Databases, Containers, and the Cloud

BRUCE MOMJIAN

This presentation explains the new options of container and cloud deployments.

https://momjian.us/presentations

Last updated: April 2023

Outline

1. Traditional Data Centers

Punch cards

https://www.flickr.com/photos/mwichary/

Electronic & Manual Storage

IBM System/360

https://www.flickr.com/photos/simplyphotography/

All Electronic

HP 9000 N-Class server

https://www.flickr.com/photos/helixblue/

Modern

Google data center

https://www.flickr.com/photos/shiver_silver/

2. Container Features

https://www.flickr.com/photos/jaxport/

What Are Containers

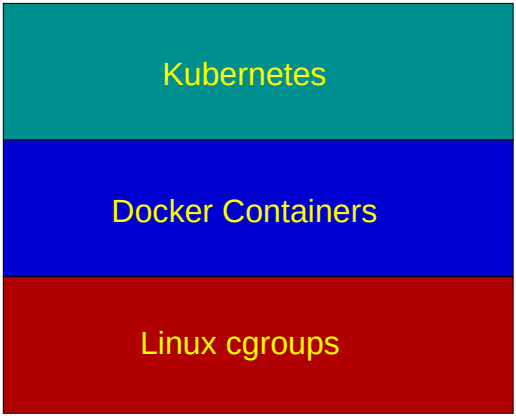

Container deployment

Application container

Process isolation

cgroups

- Process isolation
- Resources control
- CPU prioritization
- Accounting
- Freezing, checkpointing, restarting

Docker

- Executables
- Libraries overlayed using a union file system
- Specification file
- Uses cgroups
- Uses namespace/network/user isolation

Kubernetes

- Container deployment
- Scaling
- Monitoring
- Load balancing
- Stateful sets (durable storage)

Containers Using a Single Kernel

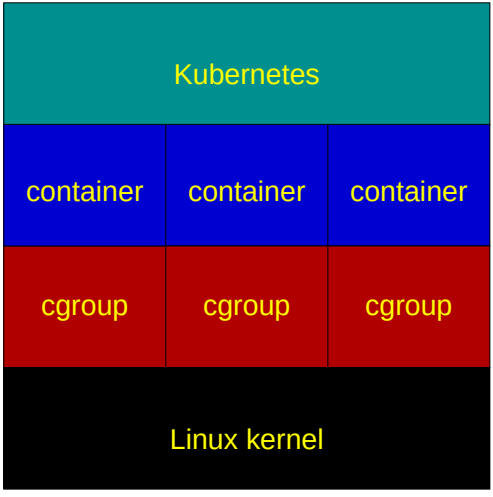

Containers Using Multiple Kernels

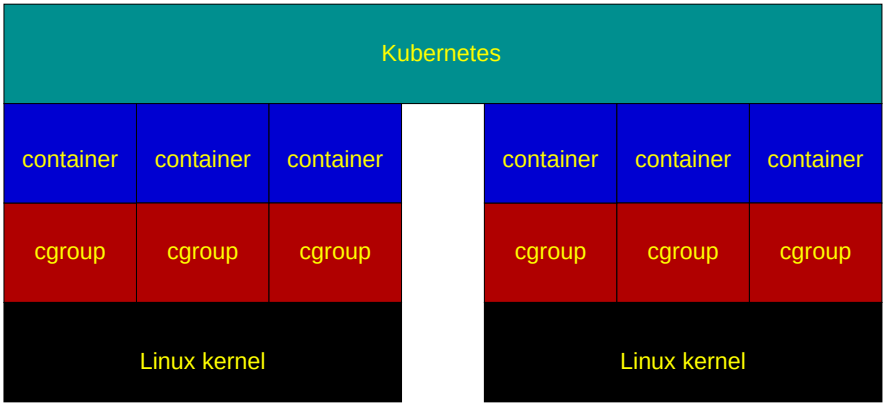

3. Database Containers

Container Capability	Benefit for Databases
rapid creation/destruction	no
less overhead than VM	no
scaling	limited
migration	limited
automated deployment	yes

Containers for Database Tooling

- Backup
- Monitoring
- Failover
- Connection pooling
- Scaling

4. Cloud Features

- No physical hardware/infrastructure to maintain
- Hardware, power, and network failures handled
- Storage recovery handled
- Increase/decrease usage easily
- Less staff time

Deployment Option	Benefit for Databases
Private servers with containers	easy deployment
Private cloud (virtual machines) with containers	above, plus different operating systems
Public cloud with self-installed software	public cloud benefits (previous slide)
Public cloud with cloud-specific software	above, plus optimizations

Cloud-Specific Software

Most database software is written for generic hardware and infrastructure.
Cloud-specific software can be optimized for:

- Storage characteristics
- High availability/fail-over
- Backup/restore
- Monitoring
- Scaling
- Persistent memory
- GPUs and FPGAs
- Single vendor to blame

6. Conclusion

- Containers ease database deployment
- Public cloud reduces the complexity of managing hardware
- Cloud-specific software leverages cloud infrastructure

https://momjian.us/presentations

https://www.flickr.com/photos/mradambrown/