

Noomaaad

An introduction to cluster schedulers

Patrick O'Connor
@dontrebootme

Lessons learned

2006-2008 Law Offices - service availability

2009-2012 CSU San Bernardino - automation

2013-2014 DreamWorks Animation - schedulers

2015+ The Walt Disney Company - simplicity

Pets

vs

Cattle

Service discovery

- Easy to query
- Service information
- Consistent


Distributed key value

- configuration data
- triggers
- externalize environment from services/code
- sourced from a repo

Cluster schedulers

- Pool resources
- Service supervision
- Scaling functionality
- State information


Cluster schedulers



MARATHON

Docker Swarm

- Simple setup
- Native Docker API
- Can run on other schedulers (e.g. Mesos)
- Tooling compatibility



Amazon ECS

- Runs on *AWS*
- Docker only
- Open source agent, not server
- Simple to get started



Google **Kubernetes**

- kubelet, kube-proxy, etcd, api server, scheduler, repl-controller
- Borg successor
- Has the notion of "pods"
- Very active community



Apache **Mesos** + Mesosphere **Marathon**

- Zookeeper, Leader and Client, Marathon leader
- Multiple frameworks can be deployed on Mesos
- Two-level scheduling



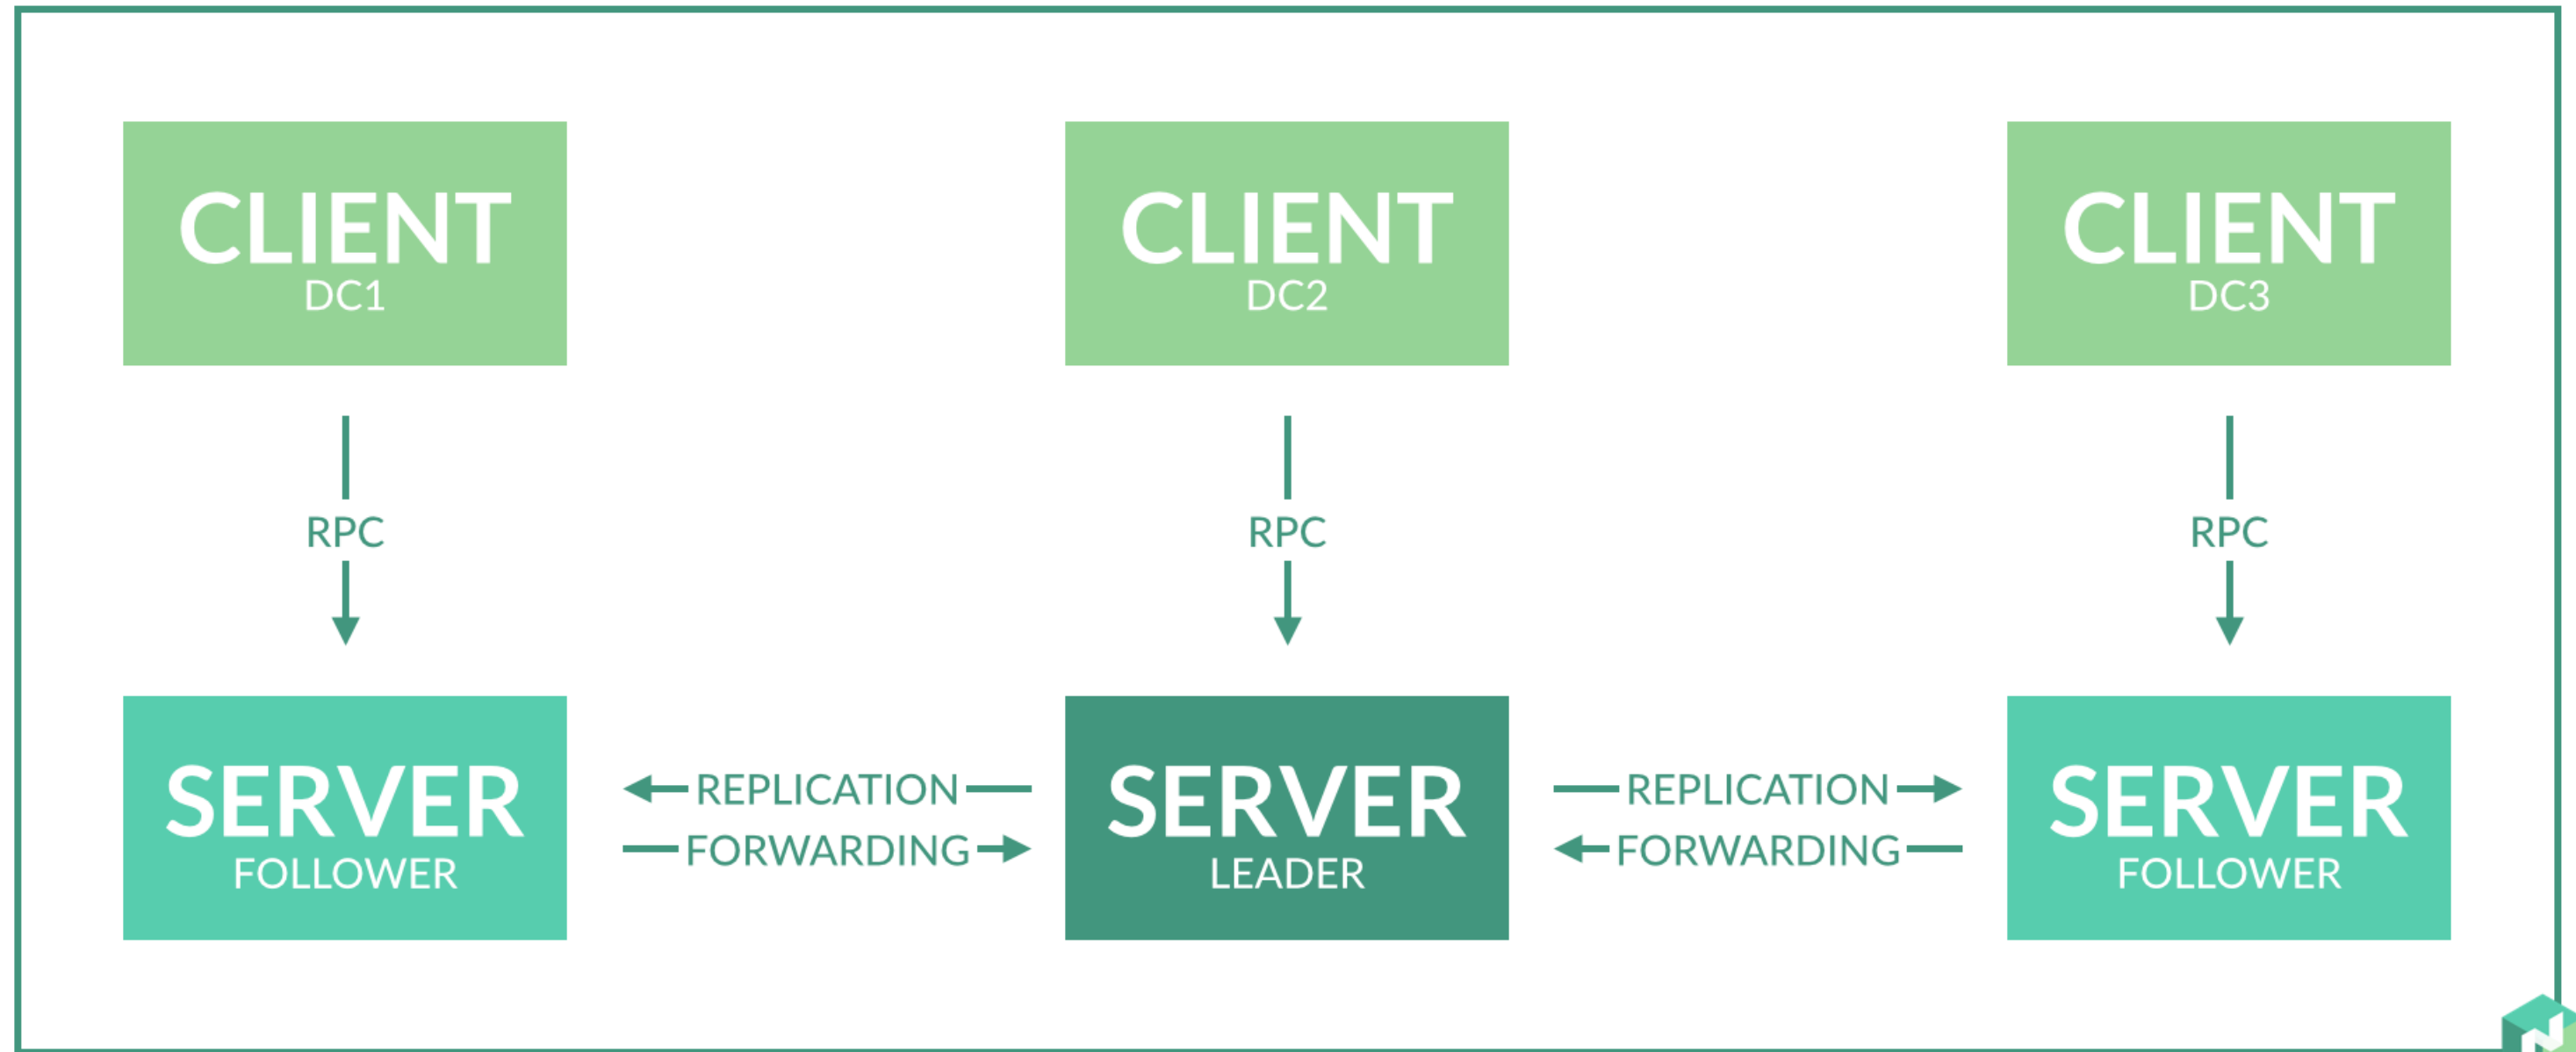
MARATHON

Hashicorp **Nomad**

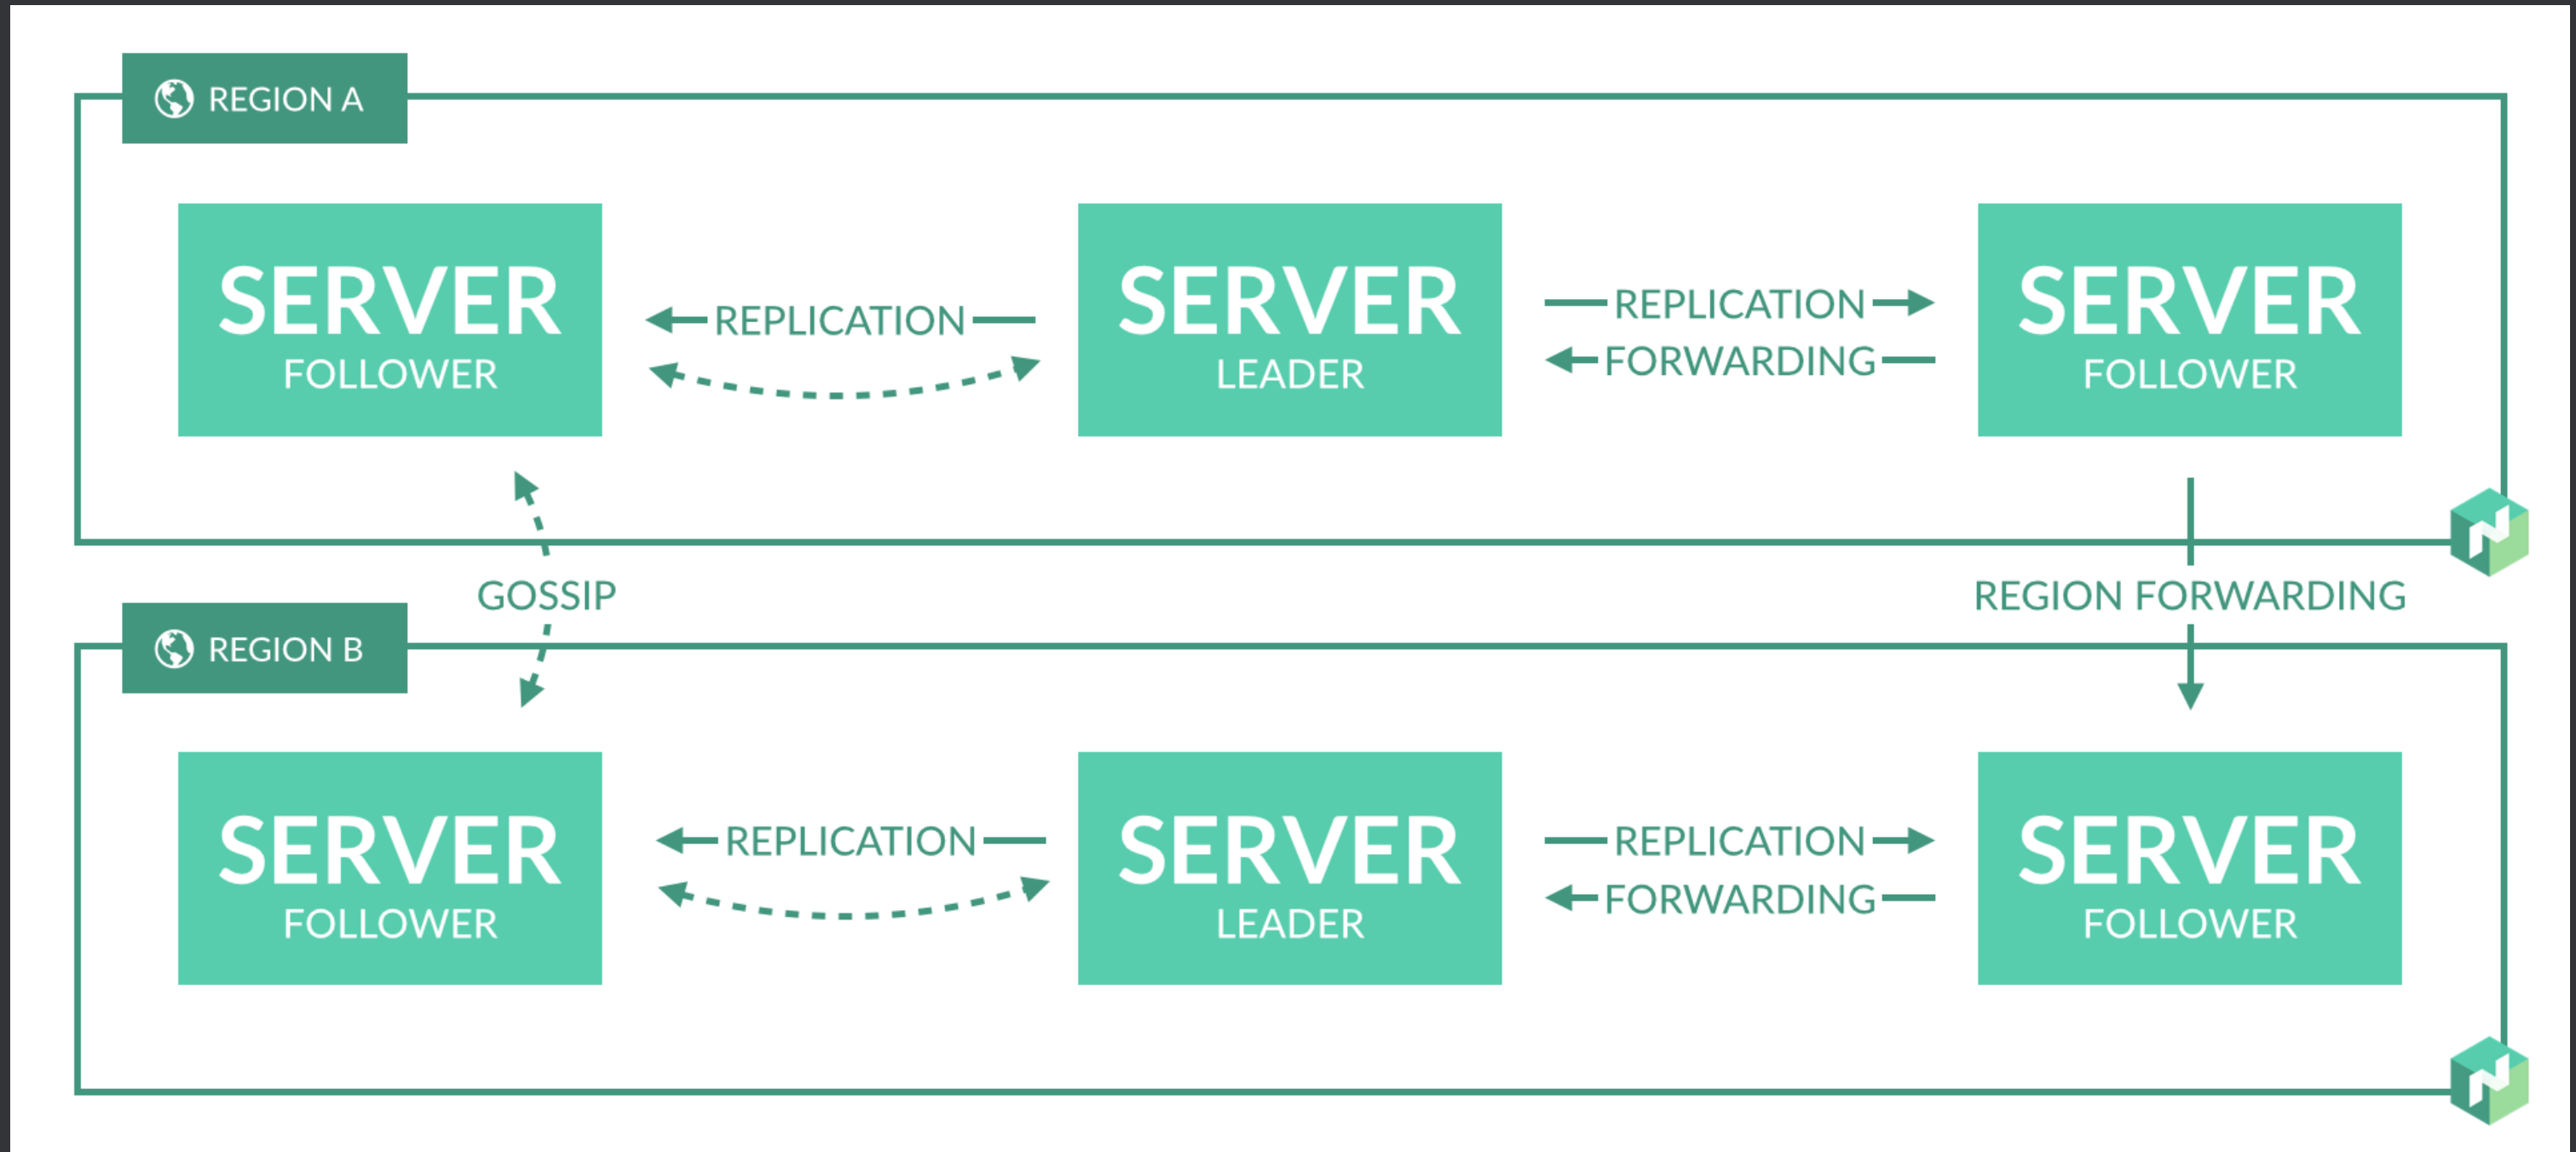
- Single binary for client *and* server
- More than just Docker
- Nomad is distributed, highly available and operationally simple
- Integrates well with Consul



Nomad Architecture *[Region]*



Nomad Architecture *[Global]*



Nomad Components

- Server
- Agent
- Job
- Task

Nomad + Consul

- Native support
- Nomad services -> Consul services (+ health checks)
- Simple

Using **Nomad**

- Deploying services
- Scaling services
- Rolling updates
- Integrating with Consul
- Load balancing deployed services

Job definition

```
job "microbot" {  
  region = "global"  
  datacenters = ["prod-east1", "prod-west1"]  
  group "demo-stack" {  
    task "microbot" {  
      driver = "docker"  
    },  
    task "microbot-db" {  
      driver = "docker"  
    }  
    ...  
  }  
}
```

Nomad is

- Operationally simple
- Distributed and highly available
- Open source
- Multi datacenter/region

Noomaaad

An Introduction to Cluster Schedulers

Patrick O'Connor
@dontrebootme