

Dynamic Infrastructure Orchestration

Jochen Lillich



Jochen Lillich

@geewiz

Sysadmin since 1993

CEO of freistil IT

freistilbox: PaaS for Drupal & WordPress

Config Management

Elimination of manual changes

Consistent configuration

Single source of truth

DRY

Weaknesses

Periodic convergence runs

Central database

Orchestration

Perform changes quickly

Handle failure

Tolerate network partitions

serf

Basics

Cluster communication tool

Developed by Hashicorp

Simple binary, written in Go

Gossip protocol

serf node 1

```
$ serf agent -node=node1 -bind=svr1
```

```
==> Starting Serf agent...
```

```
==> Starting Serf agent RPC...
```

```
==> Serf agent running!
```

```
Node name: 'svr1'
```

```
Bind addr: '192.168.0.11:7946'
```

```
...
```

```
==> Log data will now stream in as it occurs:
```

```
2014/04/09 16:05:51 [INFO] agent: Serf agent starting
```

```
2014/04/09 16:05:51 [INFO] serf: EventMemberJoin: node1  
192.168.0.11
```

```
2014/04/09 16:05:52 [INFO] agent: Received event: member-  
join
```


serf node 2

```
$ serf agent -node=node2 -bind=svr2
```

```
...
```

```
$ serf join svr1
```

```
Successfully joined cluster by contacting 1  
nodes.
```

```
$ serf members
```

```
node2 192.168.0.12:7946 alive
```

```
node1 192.168.0.11:7946 alive
```

Membership

Cluster management

Service configuration

Failure detection

Detects failure quickly

Notifies other nodes

Executes handler scripts

Events

member-join

member-reap

member-leave

user

member-failed

query

member-update

User events

Deploy application

Perform convergence run

Query information from the cluster

Example handler

```
#!/bin/bash
```

```
echo
```

```
echo "New event: ${SERF_EVENT}. Data  
follows..."
```

```
while read line; do
```

```
    printf "${line}\n"
```

```
done
```

User event

```
$ serf event deploy-app
```

```
2014/04/09 17:06:32 [INFO] agent: Received  
event: user-event: deploy-app
```

etcd

Basics

HA key-value store

Developed by CoreOS

Written in Go

Raft consensus algorithm

Pure HTTP API

Key-value store


```
/release = 9607bcfeb48905d26db8f
```

```
/cluster
```

```
  /cluster/node1 = node1.example.com
```

```
  /cluster/node2 = node2.example.com
```

Quorum



Available

Quorum




Available

Quorum



Available

Quorum



Unavailable

Writing

```
$ etcdctl set /message Hello
```

```
Hello
```

```
$ curl -L -X PUT \  
  http://127.0.0.1:4001/v2/keys/message \  
  -d value="Test"
```

```
{"action": "set", "node": {"key": "/  
message", "value": "Test", "modifiedIndex":  
15, "createdIndex": 15}, "prevNode": {"key": "/  
message", "value": "Hello", "modifiedIndex":  
2, "createdIndex": 2}}
```

Reading

```
$ etcdctl get /message
```

```
Hello
```

```
$ curl -L \
```

```
  http://127.0.0.1:4001/v2/keys/message
```

```
{"action": "get", "node": {"key": "/  
message", "value": "Test", "modifiedIndex":  
15, "createdIndex": 15}}
```


Deleting

```
$ etcdctl rm /message
```

```
$ curl -L -X DELETE \  
http://127.0.0.1:4001/v2/keys/message
```

Waiting

```
$ curl -L http://127.0.0.1:4001/v2/keys/  
release\?wait\=true
```

Waiting

```
$ curl -L http://127.0.0.1:4001/v2/keys/  
release\?wait\=true
```

```
{"action": "set", "node": {"key": "/  
release", "value": "9607bcfe", "modifiedIndex":  
16, "createdIndex": 16}, "prevNode": {"key": "/  
release", "value": "18512199", "modifiedIndex":  
14, "createdIndex": 14}}
```

Applications

Locking

Leader election

Database master

Elastic IP

Leader election




App

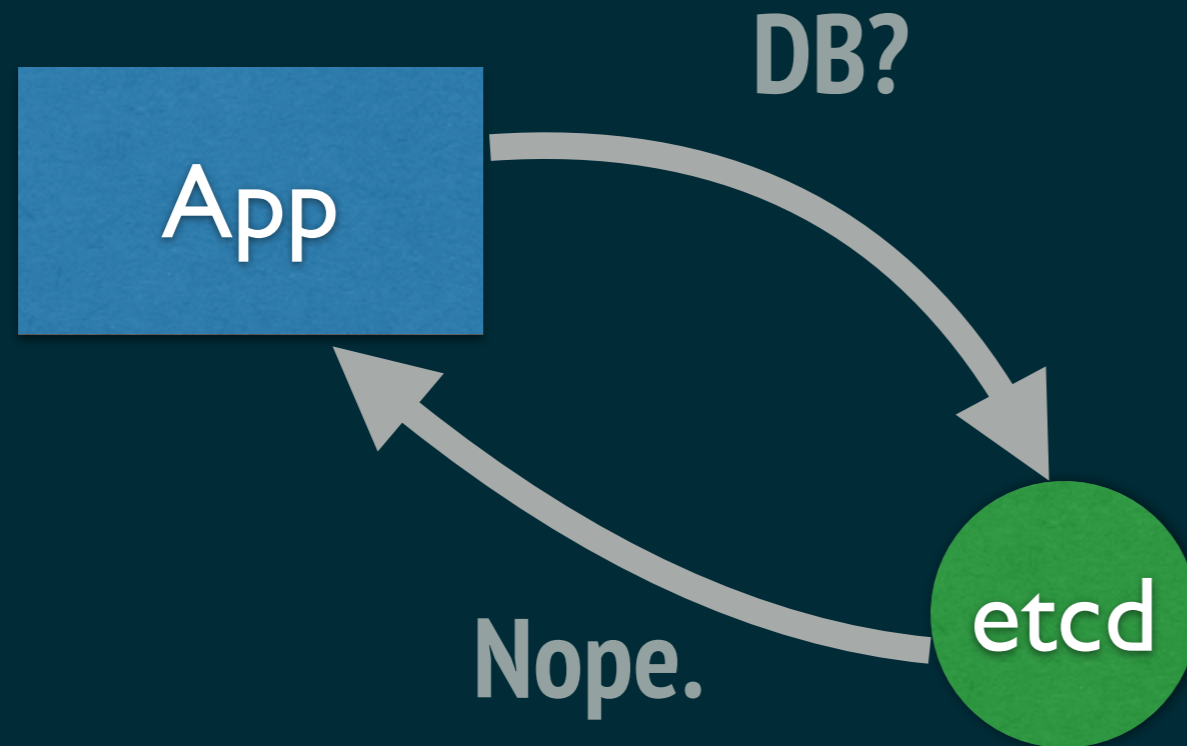
The diagram consists of a blue rectangular box on the left containing the text 'App'. To its right and slightly lower is a green circular shape containing the text 'etcd'. There are no lines or arrows connecting the two shapes.

etcd

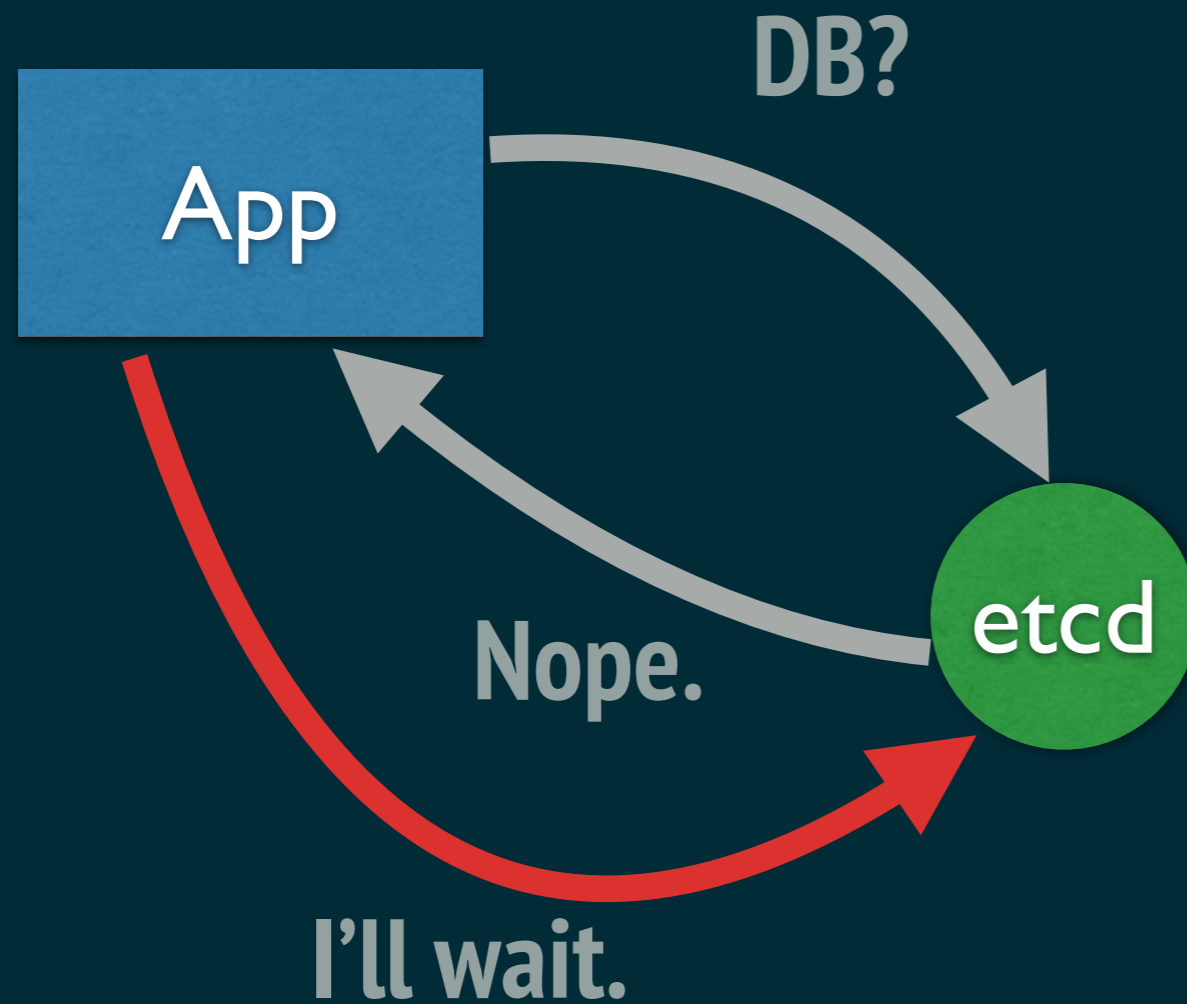
Leader election




Leader election



Leader election




Leader election



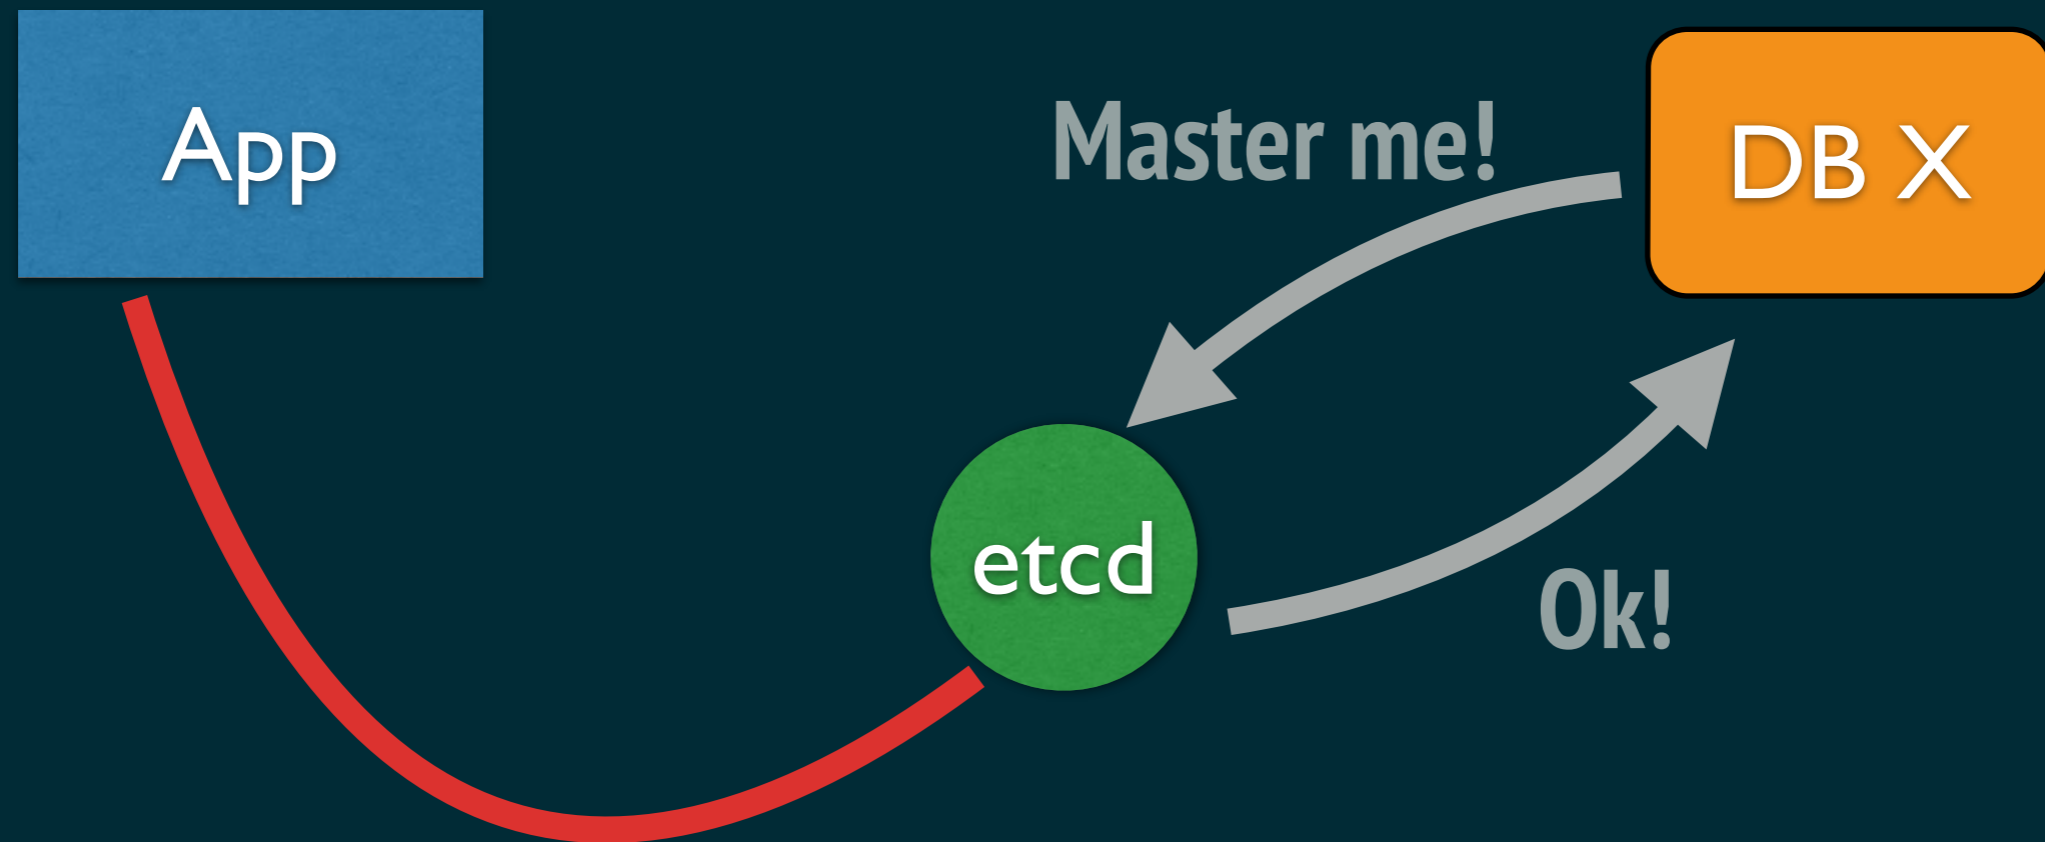
Leader election




Leader election




Leader election




Leader election



Leader election

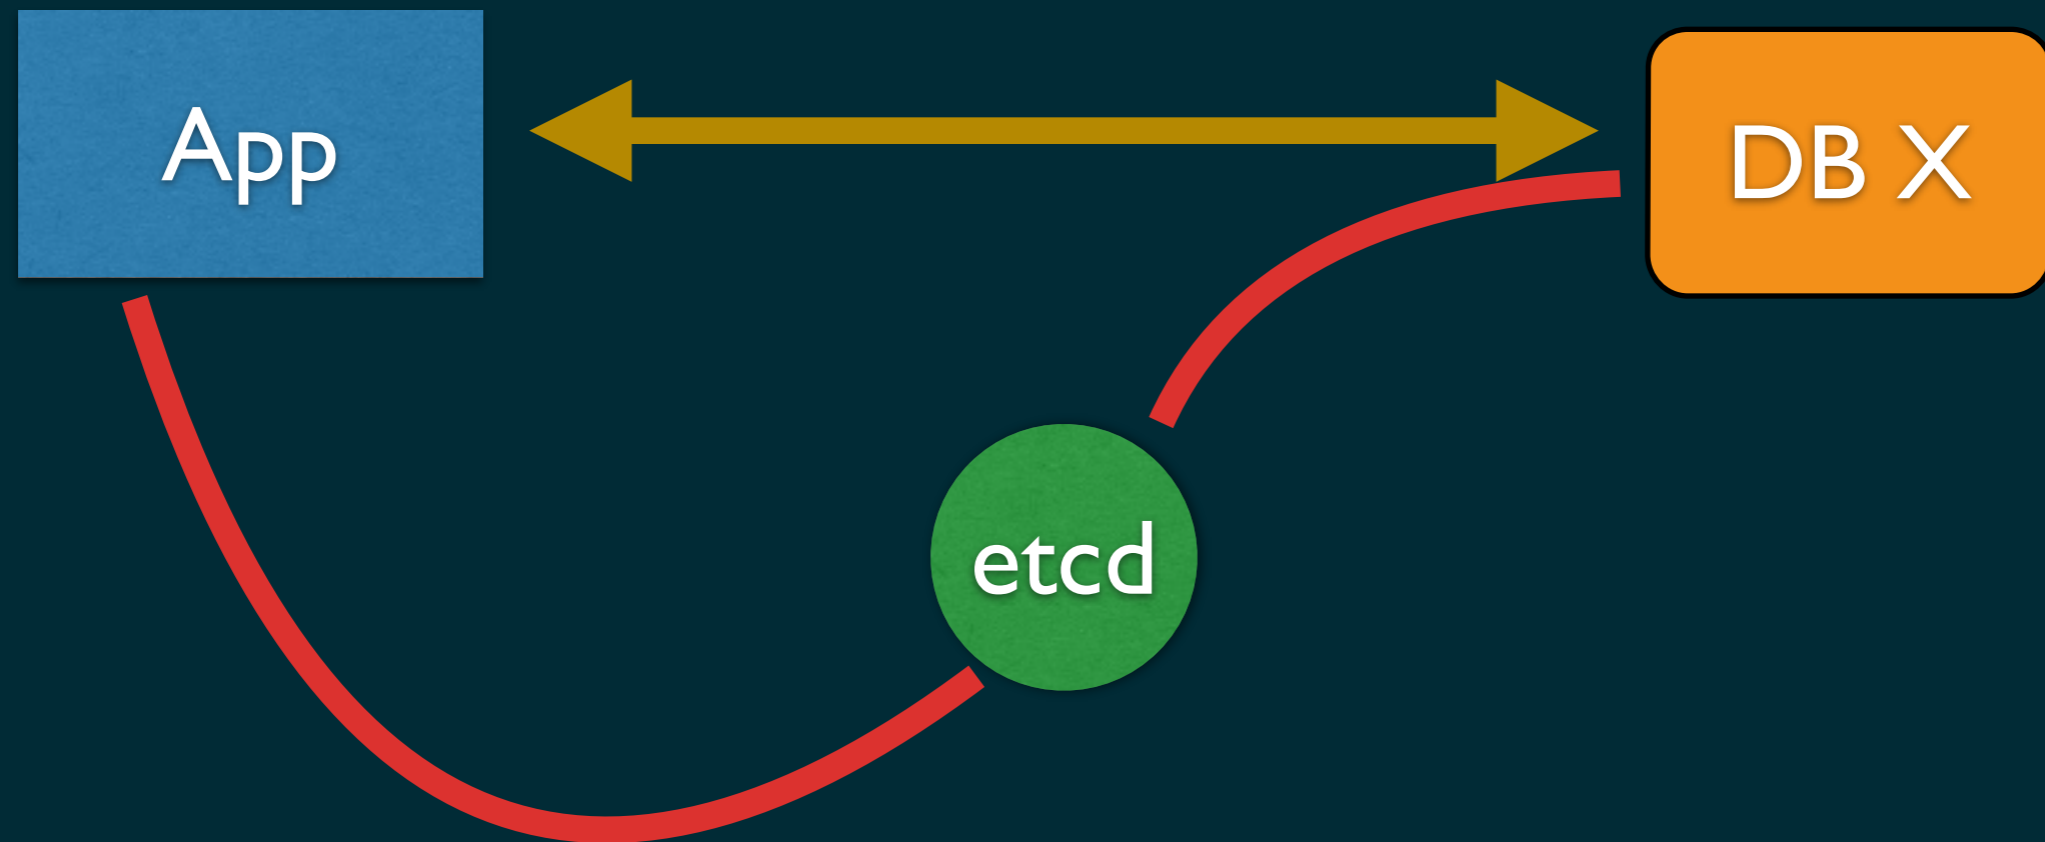


Leader election

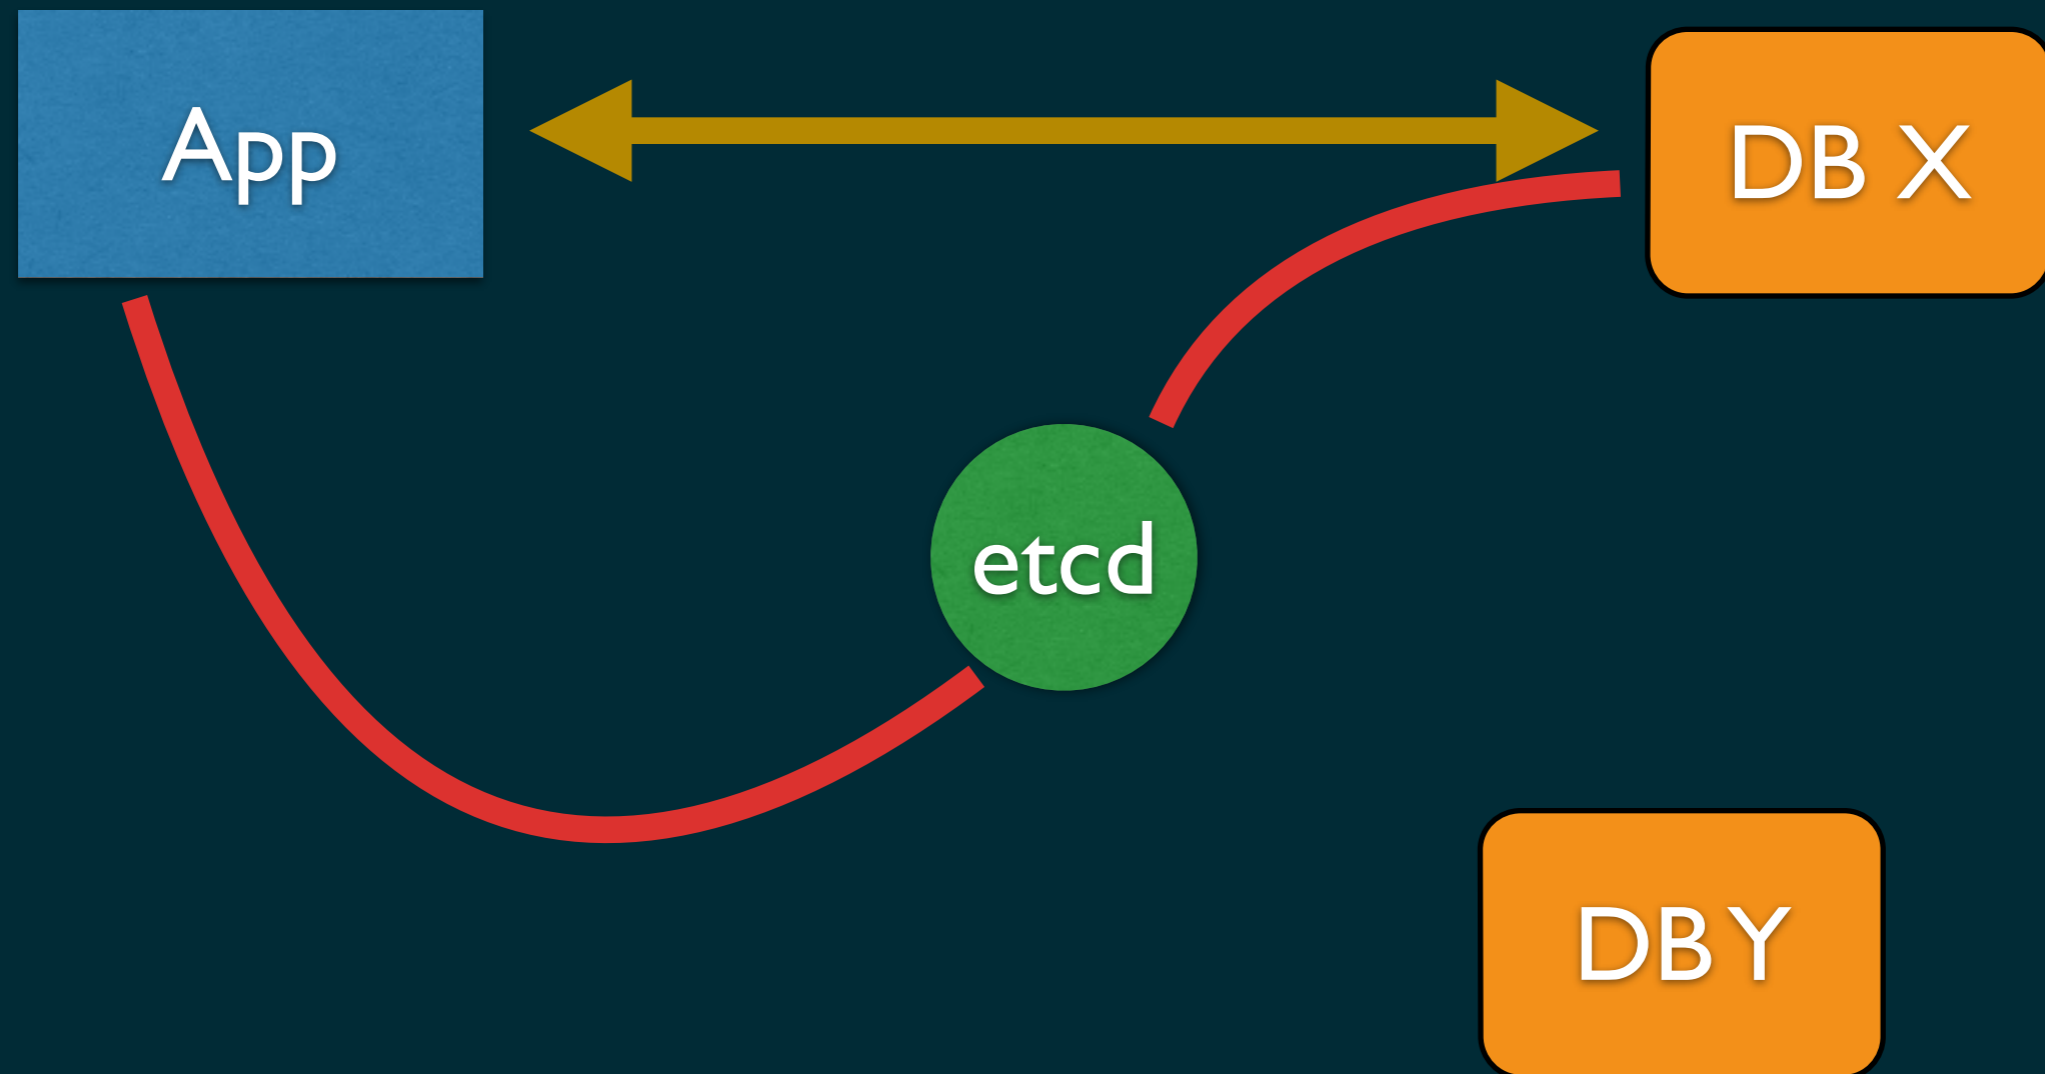


DB X is master.


Leader election



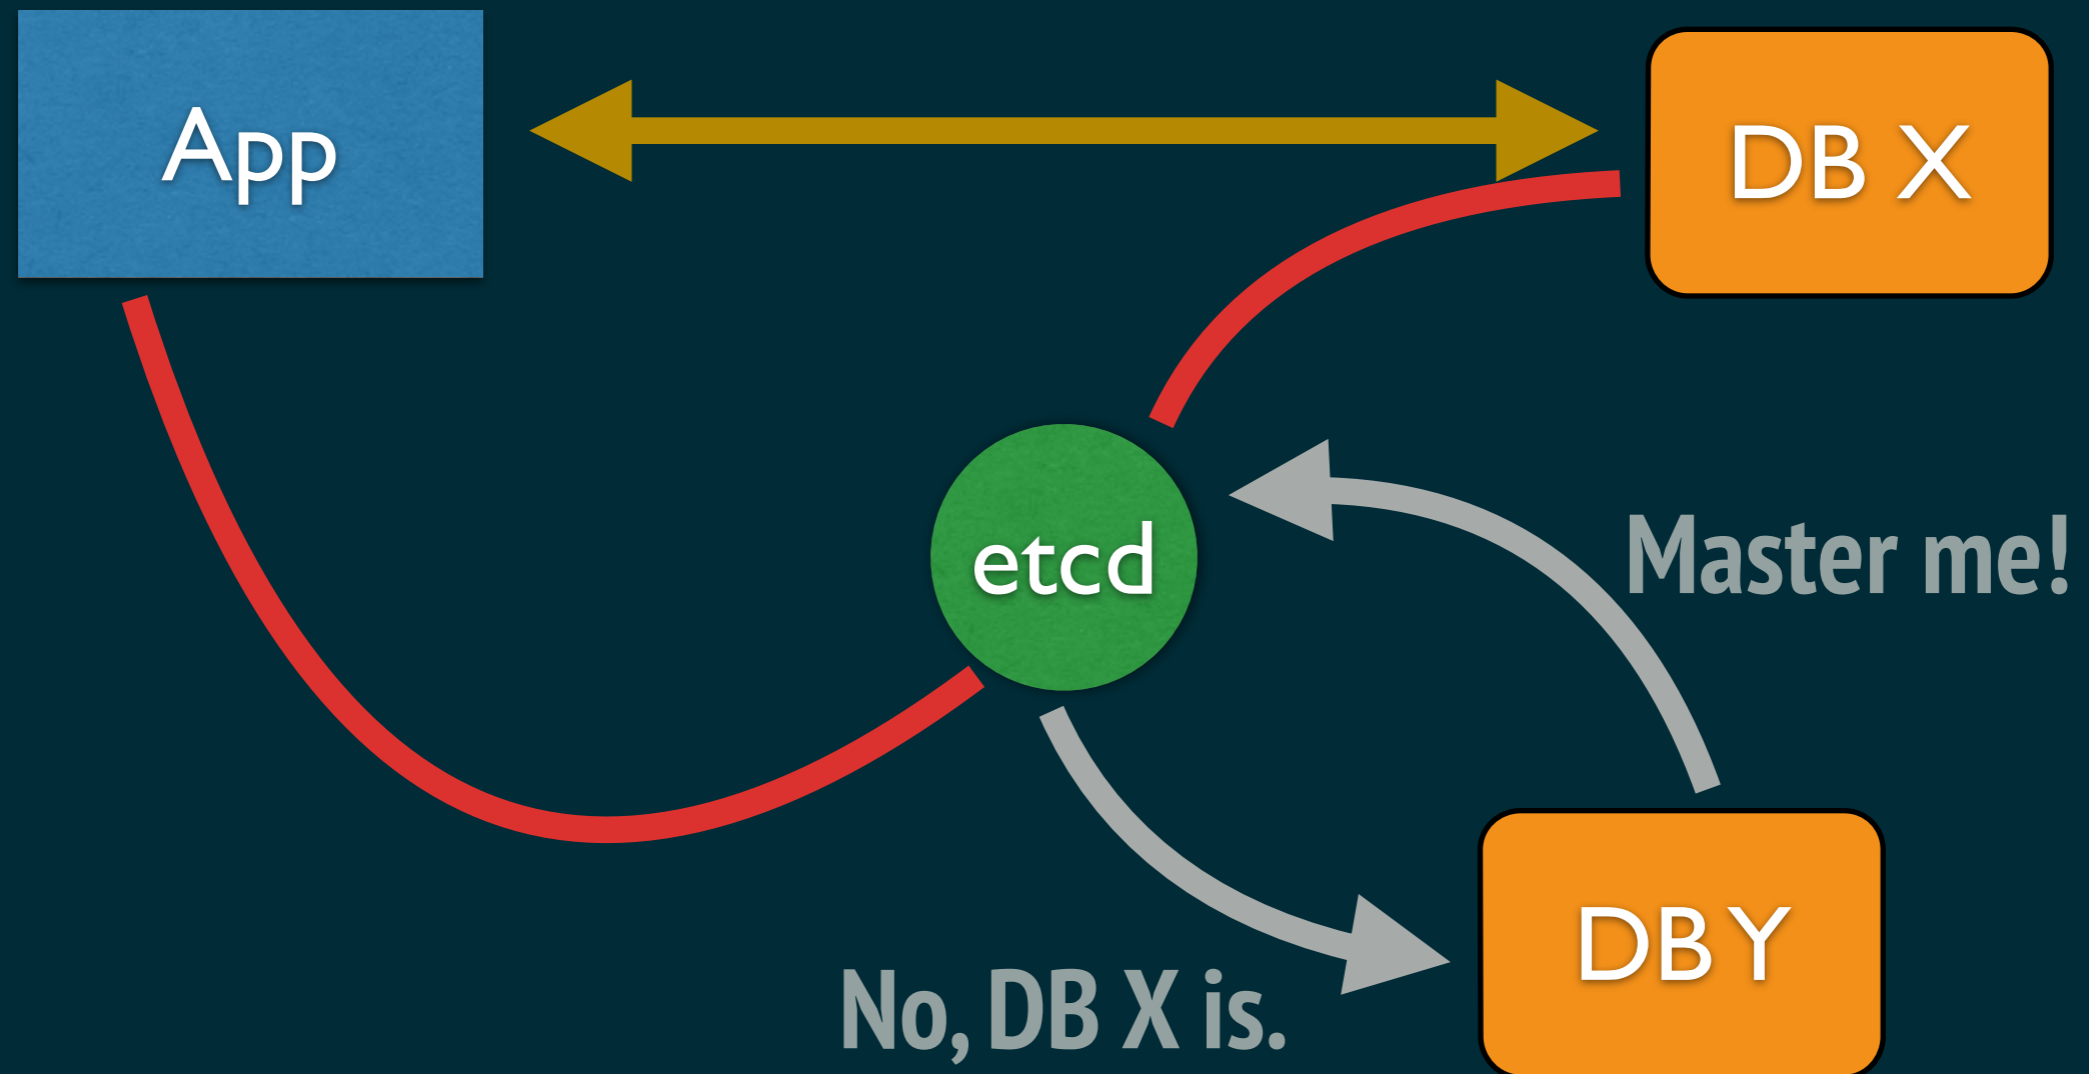
Leader election




Leader election




Leader election




Leader election




Leader election




Leader election



Leader election




Leader election



Leader election



Leader election



Summary

Thanks!

@geewiz

jochen@freistil.it