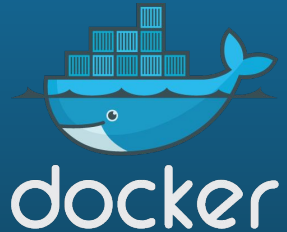


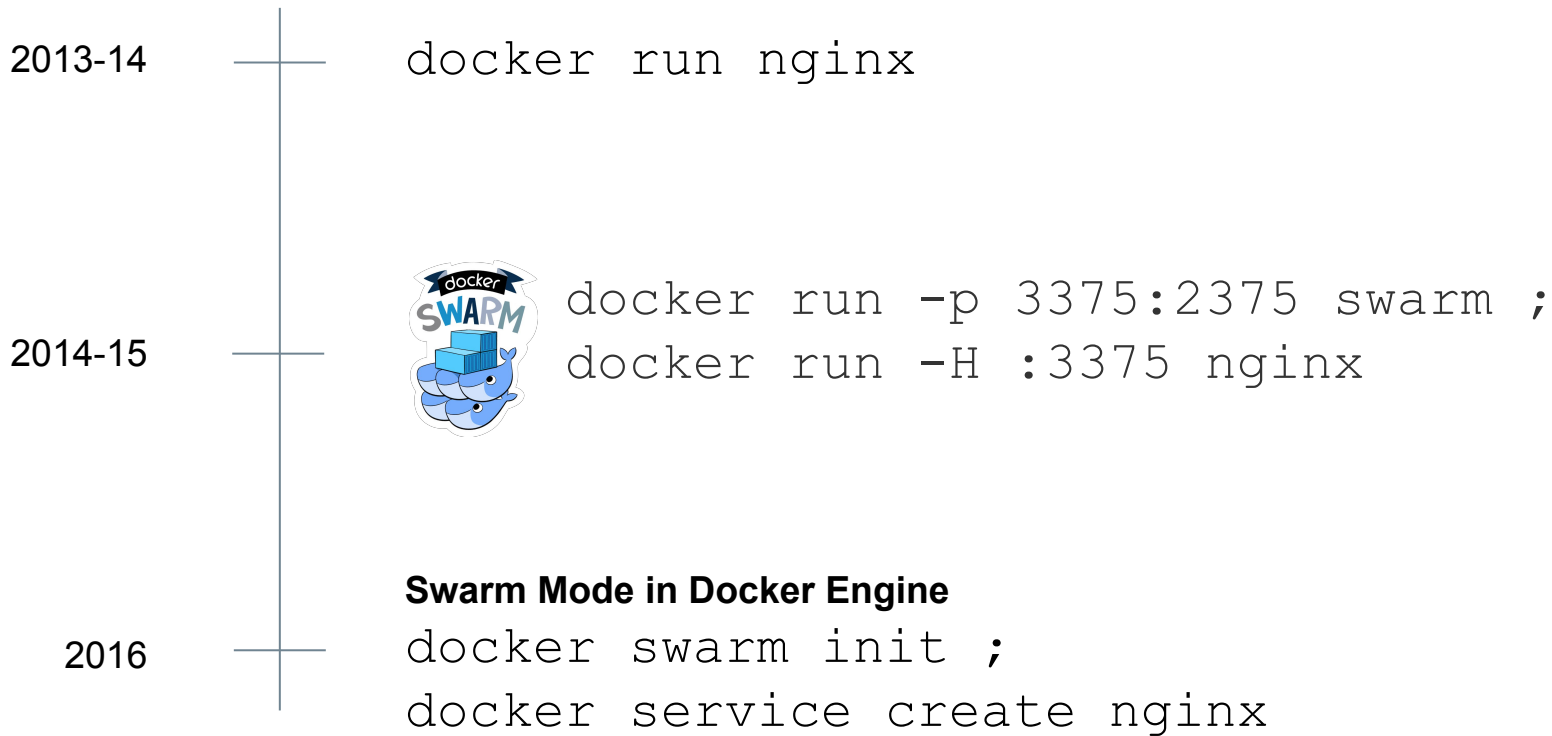
docker service is the new docker run

Getting Started with Docker Clustering

Mike Goelzer / mgoelzer@docker.com / @mgoelzer
Docker Inc.



docker service is the new docker run



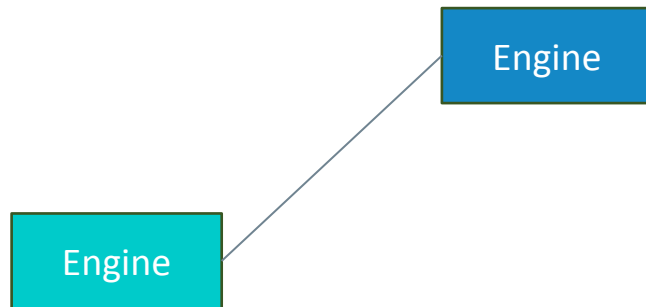
Features Walkthrough

Swarm Mode

Engine

```
■ $ docker swarm init
```

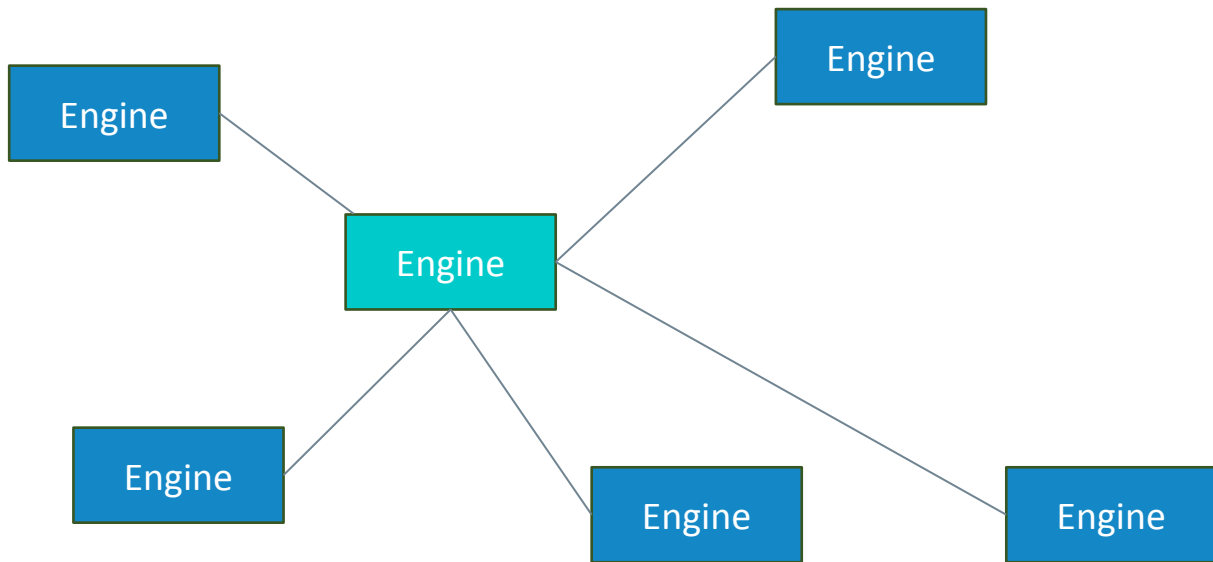
Swarm Mode




■ `$ docker swarm init`

■ `$ docker swarm join <IP of manager>:2377`

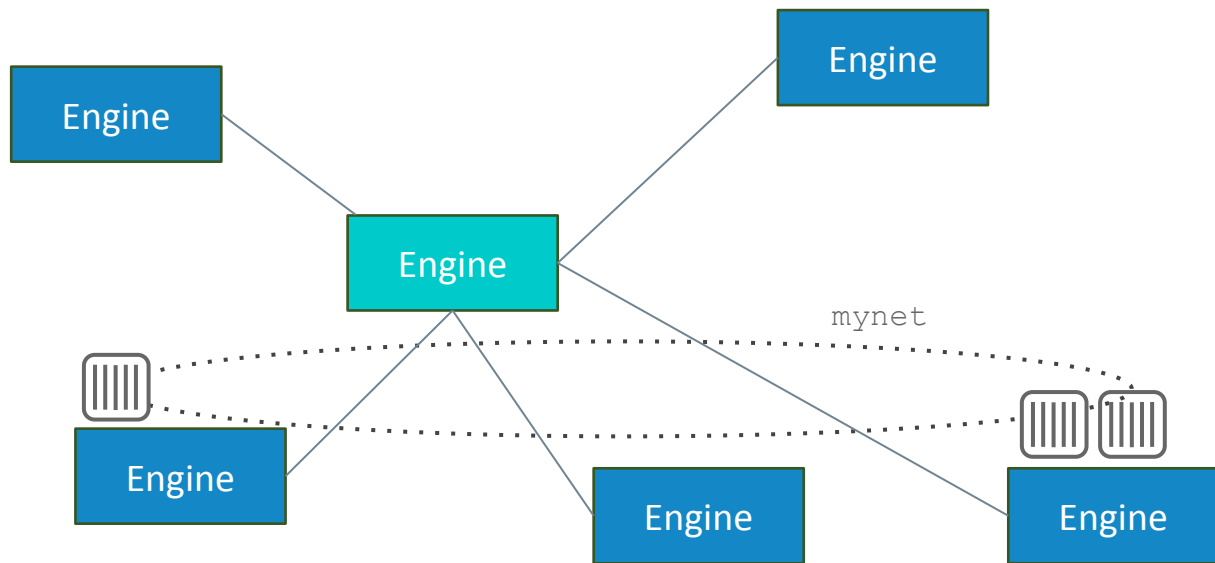
Swarm Mode



 `$ docker swarm init`

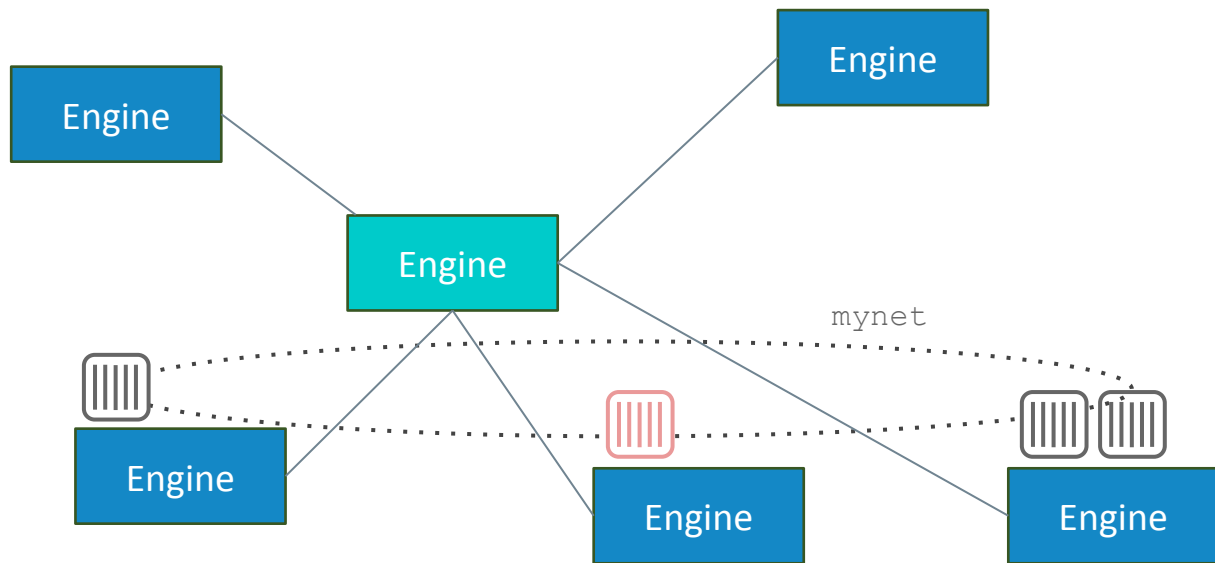
 `$ docker swarm join <IP of manager>:2377`

Services



```
🗄️ $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

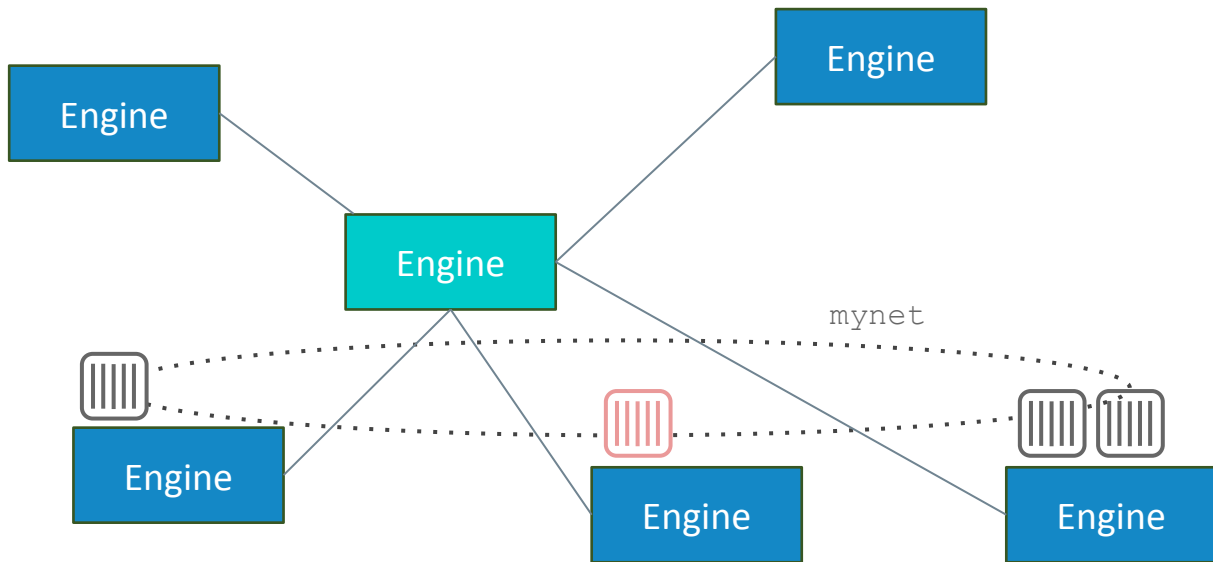
Services



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

```
📦 $ docker service create --name redis --network mynet redis:latest
```

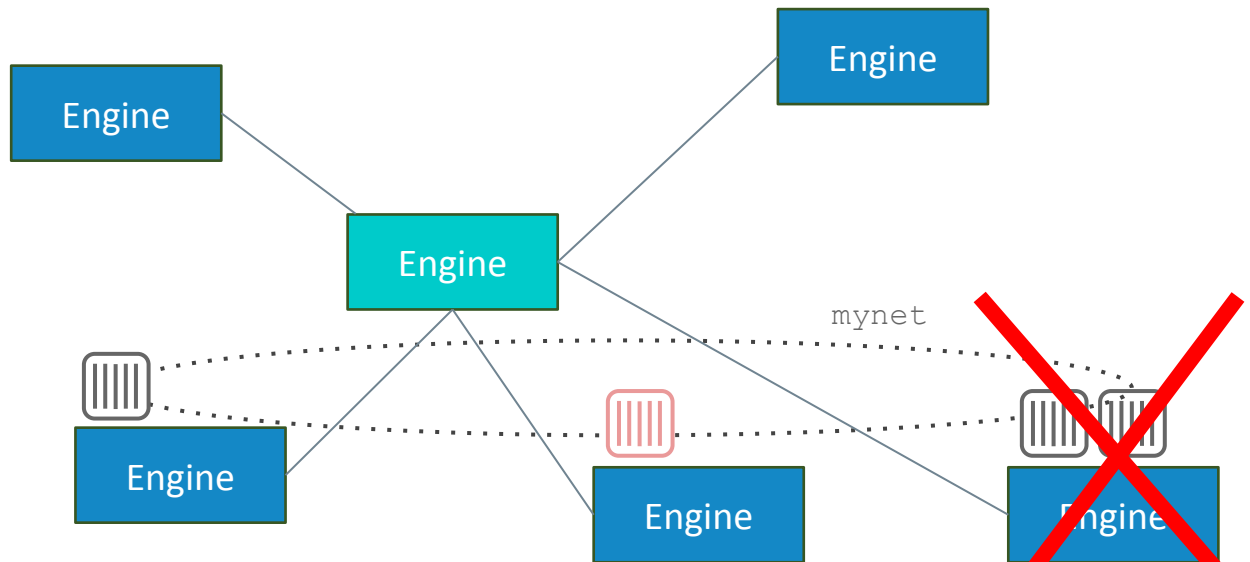

Node Failure



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

```
📦 $ docker service create --name redis --network mynet redis:latest
```

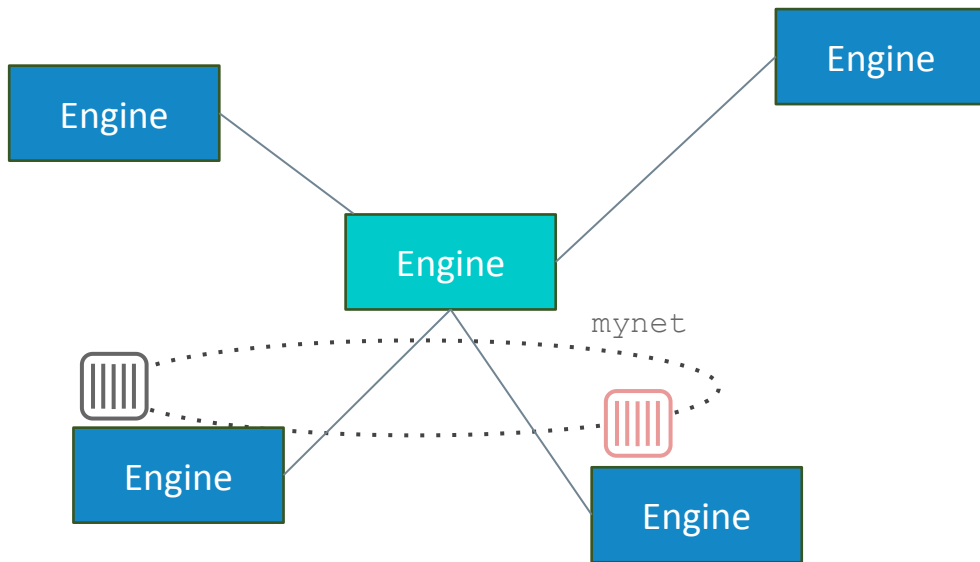
Node Failure



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

```
📦 $ docker service create --name redis --network mynet redis:latest
```

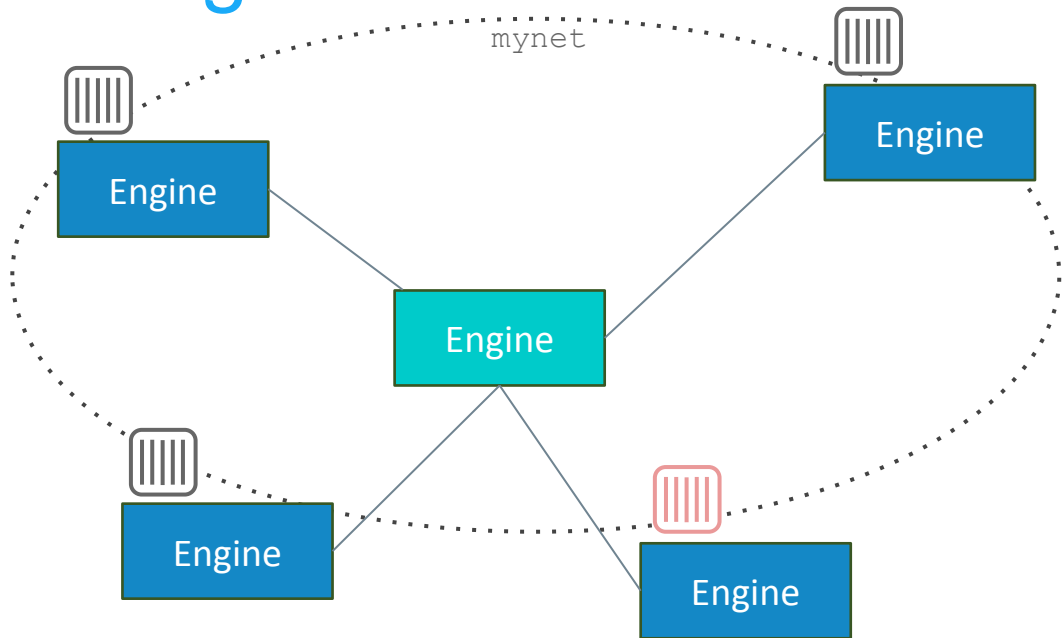
Desired State \neq Actual State



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

```
📦 $ docker service create --name redis --network mynet redis:latest
```

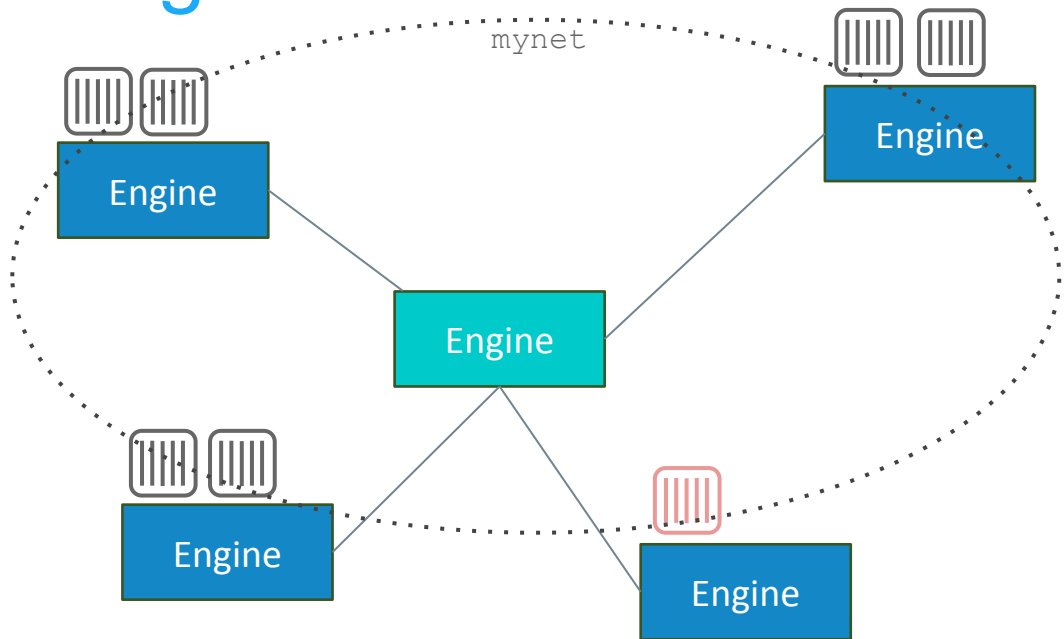
Converge Back to Desired State



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

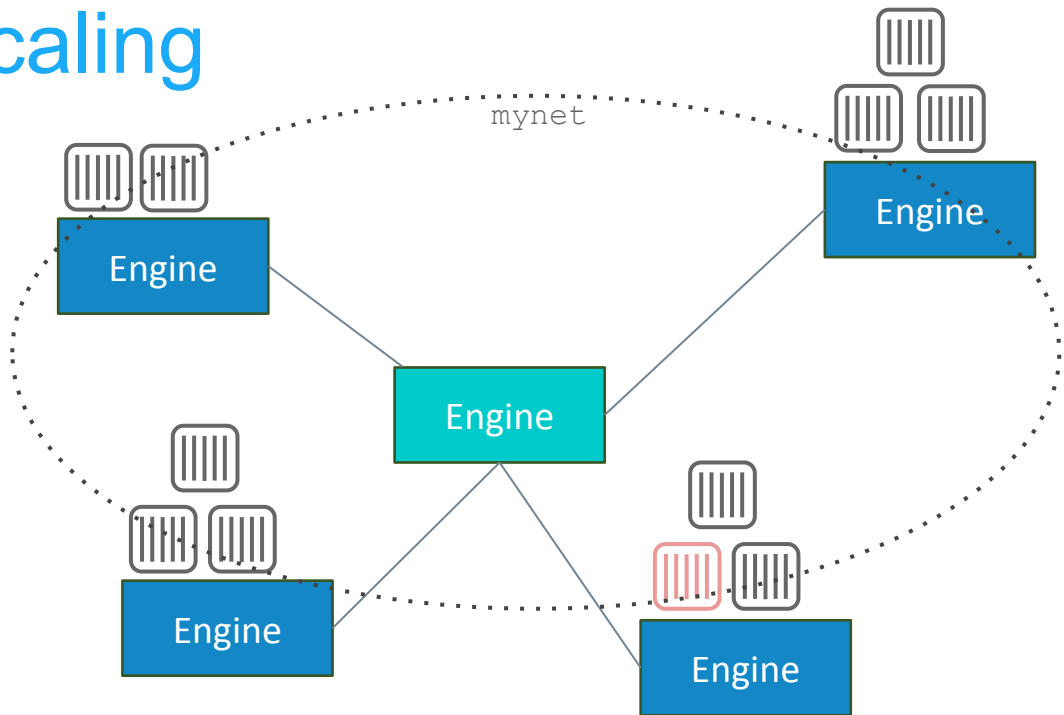
```
📦 $ docker service create --name redis --network mynet redis:latest
```

Scaling



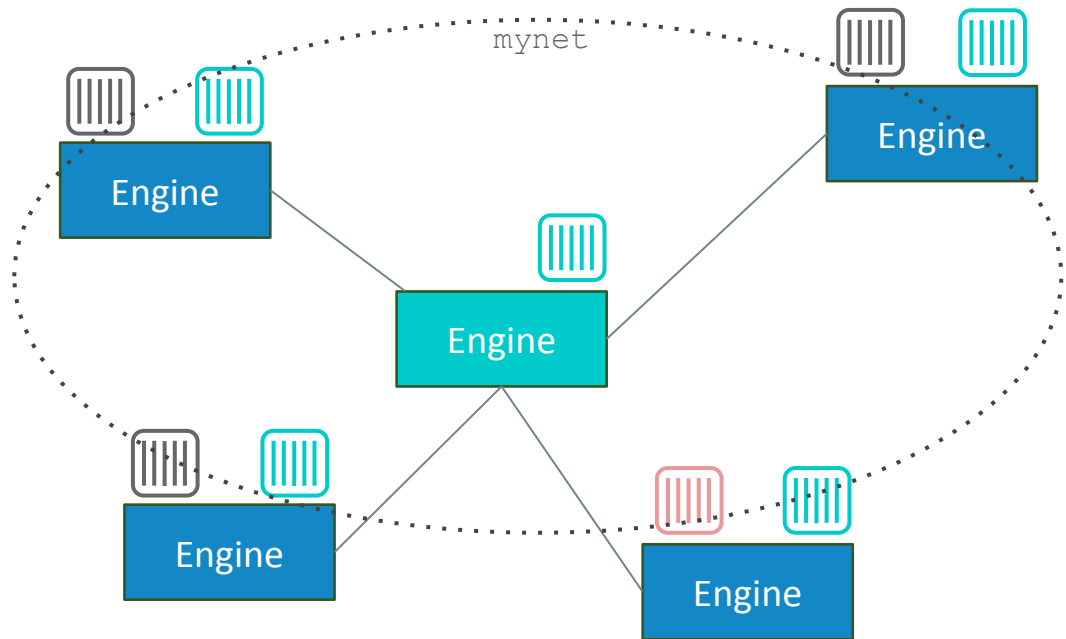
```
$ docker service update --replicas 6 frontend
```


Scaling



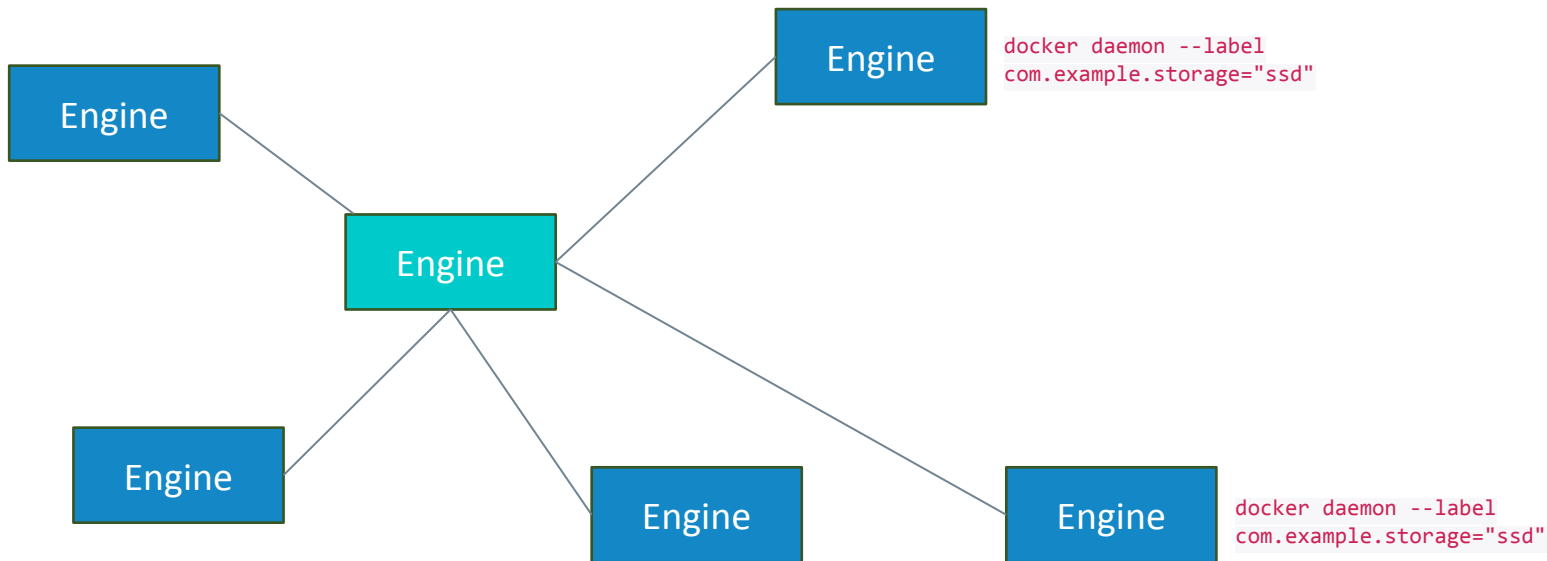
```
$ docker service update --replicas 10 frontend
```

Global Services

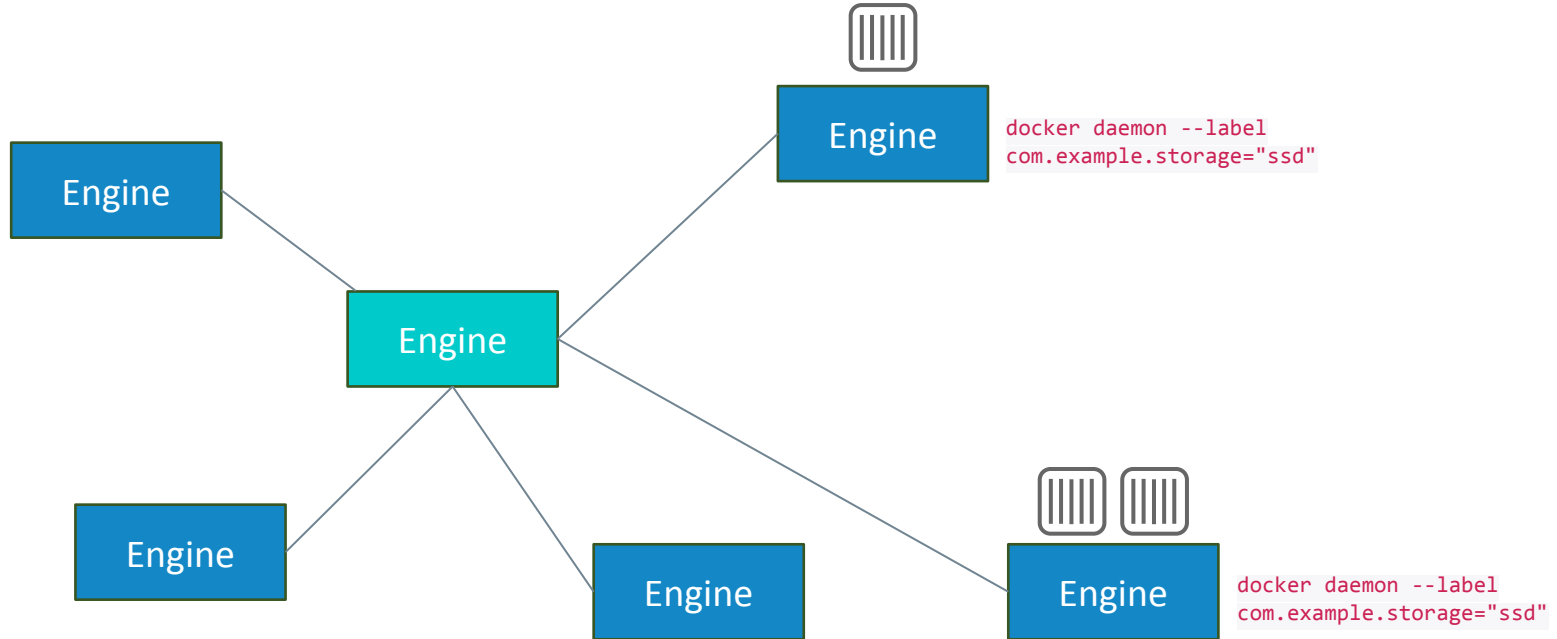


```
 $ docker service create --mode=global --name prometheus  
prom/prometheus
```

Constraints

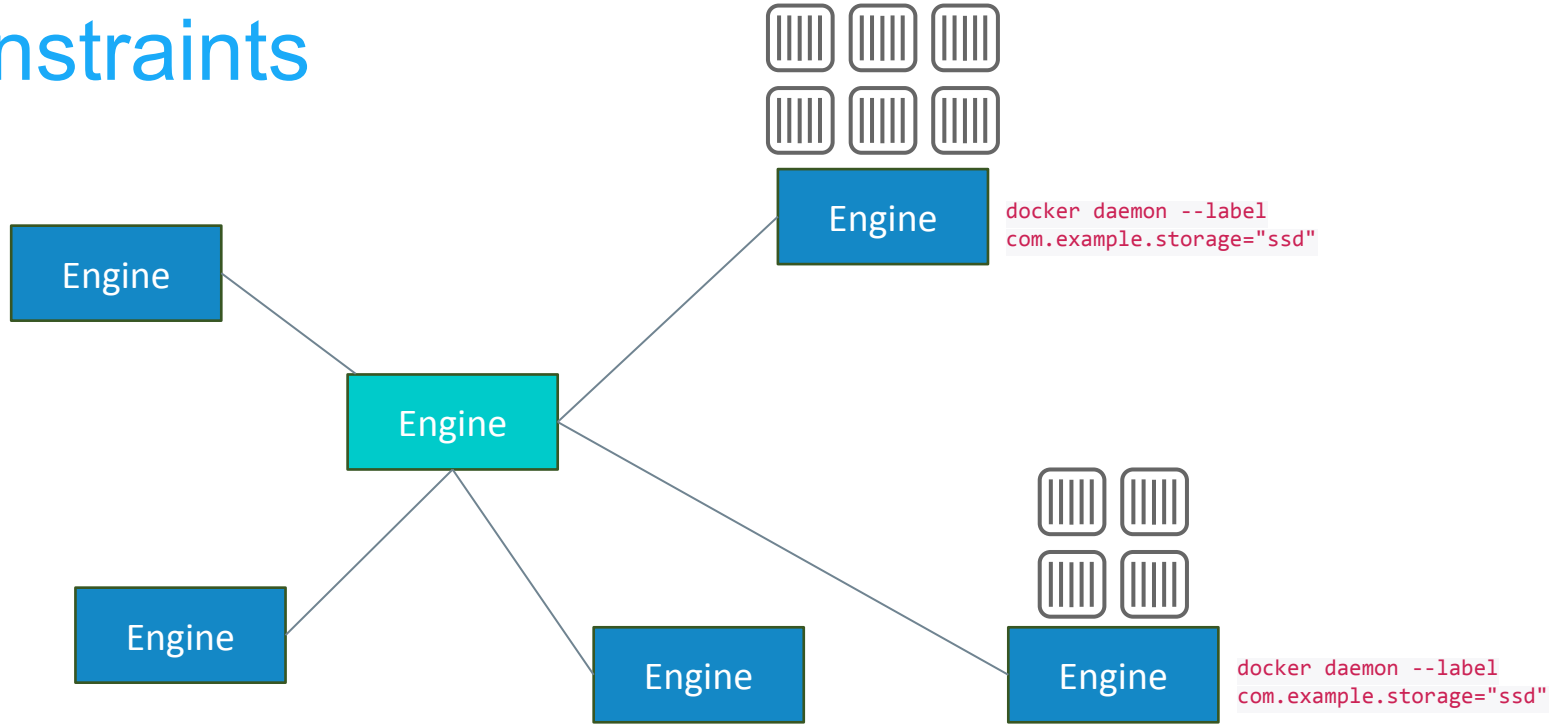


Constraints



```
🗄️ $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 --constraint engine.labels.com.example.storage==ssd  
frontend:latest
```

Constraints



```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 --constraint engine.labels.com.example.storage==ssd  
frontend:latest  
$ docker service update --replicas 10 frontend
```

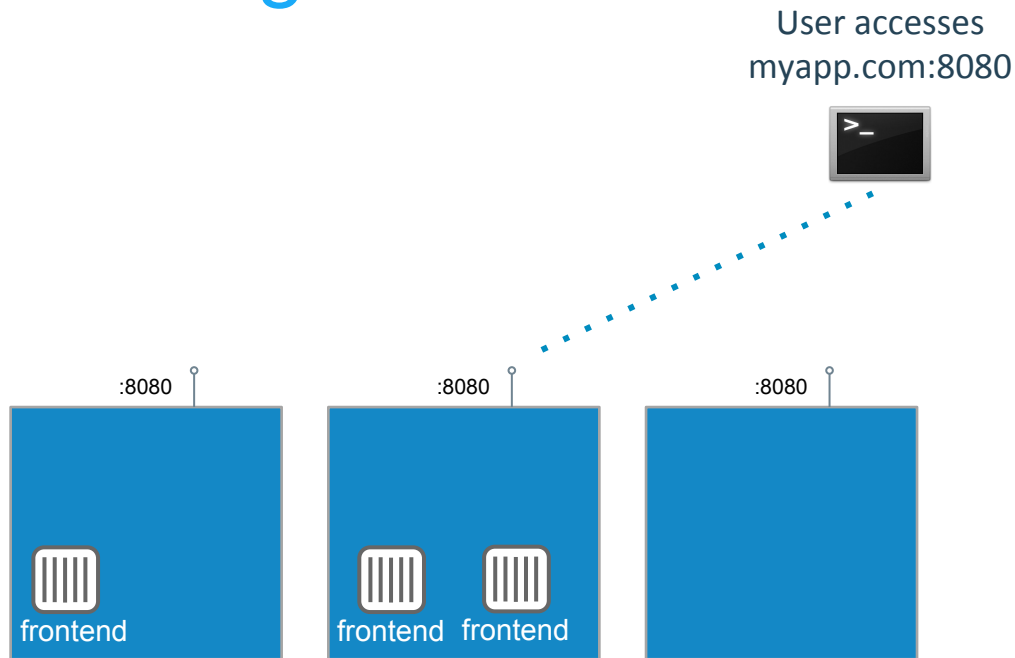
Container Health Check in Dockerfile

```
HEALTHCHECK --interval=5m --timeout=3s  
  --retries 3  
CMD curl -f http://localhost/ || exit 1
```

Check web server every 5 minutes, require < 3 sec latency.
>= 3 consecutive failures sets unhealthy state

Coming soon: health checks in official images

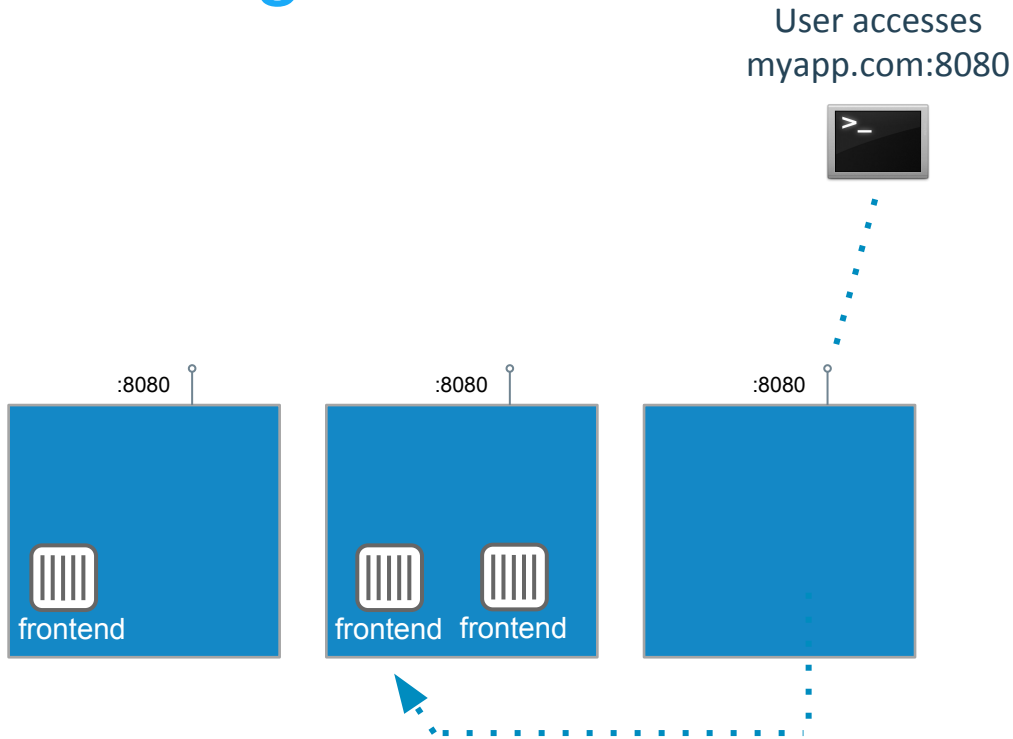
Routing Mesh



- Operator reserves a swarm-wide ingress port (8080) for `myapp`
- Every node listens on 8080
- Container-aware routing mesh can transparently reroute traffic from Worker3 to a node that is running container
- Built in load balancing into the Engine
- DNS-based service discovery

```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend:latest
```

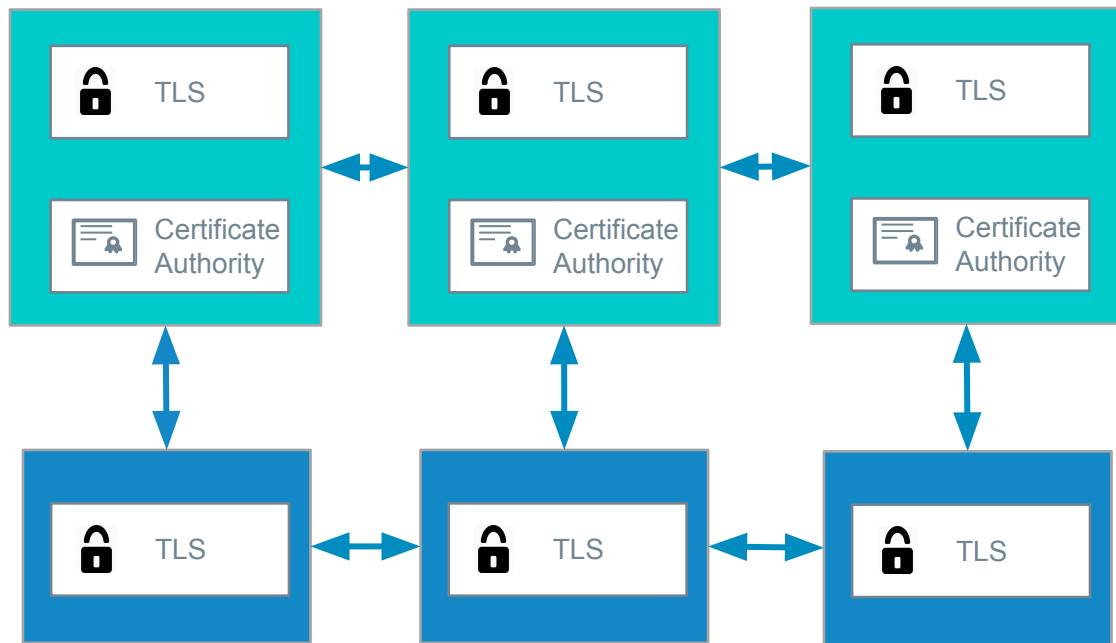
Routing Mesh: Published Ports



- Operator reserves a swarm-wide ingress port (8080) for `myapp`
- Every node listens on 8080
- Container-aware routing mesh can transparently reroute traffic from third node to a node that is running container
- Built in load balancing into the Engine
- DNS-based service discovery

```
📦 $ docker service create --replicas 3 --name frontend --network mynet  
-p 8080:80 frontend_image:latest
```

Secure by default with end-to-end encryption



- Out-of-the-box TLS encryption and mutual auth
- Automatic cert rotation
- External or self-signed root CA
- Cryptographic node identity

Scale: 2,000 Nodes and Counting

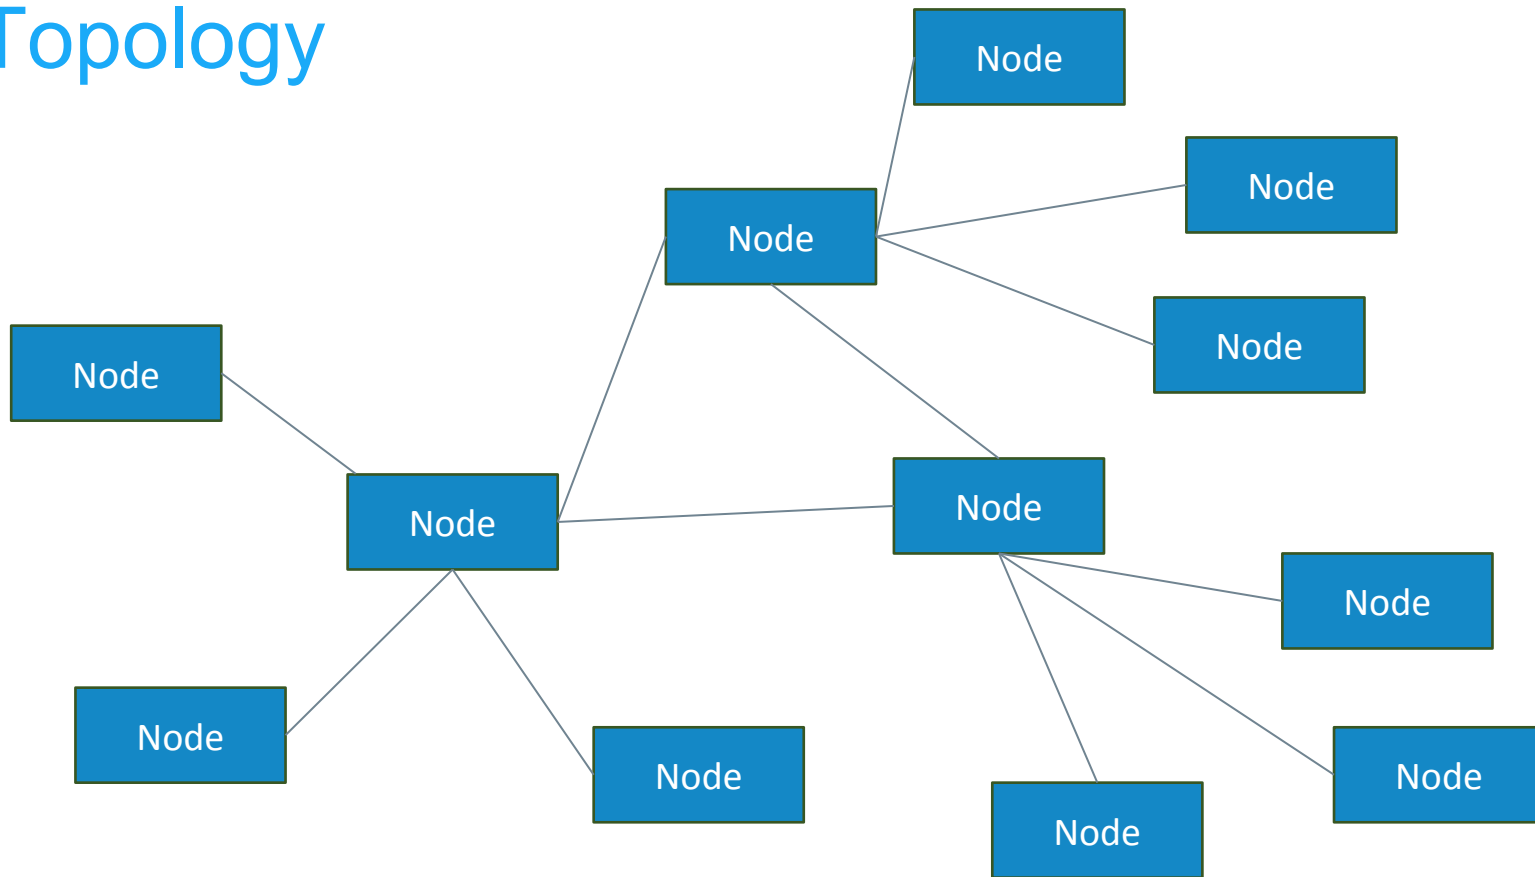
- For now: community testing, crowd-sourced nodes, not funded by Docker
- Credit to: Chanwit Kaewkasi, Suranaree University of Technology (SUT), Thailand
- Results:
 - 2,384 nodes
 - 96,287 containers
 - Manager CPU/memory $\lesssim 15\%$
 - Test stopped because 3rd-party monitoring failed
- <https://github.com/swarm2k/swarm2k>



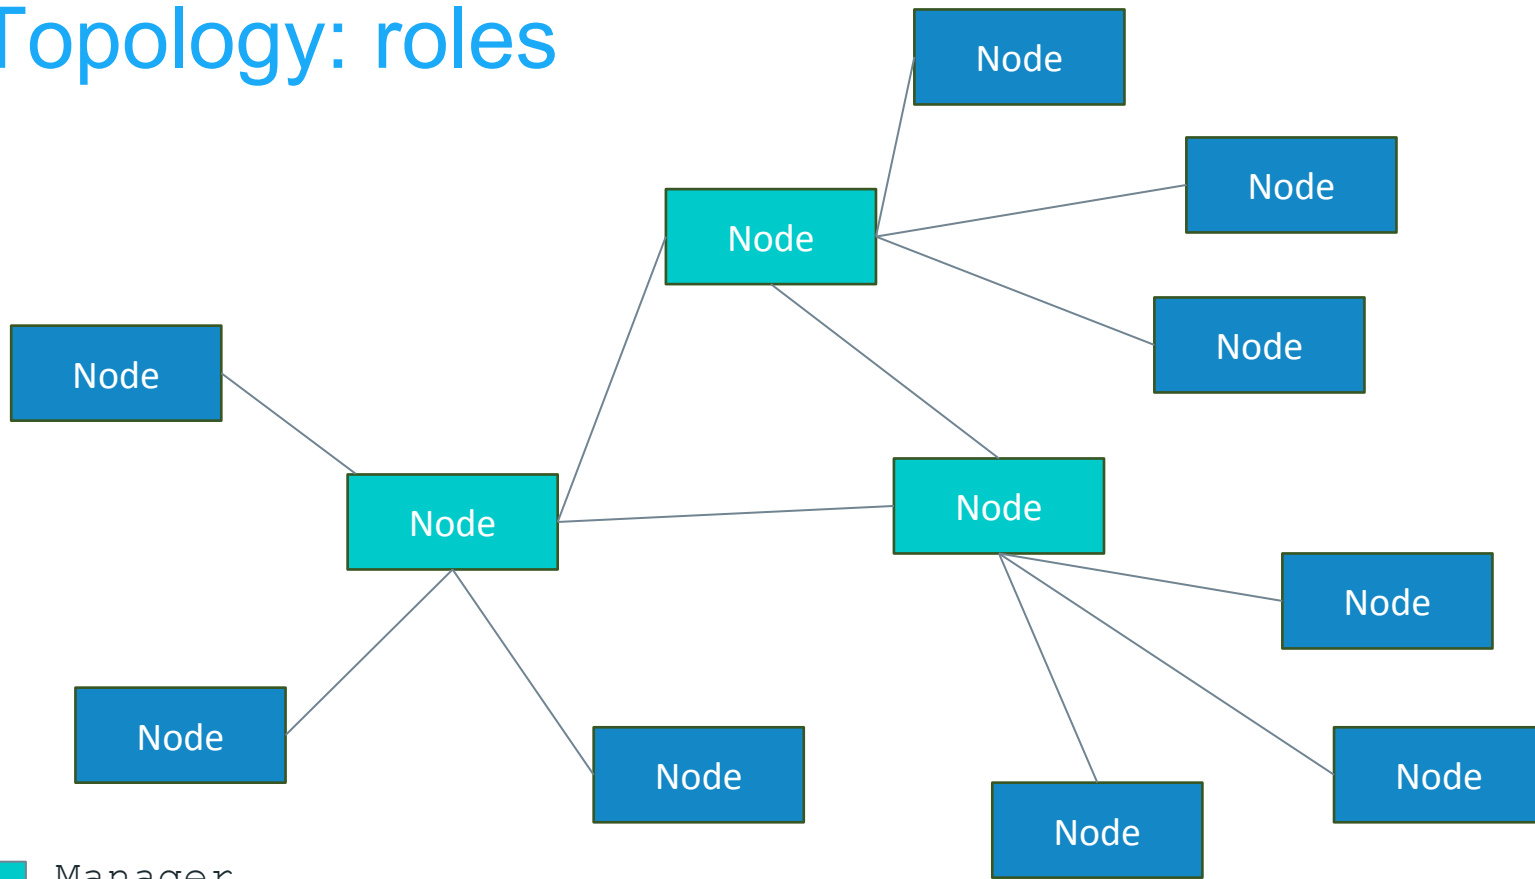
@chanwit

Deep Dive: Topology

Topology



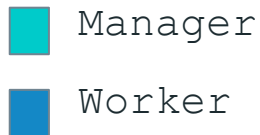
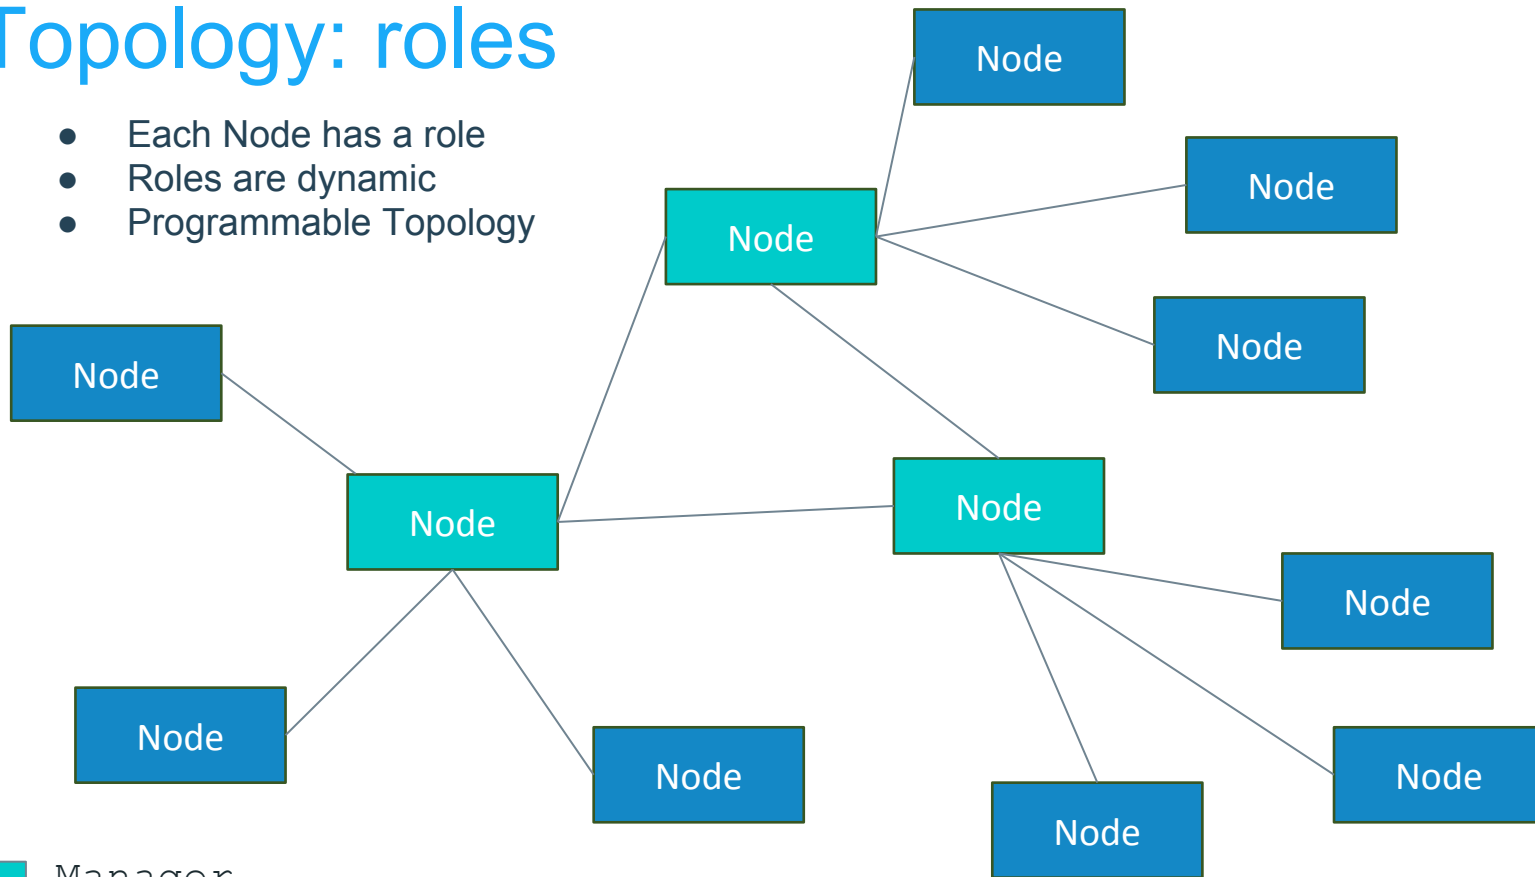
Topology: roles



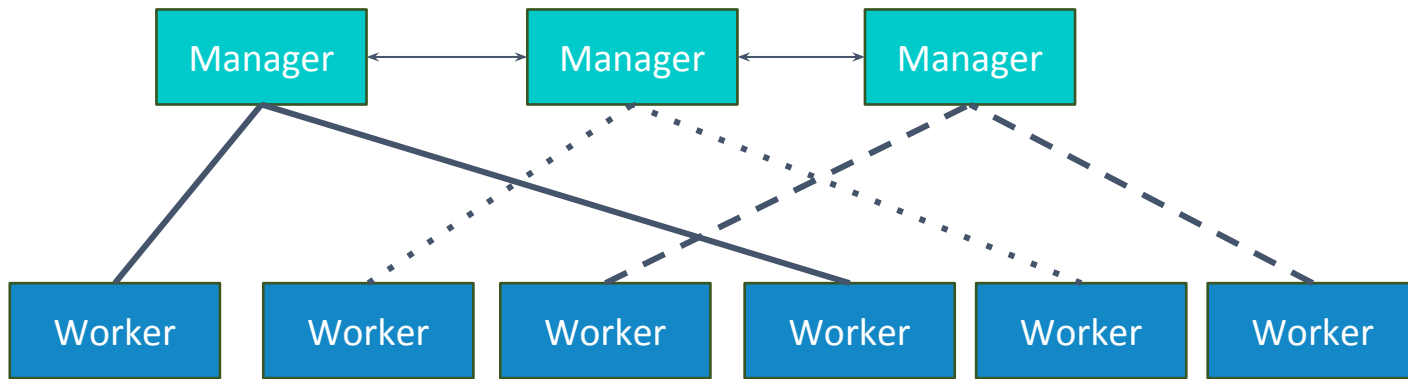
- Manager
- Worker

Topology: roles

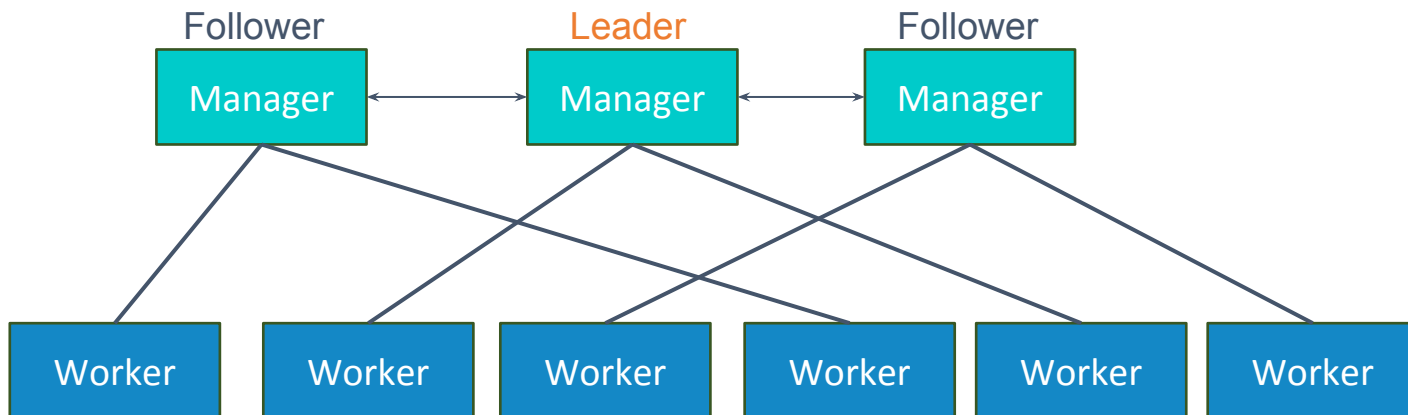
- Each Node has a role
- Roles are dynamic
- Programmable Topology



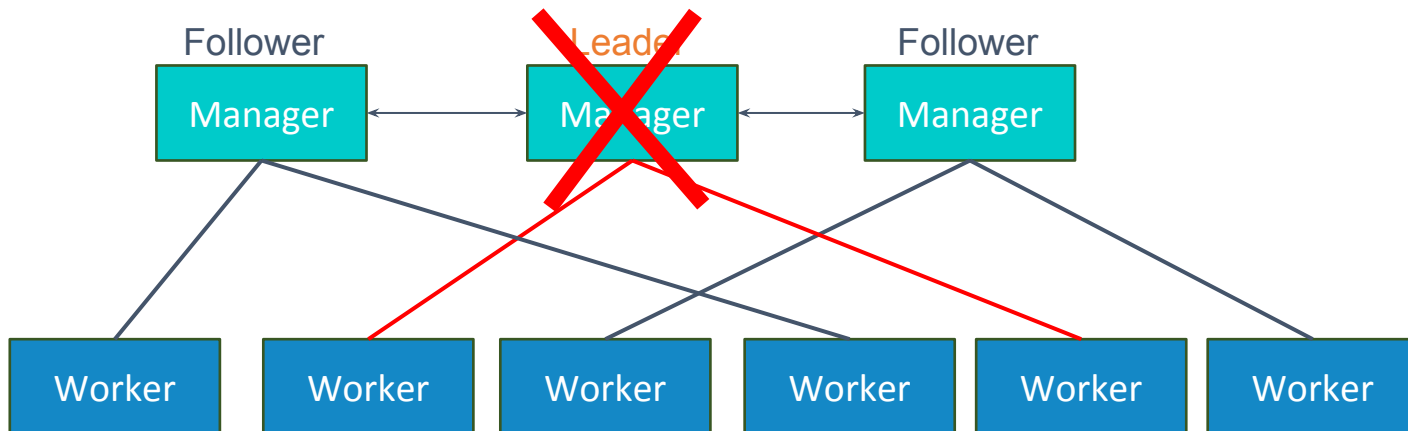
Topology: scaling model



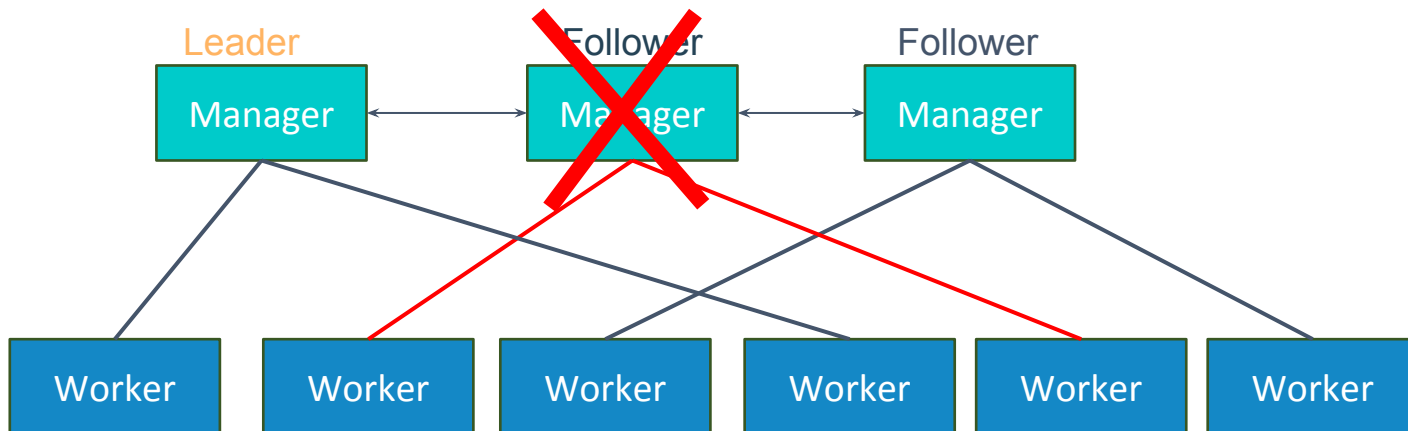
Topology: High Availability



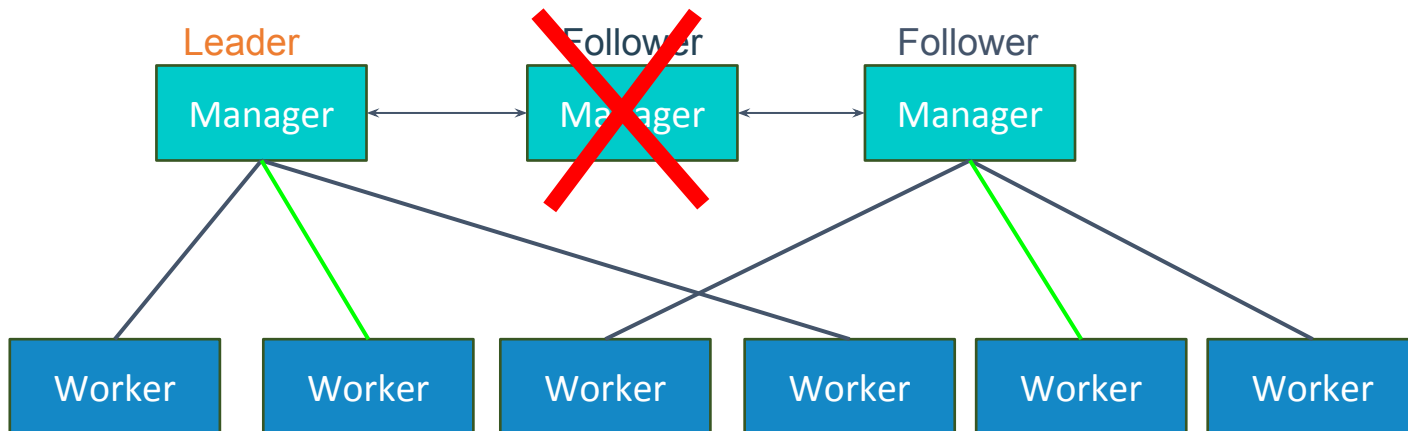
Topology: High Availability



Topology: High Availability



Topology: High Availability



DEMO



docker

Victor Vieux

vieux@docker.com / @vieux

Mike Goelzer

mgoelzer@docker.com / @mgoelzer