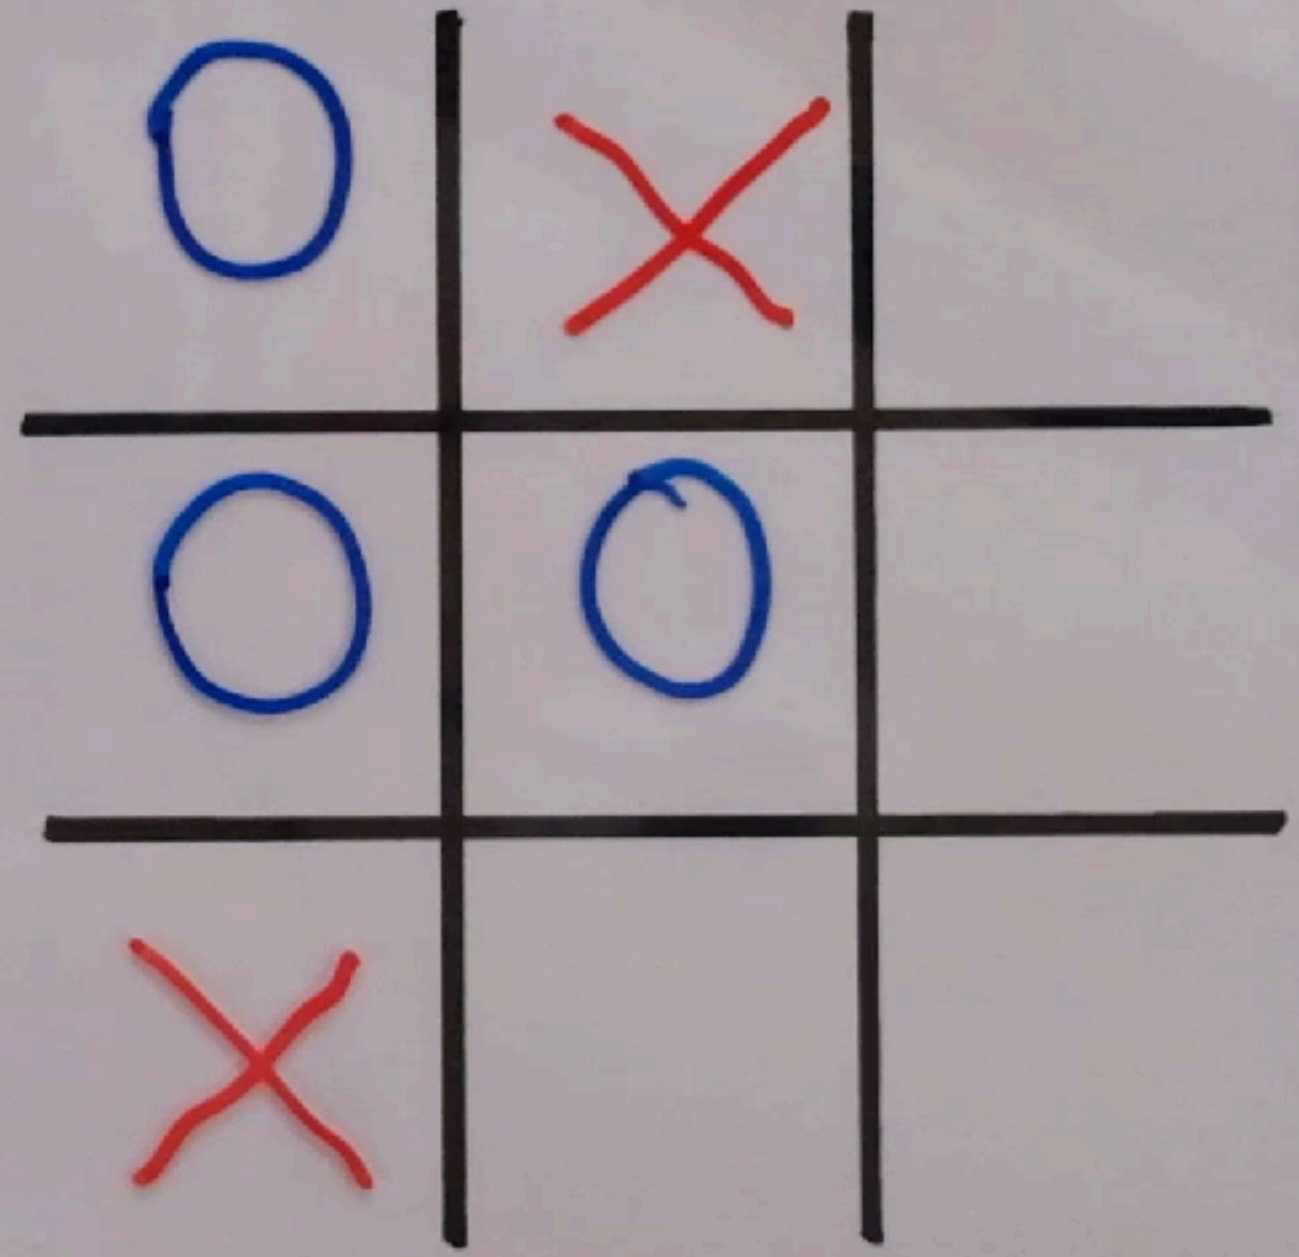
X

O



O

O

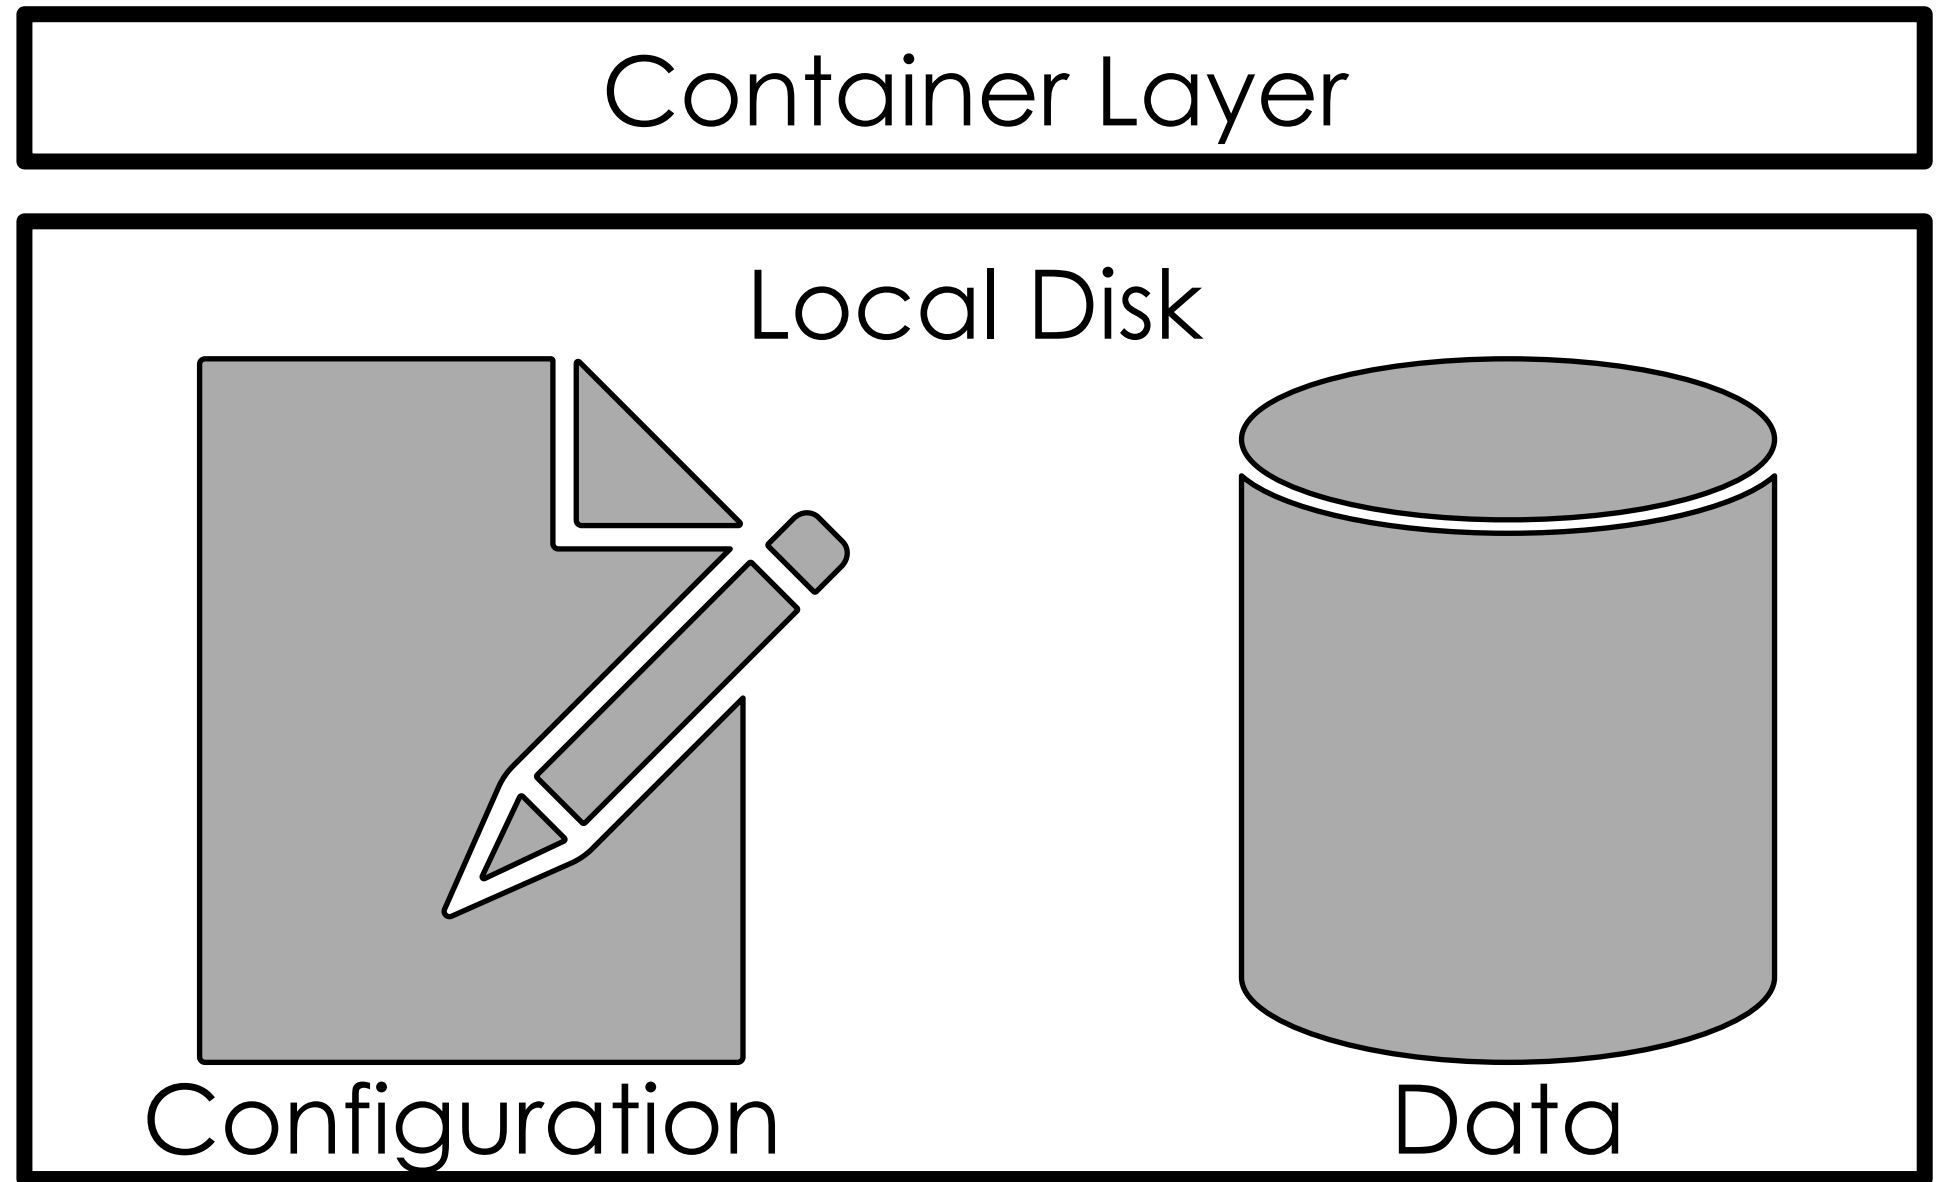

X



Database Containers



Database Containers



```
CONTAINER_NAME=punxsutawney_test
```

```
mkdir -p /data/${CONTAINER_NAME}/data
```

```
docker volume create \  
  --opt type=none \  
  --opt o=bind \  
  --opt device=/data/${CONTAINER_NAME}/data \  
  ${CONTAINER_NAME}_data
```

An **Image** is the game board and rules.

Containers host and persist games.

Database Image Rule #1:

If there is a database, start it.

If not, create a new database.

Database Container Startup

```
docker run -d \  
  --name ${CONTAINER_NAME} \  
  -e ORACLE_SID=ORCL \  
  -v /data/${CONTAINER_NAME}:/opt/oracle/oradata \  
  oracle/database:19.3.0-ee
```

Clone the Data Directory

```
cp -rp /data/punxsutawney/ \
      /data/${CONTAINER_NAME}
```

Clone the Data Directory

```
cp -rp /data/punxsutawney/ \
      /data/${CONTAINER_NAME}
```

```
cp -rp /NFS/gold_data/punxsutawney/ \
      /data/${CONTAINER_NAME}
```

Clone the Data Directory


```
cp -rp /data/punxsutawney/ \
    /data/${CONTAINER_NAME}
```

```
cp -rp /NFS/gold_data/punxsutawney/ \
    /data/${CONTAINER_NAME}
```

```
cp -rp /NFS/gold_data/punxsutawney/2023-02-02.0600/ \
    /data/${CONTAINER_NAME}
```

If data can be codified, **data is versionable.**


Versioned Data



2023-02-02

Wake up campers!
Phil? Phil Connors?
Six more weeks of winter?
Blizzard is coming!

Versioned Data






2023-02-02

```
select *  
  from knowledge  
 where fname = 'Phil'  
    and lname = 'Connors';
```

ARMORED_TRUCKS:	NULL
LAKES_AND_RIVERS:	NULL
IMPRESSIVE_PIANO:	NULL
ICE_SCULPTING:	NULL

Versioned Data





Punxsutawney

Town in Pennsylvania

Punxsutawney is a borough in southern Jefferson County, Pennsylvania, United States. Punxsutawney is known globally for its annual Groundhog Day celebration held each February 2, during which thousands ... [Wikipedia](#)

Area: 3.429 mi²

Elevation: 1,230'


Weather: 33°F (1°C), Wind W at 6 mph (10 km/h), 83% Humidity

[More on weather.com](#)


Population: 5,725 (2021)

Apply **lean principles** to versioned data.

Versioned Data



+



+



Entrypoints

- Generally adapted but *not* an OCI standard
- Implementation varies by image, vendor/author
 - Setup only
 - Setup & startup


Entrypoints

```
CONTAINER_NAME=BANK-ROBBER
```

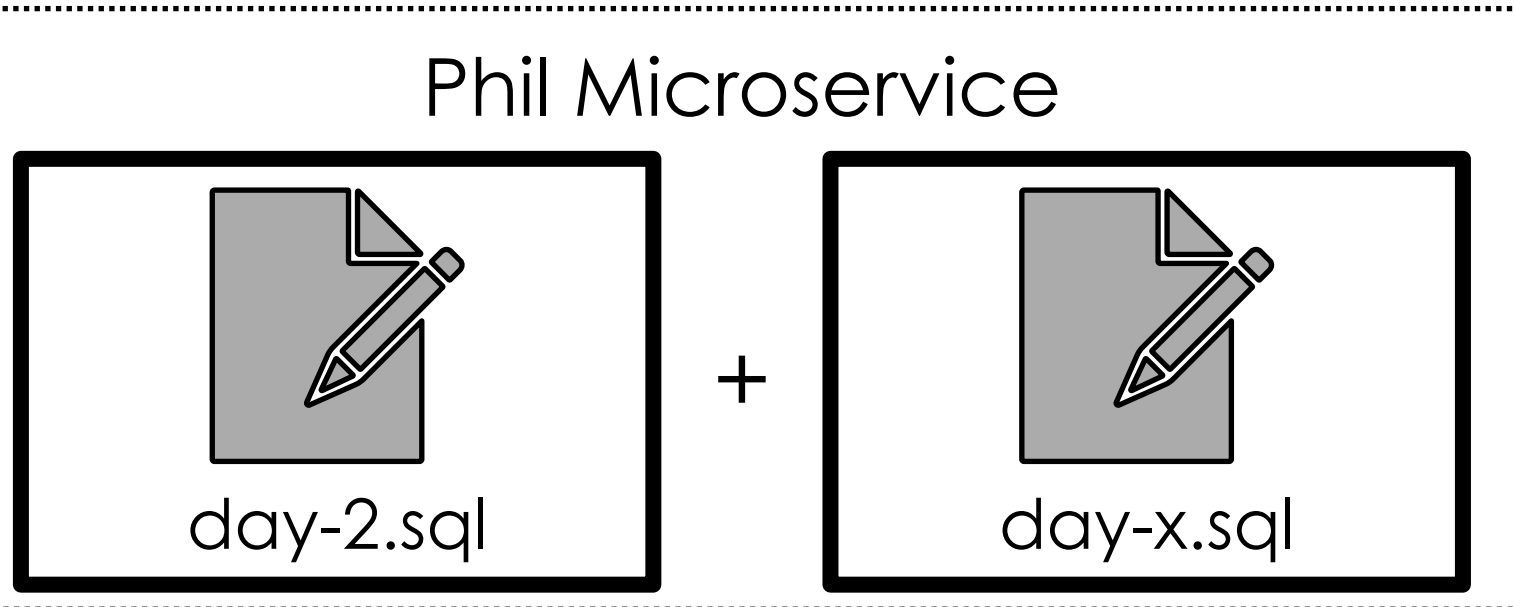
```
cp -r /data/punxsutawney_orig/ /data/${CONTAINER_NAME}
```

```
docker run -d \  
  --name ${CONTAINER_NAME} \  
  -e ORACLE_SID=ORCL \  
  -v /data/${CONTAINER_NAME}:/opt/oracle/oradata \  
  -v /phil/${CONTAINER_NAME}:/docker-entrypoint-initdb.d \  
  oracle/database:19.3.0-ee
```

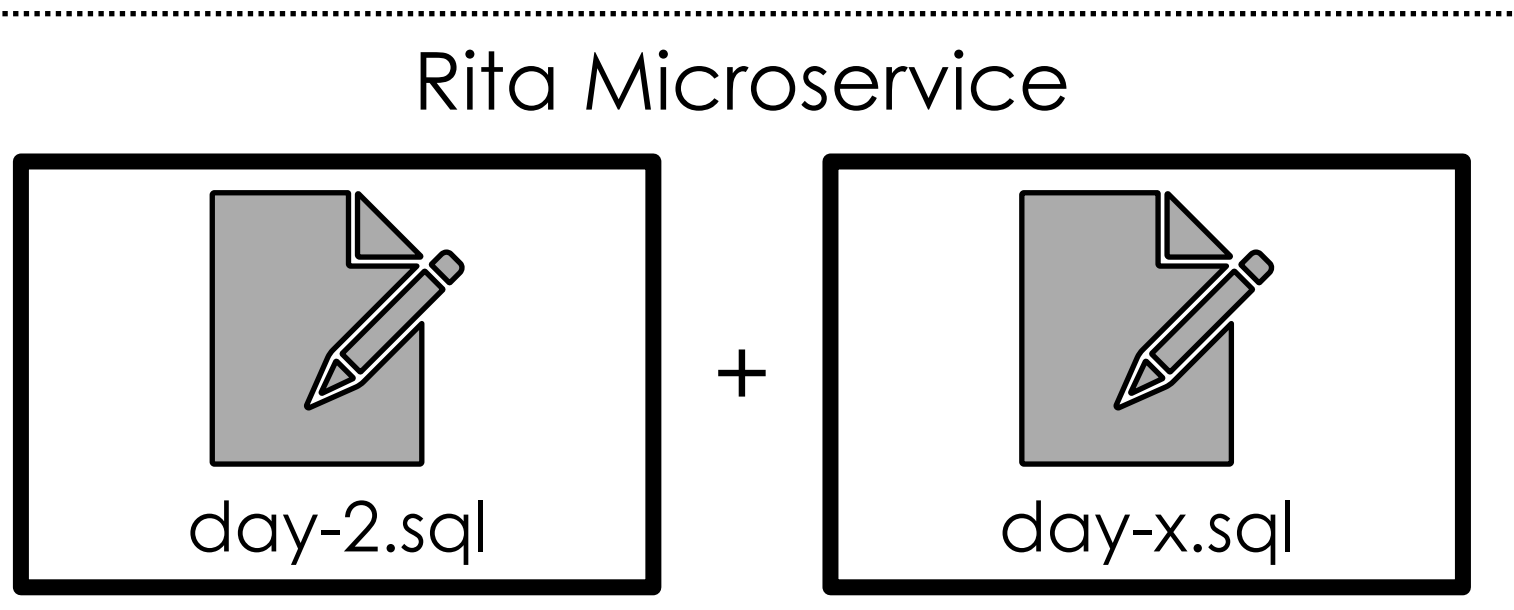

Versioned Data



+



+




Save **ancestor data** as a **dataset**.

Apply **versioned** SQL to create **descendants**.

Phil Connors:

- main
 - └ day-2
 - └ bank-robber
 - └ piano-student
 - └ redemption



Phil wakes up

Encounters Ned

Tape segment

Blizzard!


6:00 AM

7:00 AM

7:20 AM

9:00 AM

Phil Connors:
main



Data as Code:

Standardizes data

Decouples DBAs from the development lifecycle

Accelerates & streamlines development

Assures test reliability and quality



oraclesean.com



<https://www.linkedin.com/in/soscott/>



[@oraclesean](https://twitter.com/oraclesean)



<https://github.com/oraclesean>



sean.scott@viscosityna.com



https://youtube.com/@oracle_sean