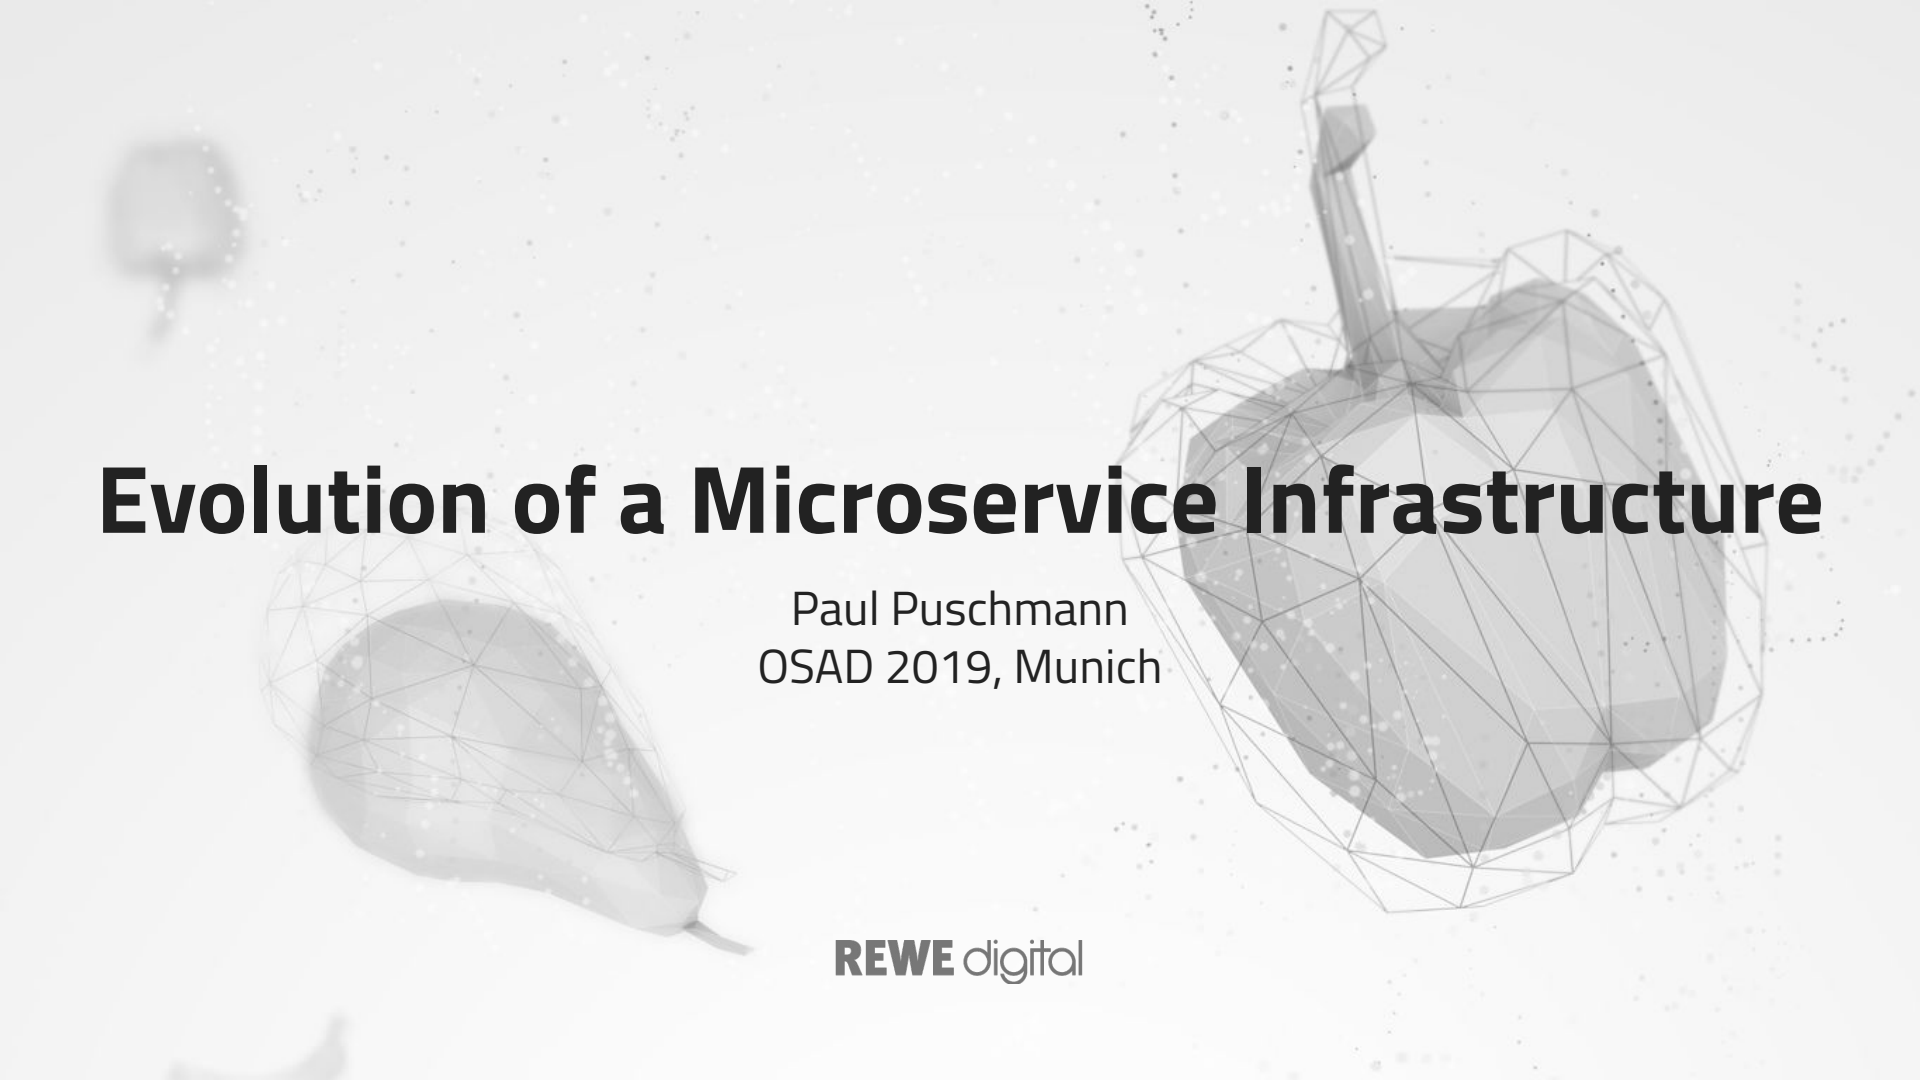

Evolution of a Microservice Infrastructure

Paul Puschmann

OSAD 2019, Munich

REWE digital

What do we actually run?

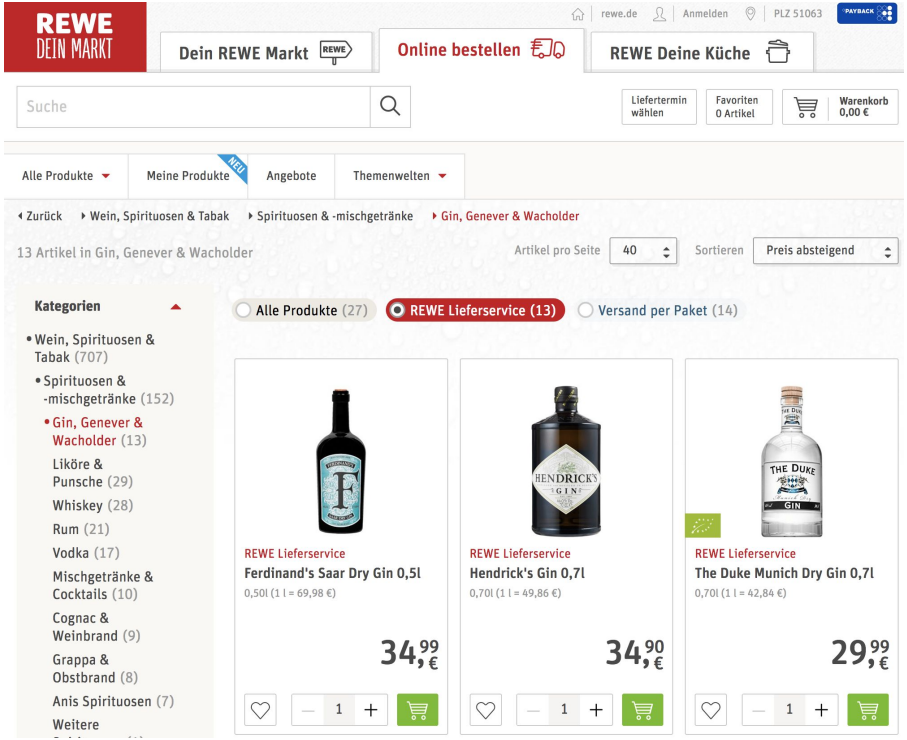

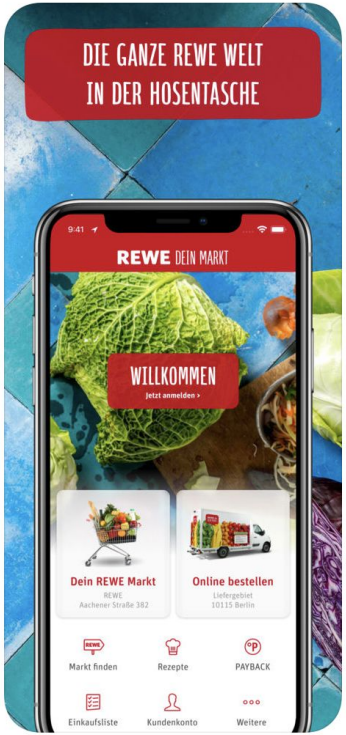

REWE digital

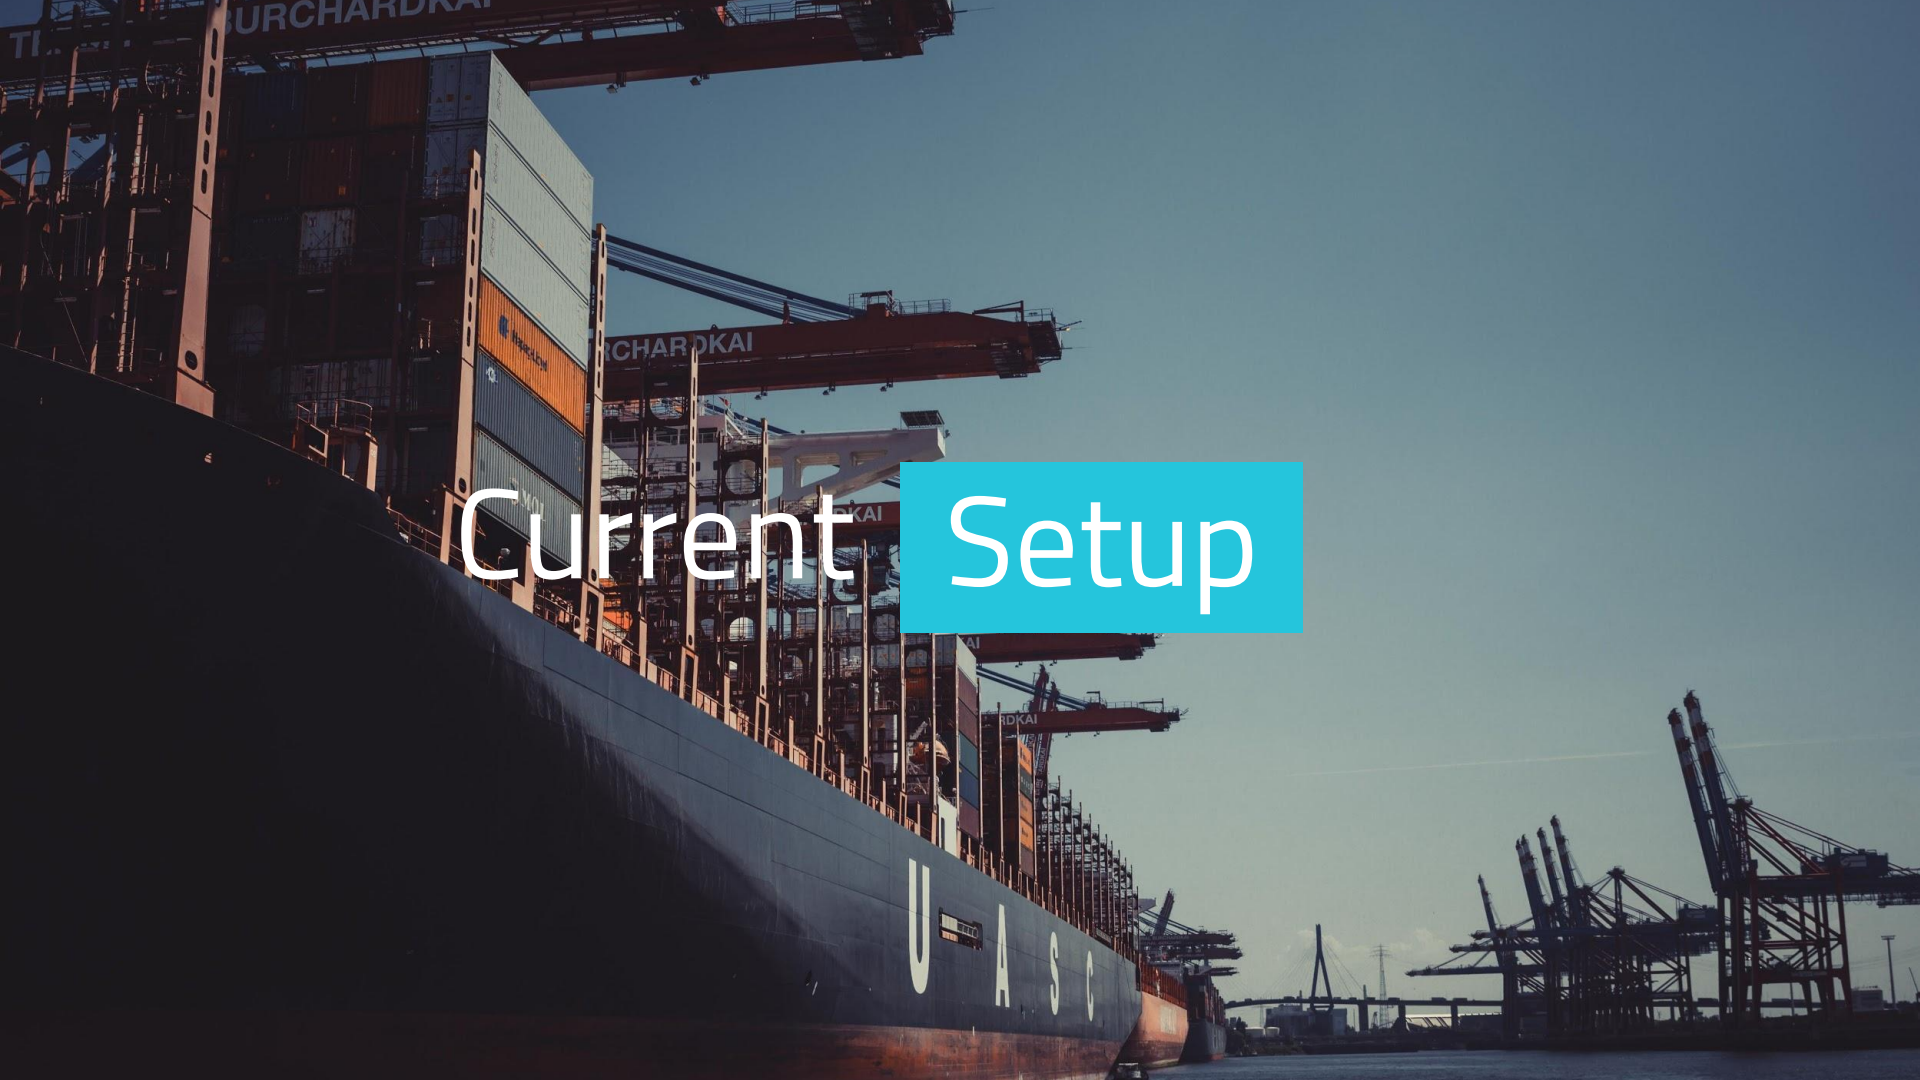
Current Setup

Recap

The state of 2018

We’re operating a custom Docker-Environment consisting of:

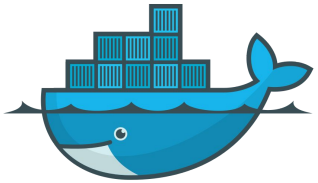

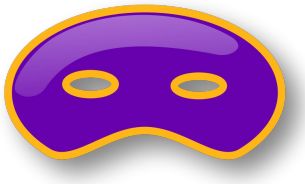

Recap

The state of 2018

We’re operating a custom Docker-Environment consisting of:

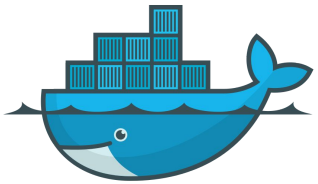

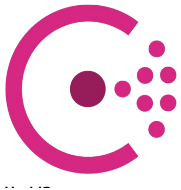

Everything was cool. Developers can bring Code live. All is well.

... and looks like this

for infrastructure provisioning

One repository for infrastructure-configuration

Ansible, Vagrant, Terraform, ... executed via Jenkins.

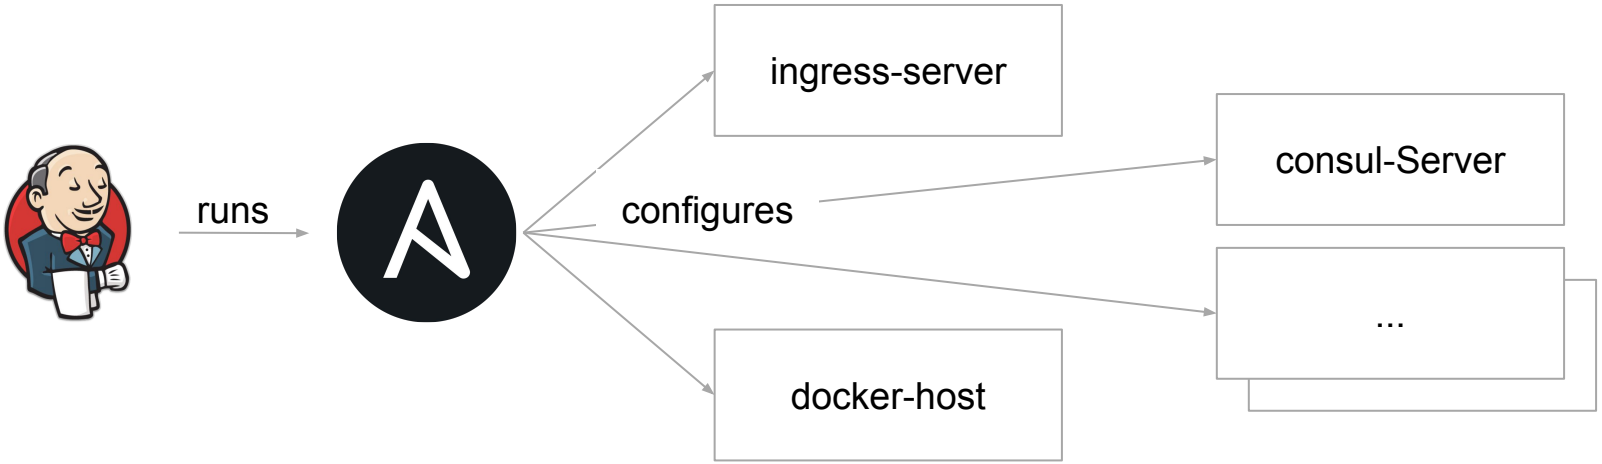

... and works like this

for deployment of services

One central repository for service-deployments

- Used on every Team-Jenkins as external resource
- Teams provide a “service-descriptor.yaml” for each service
- “service-descriptor.yaml” gets updated with environment-specific variables
- containers get started with environment of “service-descriptor.yaml”
- standardised deployment is ensured

Example

```
---
service_name: "example"
service_version: "1.2.1"
squad: "Example-Squad"
team: "Example-Team"
num_instances: 3
prometheus_enabled: "true"
prometheus_path: "/metrics/prometheus"
service_memory: 1536
service_configuration:
  JAVA_META_SIZE_TO_HEAP_QUOTA: 40
  # Example DB
  DATASOURCES_SHOP_JDBCURL: "jdbc:postgresql://{{ psql_cluster_master }}:5432/{{ db_name_example }}"
  DATASOURCES_SHOP_USERNAME: "{{ db_user_example }}"
  DATASOURCES_SHOP_PASSWORD: "{{ db_password_example }}"
  ...
```

Recap

The state of 2018

We’re operating a custom Docker-Environment consisting of:

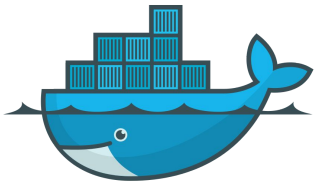

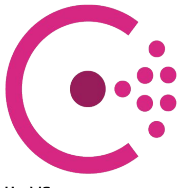

Everything was cool. Developers can bring Code live. All is well.


All is fine
kobo
Photo by Perfecto Capucine from Pexels

... and looks like

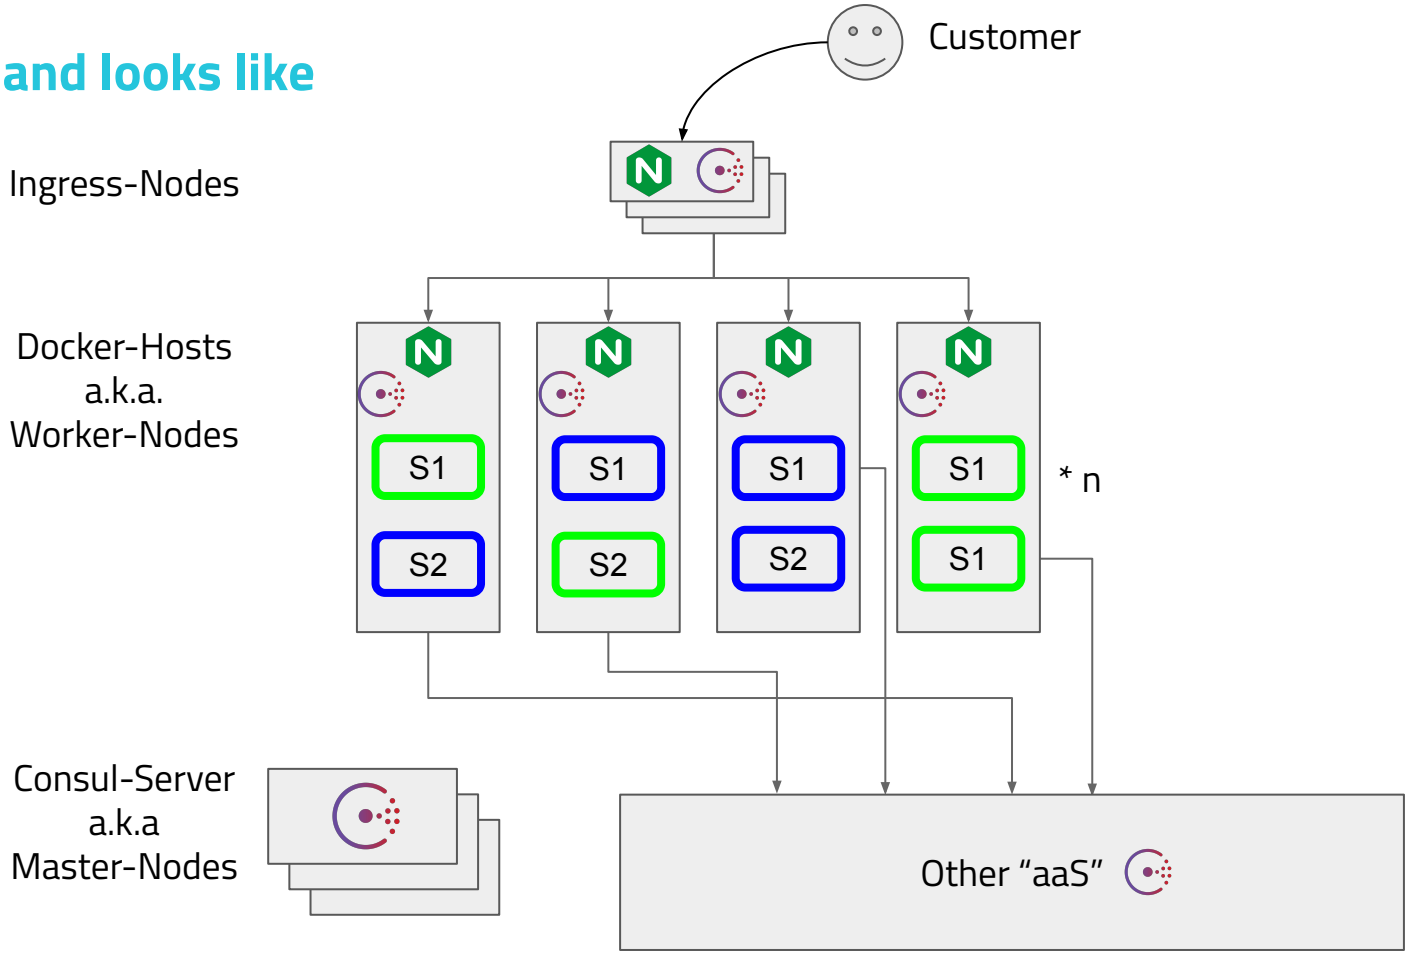

... and looks like

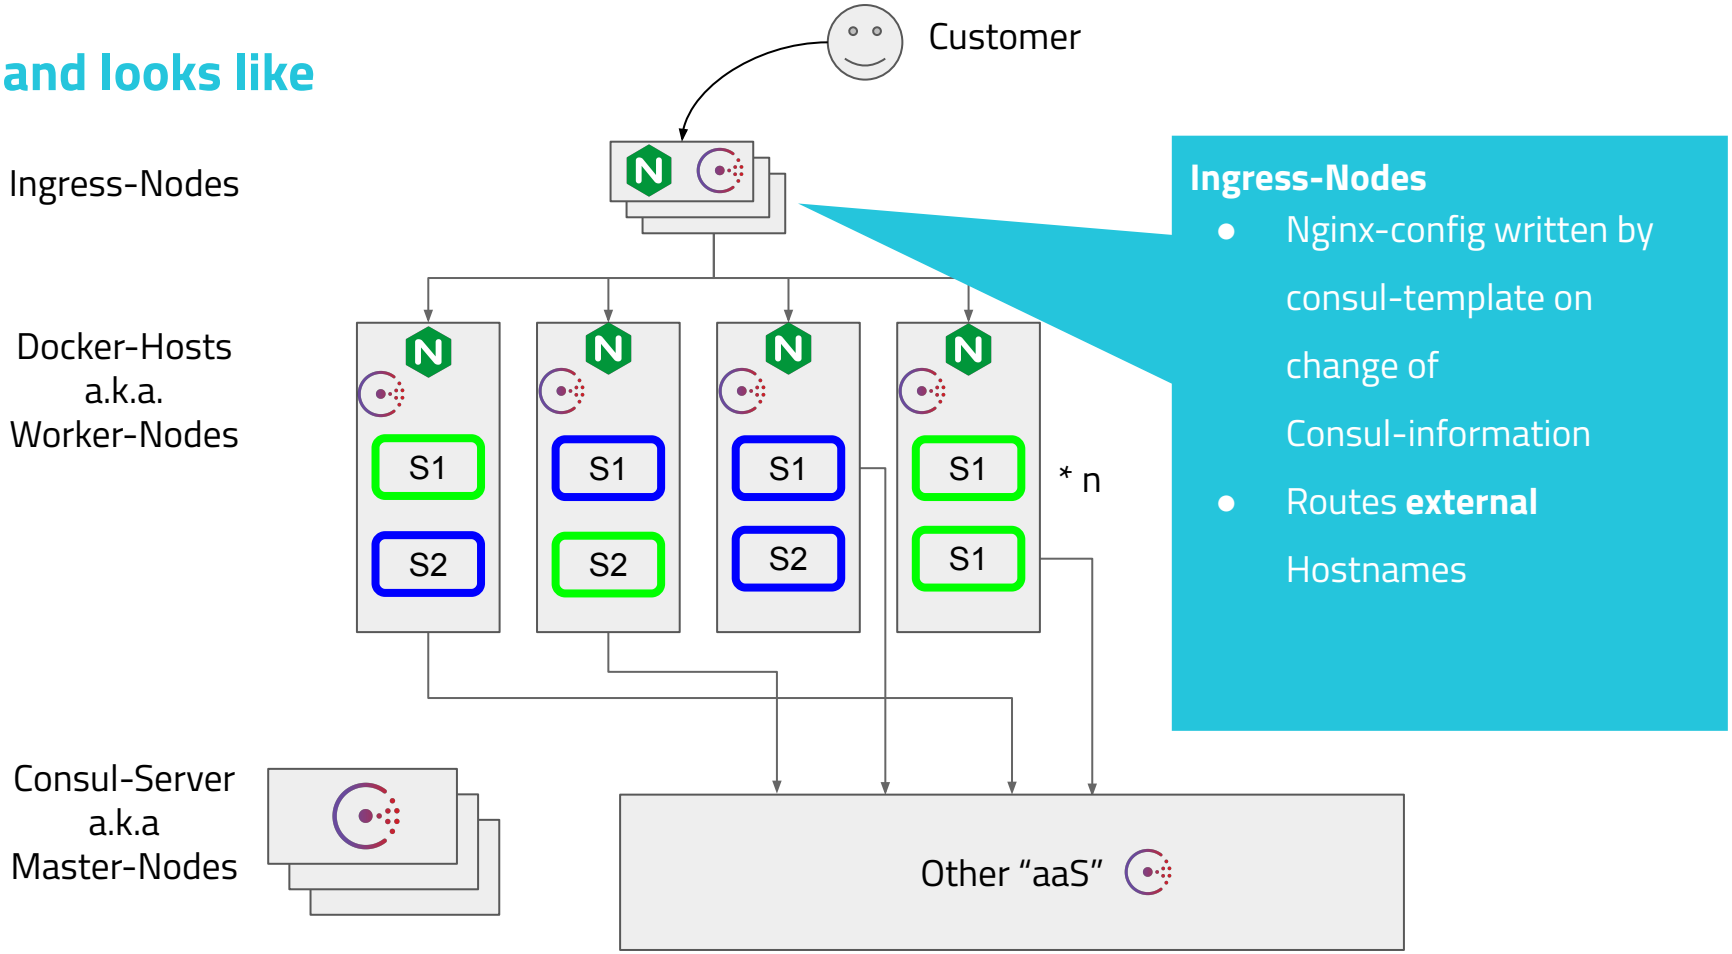

... and looks like

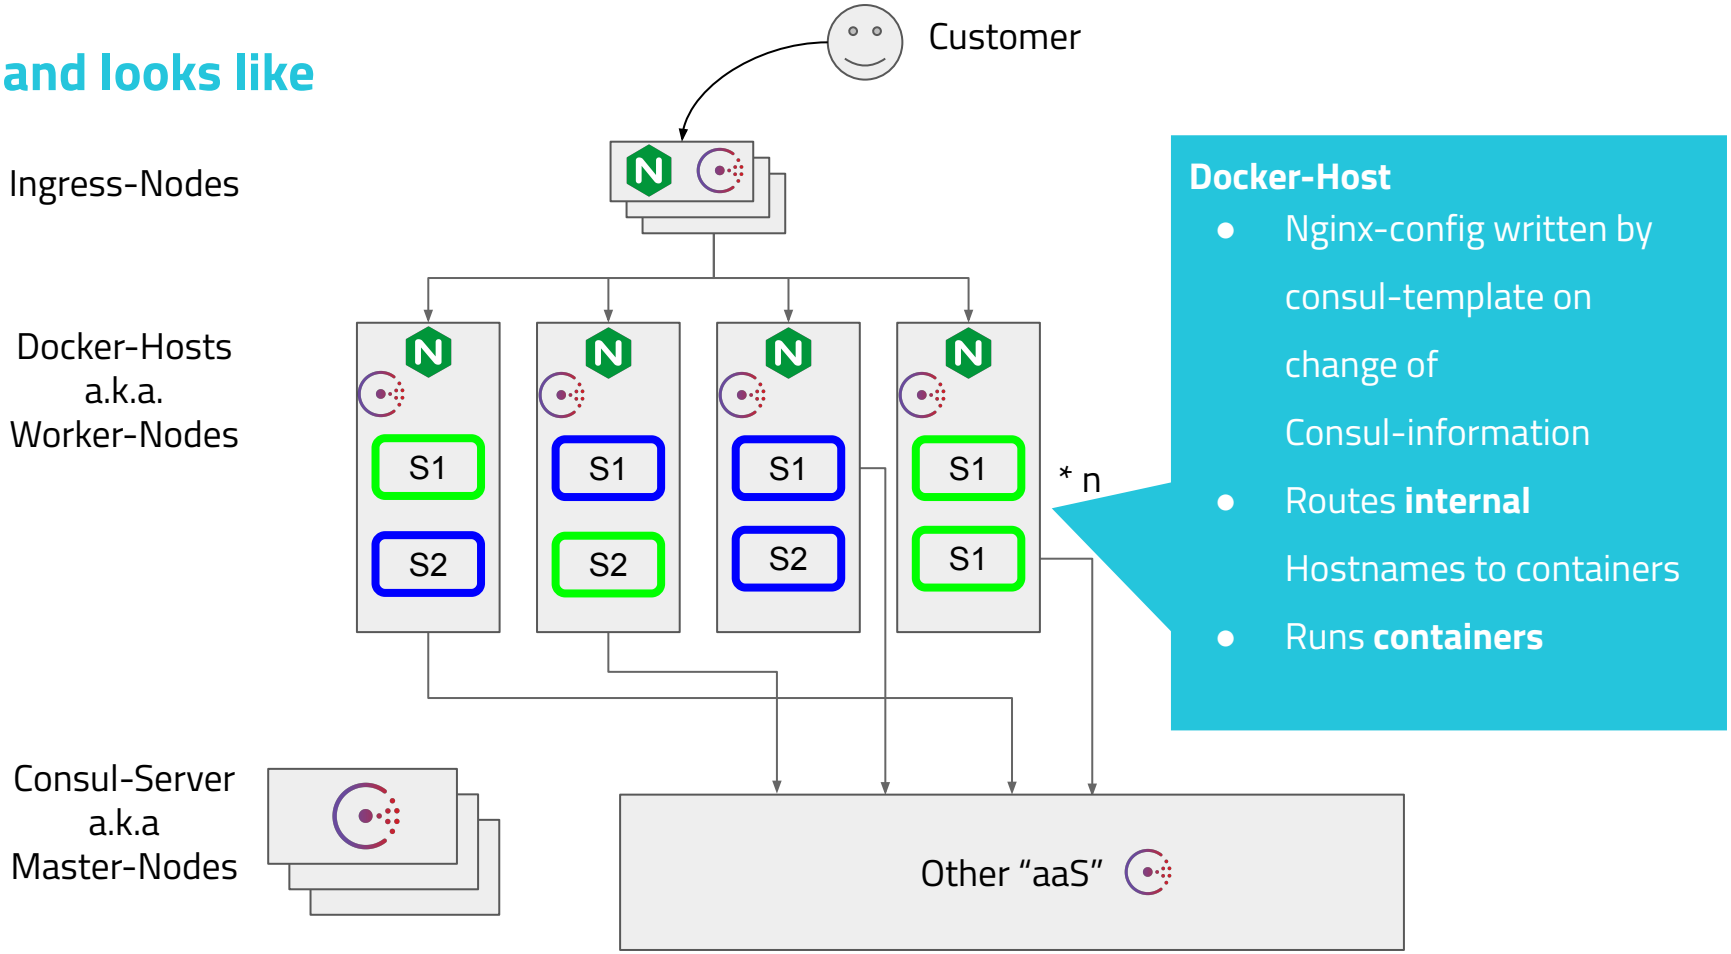

... and looks like

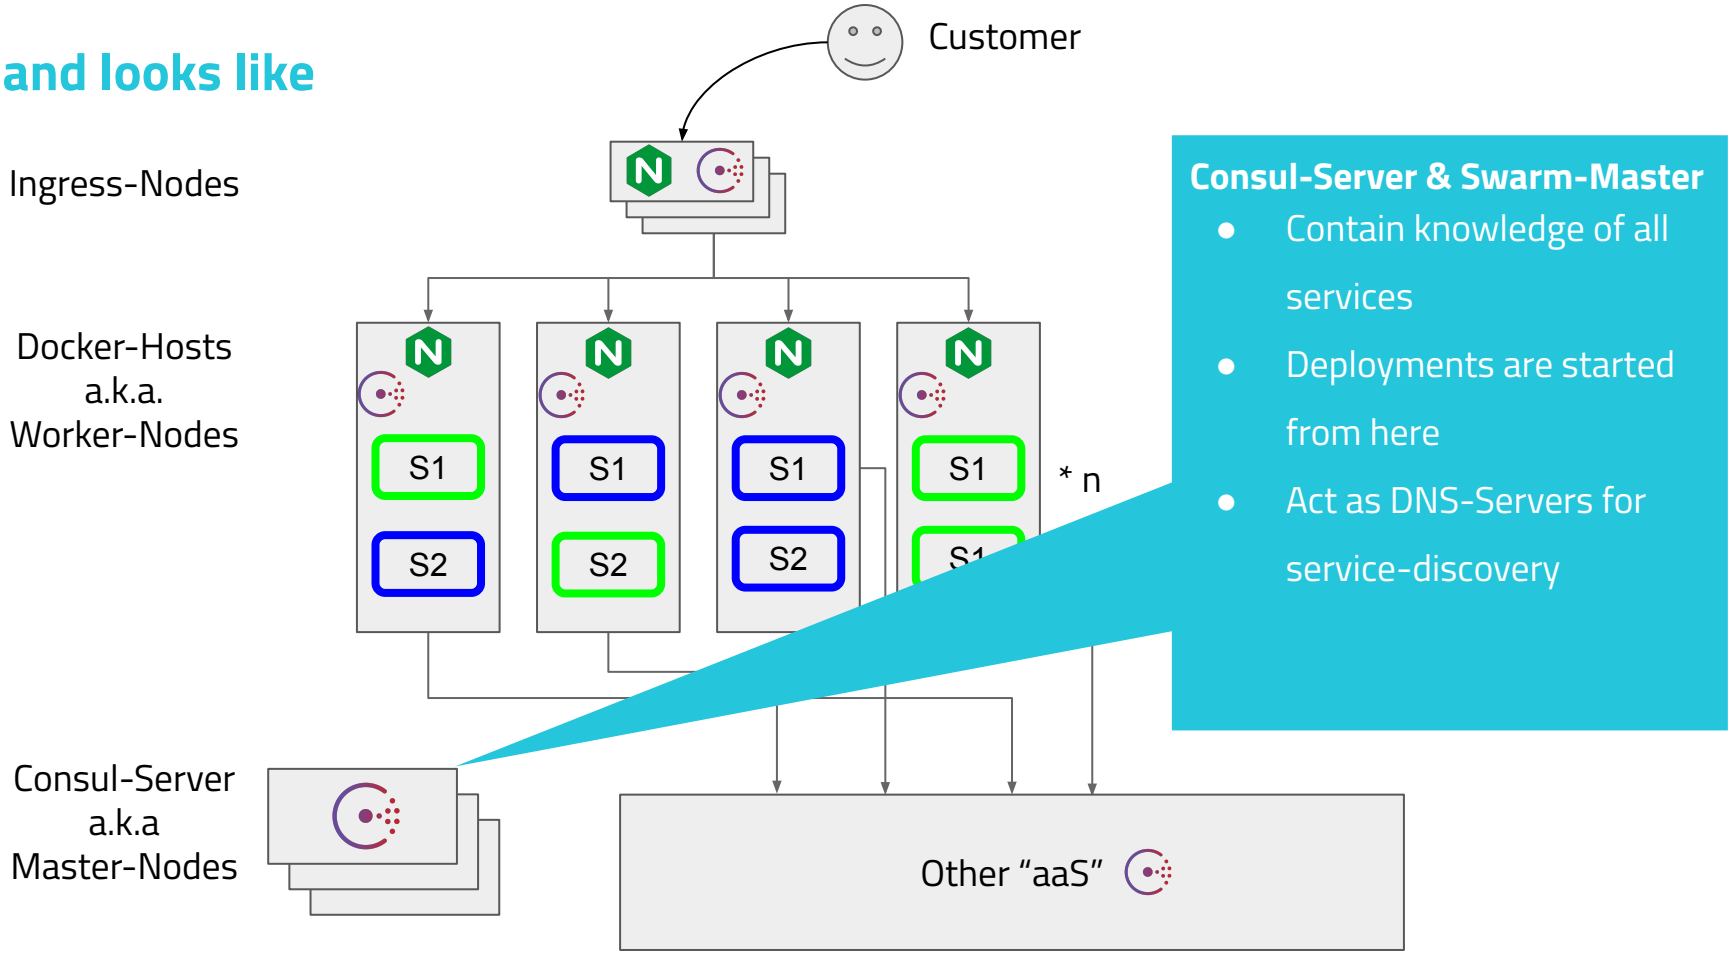

... and looks like

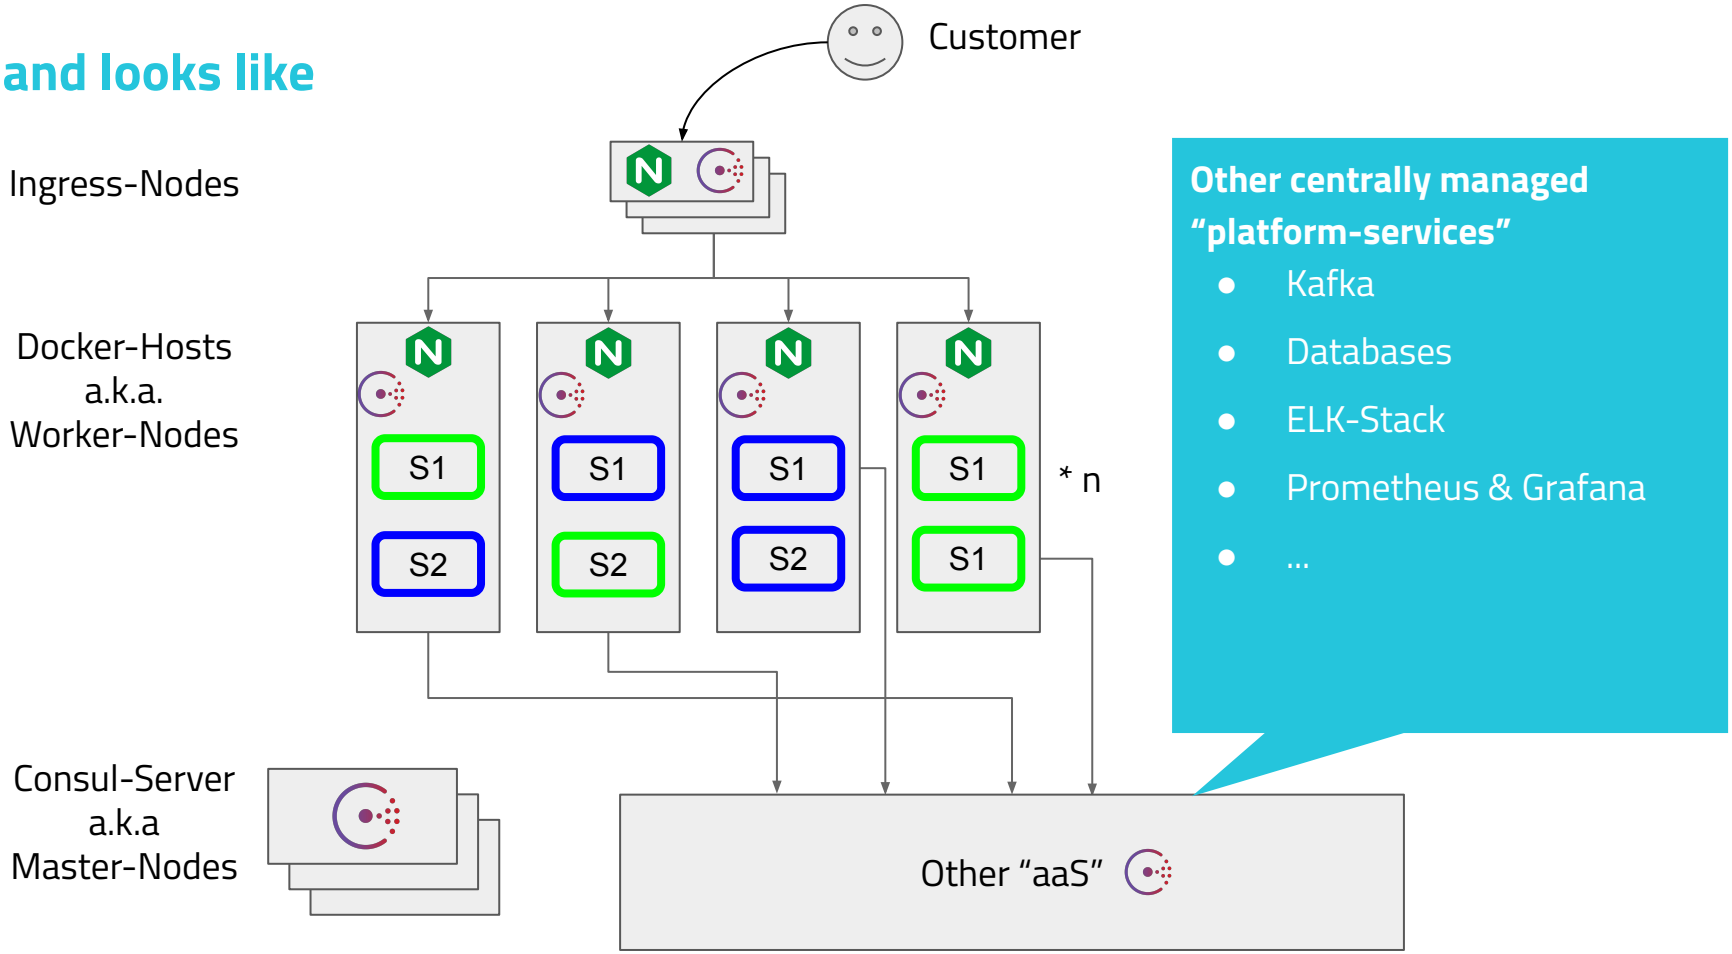

Request routing

how can services be addressed

- Both colors have the same DNS record
    - Consul will return IPs for all hosts where the Service is running
- Nginx running on each Worker Node
    - routes to colour depending on used port


Routing
Problems

Problems with Nginx

increased with the size of the environment

- There are requests which never reached their destination

Always happened at the time of deployments

... and looks like

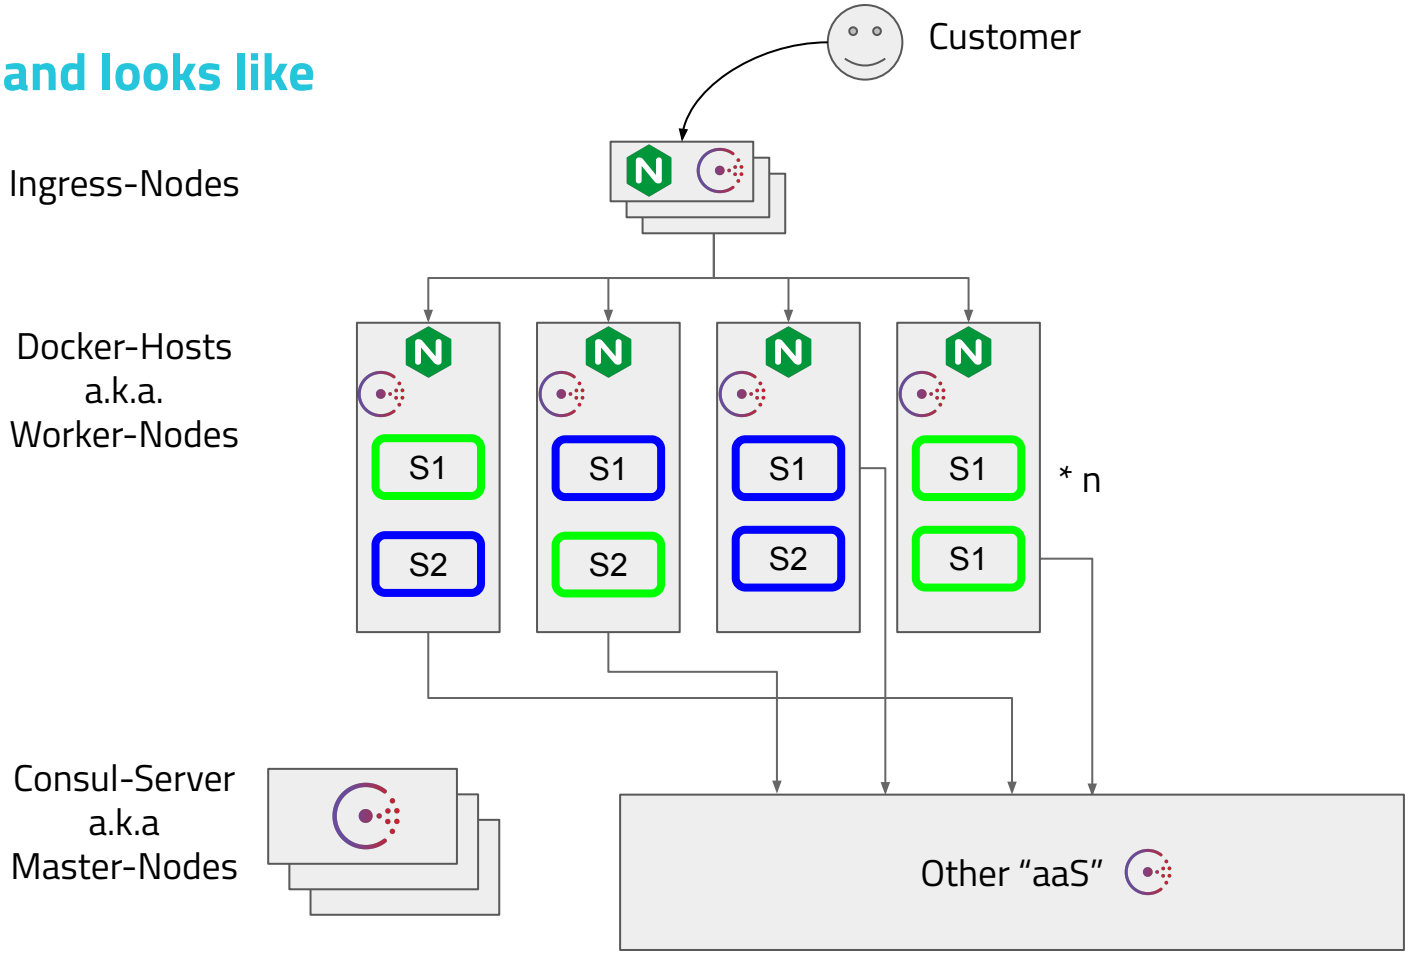

Problems with Nginx

increased with the size of the environment

- There are requests which never reached their destination
- Always happened at the time of deployments
- Consul-template would **reload all** Nginx instances

  at the same time
- What happens at a reload?

Problems with Nginx

looking for solutions

Look for different reverse proxy

- No reload on config change (optional)
- Dynamic configuration (optional)
- Robust connections to the client

Problems with Nginx

possible replacements

Traefik

- Dynamically configurable
- Live reloading of configuration
- Lots of metrics
- Nice web ui
- Single Go binary

Since Traefik 2.x:

- independent configuration of frontend & backend
    - mix consul service-discovery with file-based configuration

Traefik

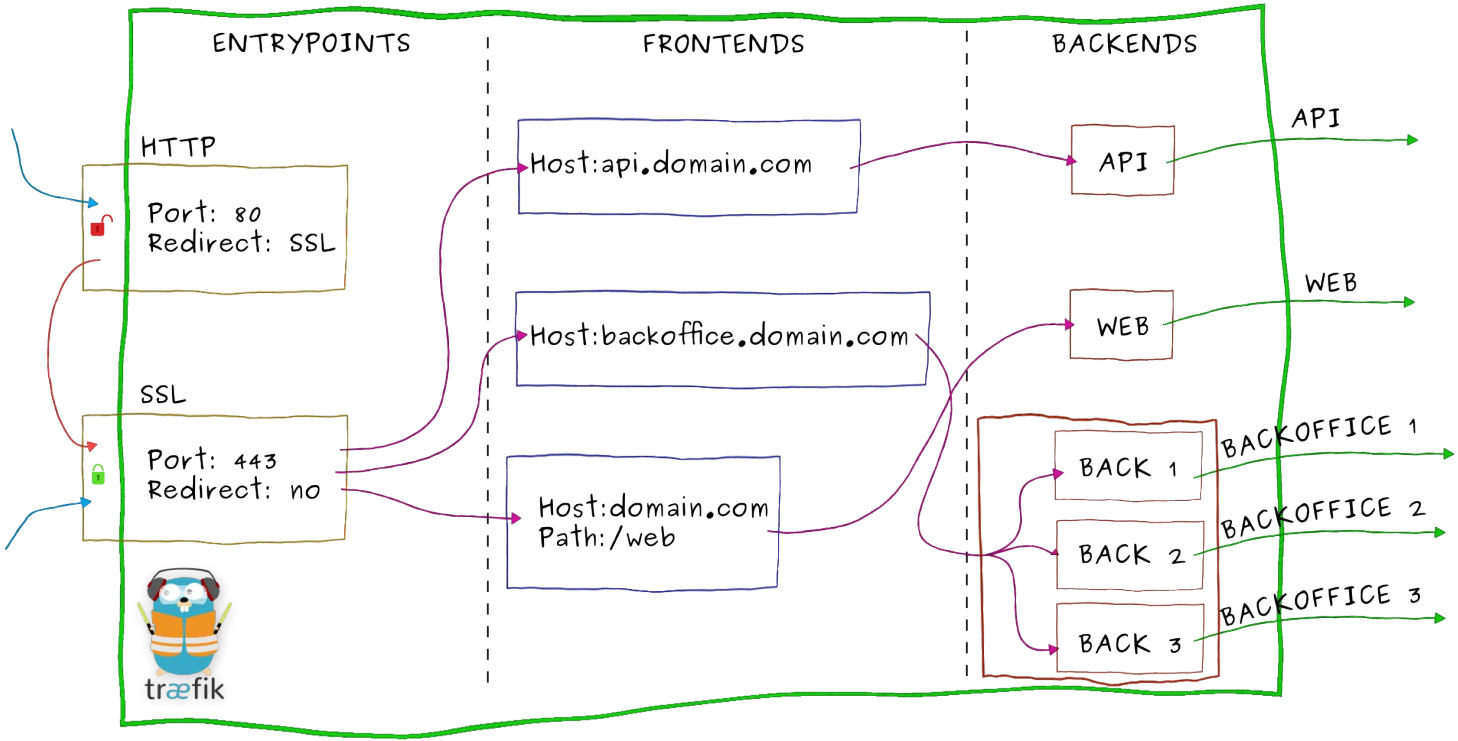
ENTRYPOINTS
FRONTENDS
BACKENDS
HTTP
Port: 80
Redirect: SSL
SSL
Port: 443
Redirect: no
Host:api.domain.com
Host:backoffice.domain.com
Host:domain.com
Path:/web
API
WEB
BACK 1
BACK 2
BACK 3
API
WEB
BACKOFFICE 1
BACKOFFICE 2
BACKOFFICE 3
træfik

Traefik

how to migrate

1. Install alongside Nginx on Worker and Ingress Nodes
    - listen on different ports
2. Check that configured routes are correct and work
3. Change port mapping host by host -> Traefik is active
4. Remove Nginx

Traefik

how to migrate

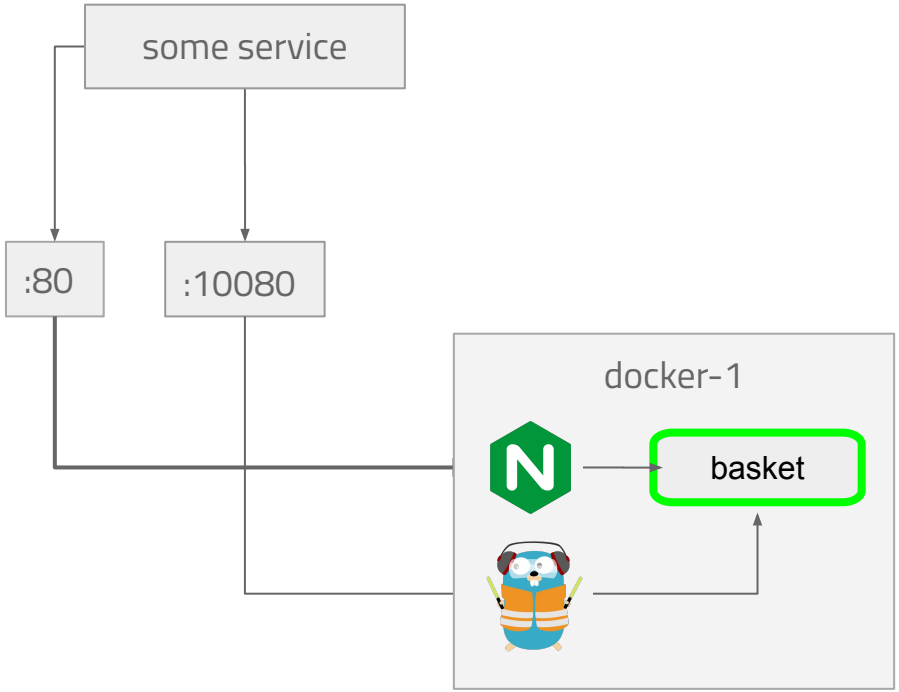

Traefik

how to migrate

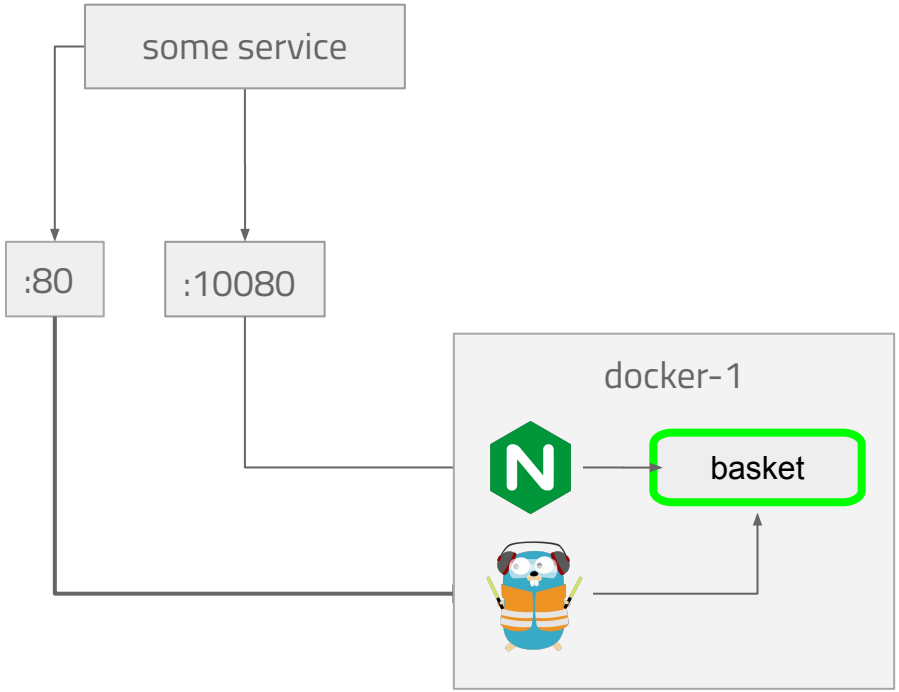

Traefik

how to migrate

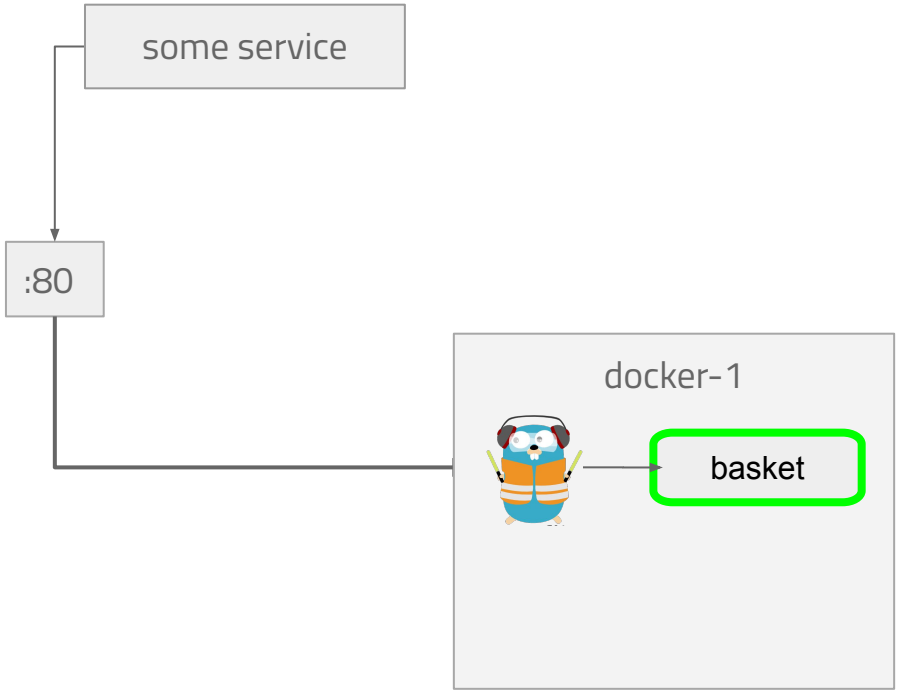

Traefik

Benefits

- Keepalive and connection problems immediately went away
- Almost real time data about service response time
- Web UI to check routes
- Rich access logs

Traefik

Benefits

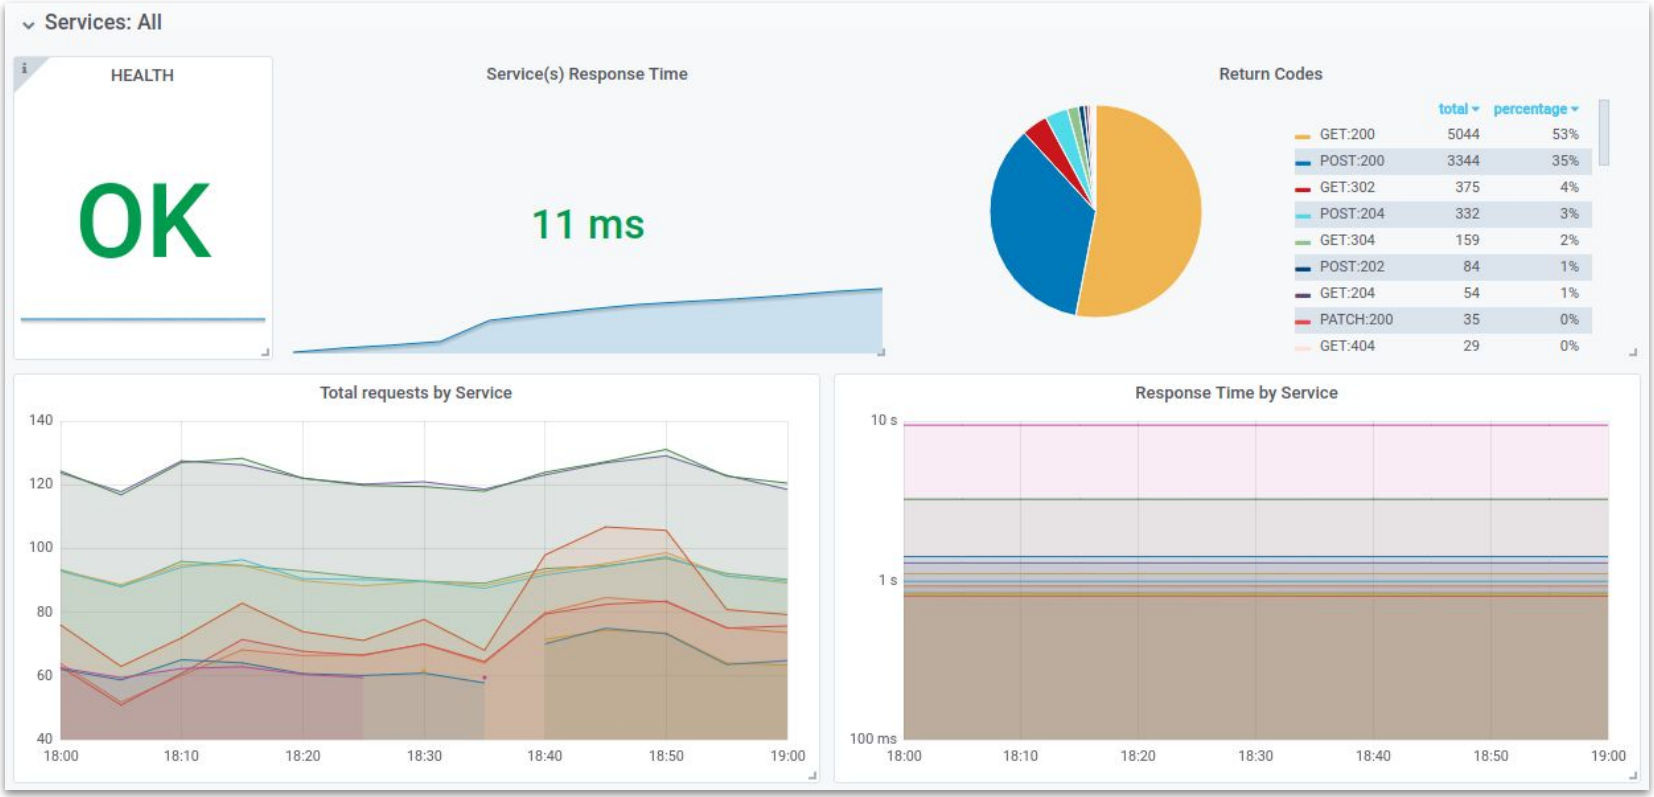

Traefik

Benefits

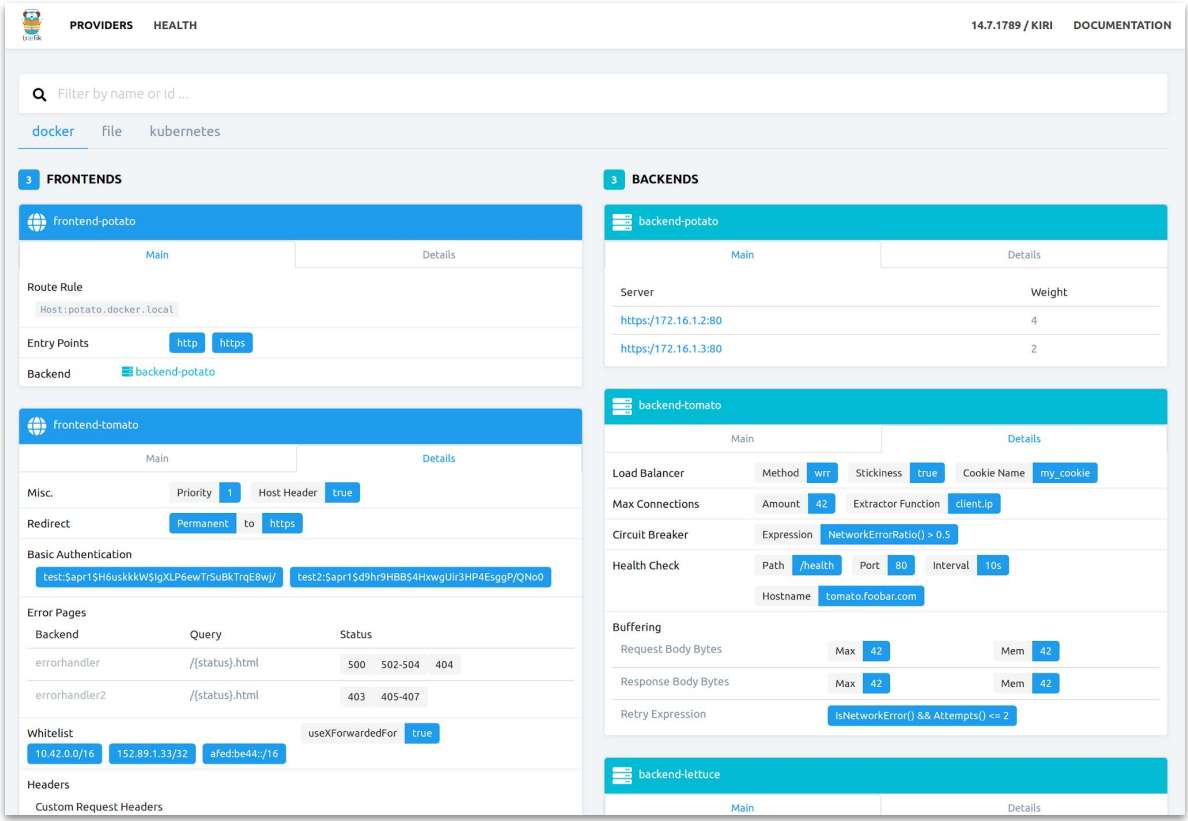

REWE digital


Container Problems

Problems with standalone Swarm

also increased increased with increasing workload

- Poor container spread
  - all service instances running on one host
- No self healing
- Manual node draining (e.g. for maintenance)
  - we’re still dependent on docker-compose files
- Only few metrics

Swarm...

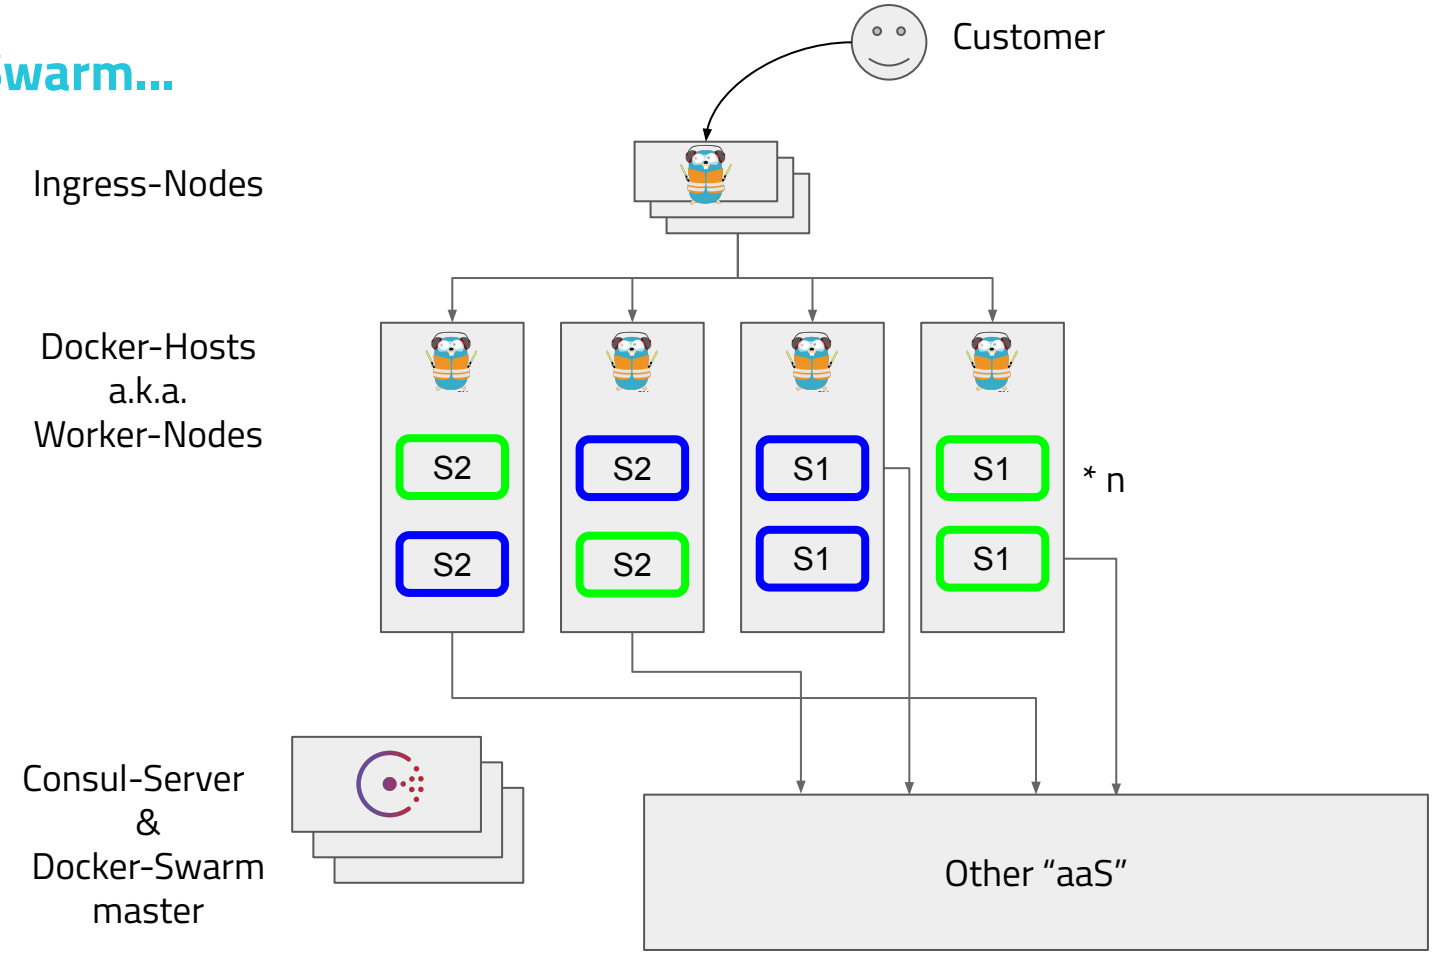
Customer
Ingress-Nodes
Docker-Hosts
a.k.a.
Worker-Nodes
S2
S2
S2
S2
S1
S1
S1
S1
* n
Consul-Server
&
Docker-Swarm
master
Other “aaS”

We want this

- self healing
- proper container spread
- metrics
- resource limits (optional)
- stateless docker-host

Possible replacements

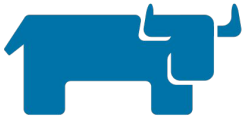

RANCHER

HashiCorp
Nomad

kubernetes

Nomad

- Seamless Consul integration
    - almost no setup needed
- Self healing
- Bin packing
- Single Go binary
- Nice Web UI
- (Memory) Limits enforced by default
- Token-based ACL

Nomad

Benefits

- Not limited to Docker
  - Rkt and LXC
- Not limited to Containers
  - Jar files
  - Binaries
  - VMs

Swarm...

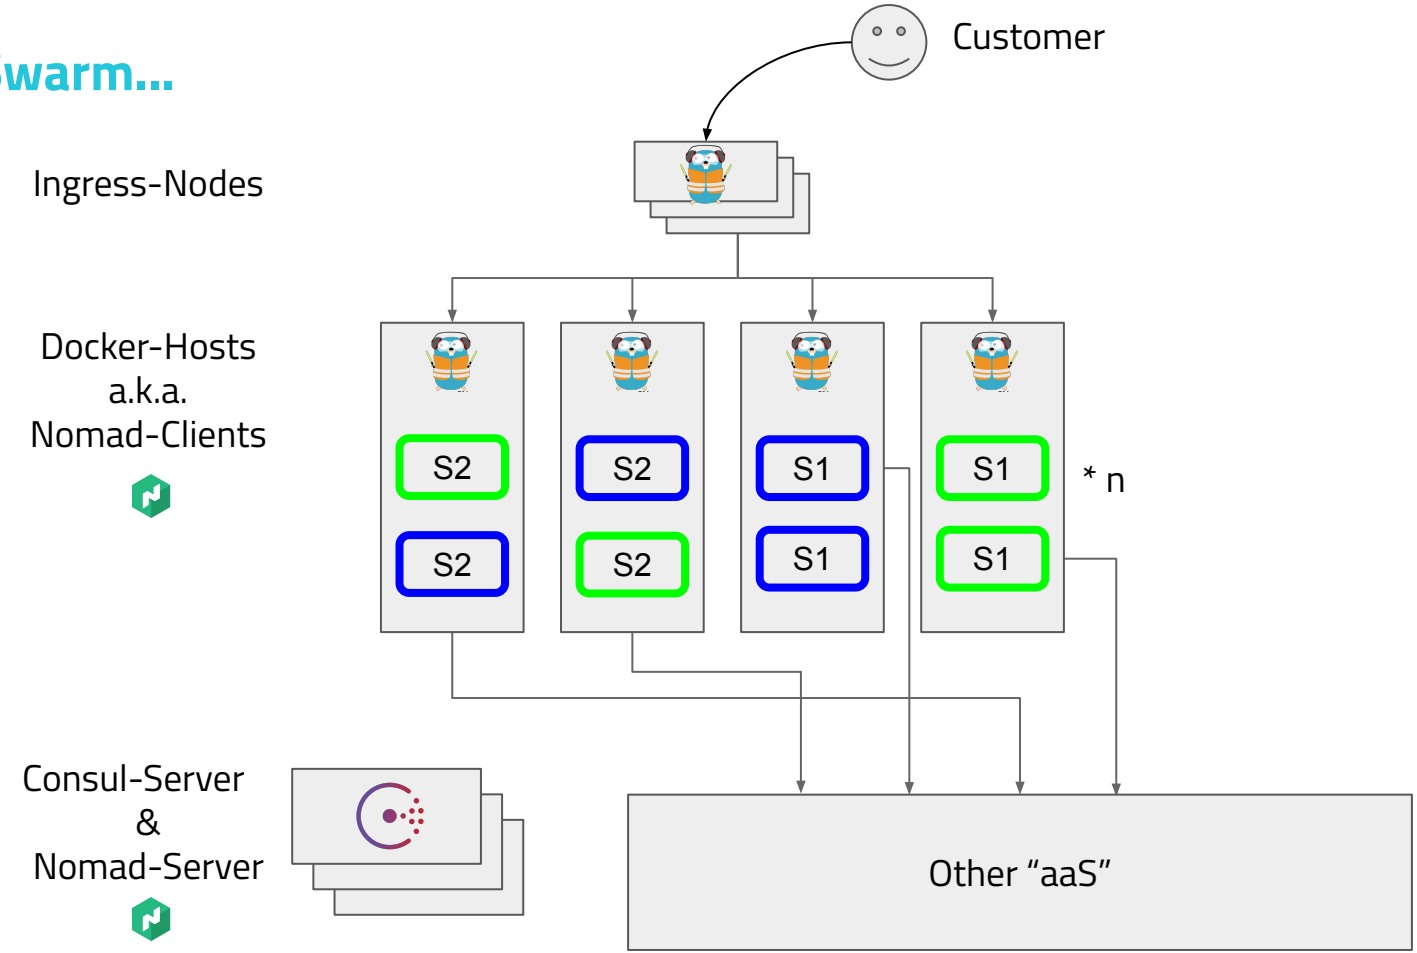
Customer
Ingress-Nodes
Docker-Hosts
a.k.a.
Nomad-Clients
S2
S2
S2
S2
S1
S1
S1
S1
* n
Consul-Server
&
Nomad-Server
Other “aaS”

Nomad

Benefits

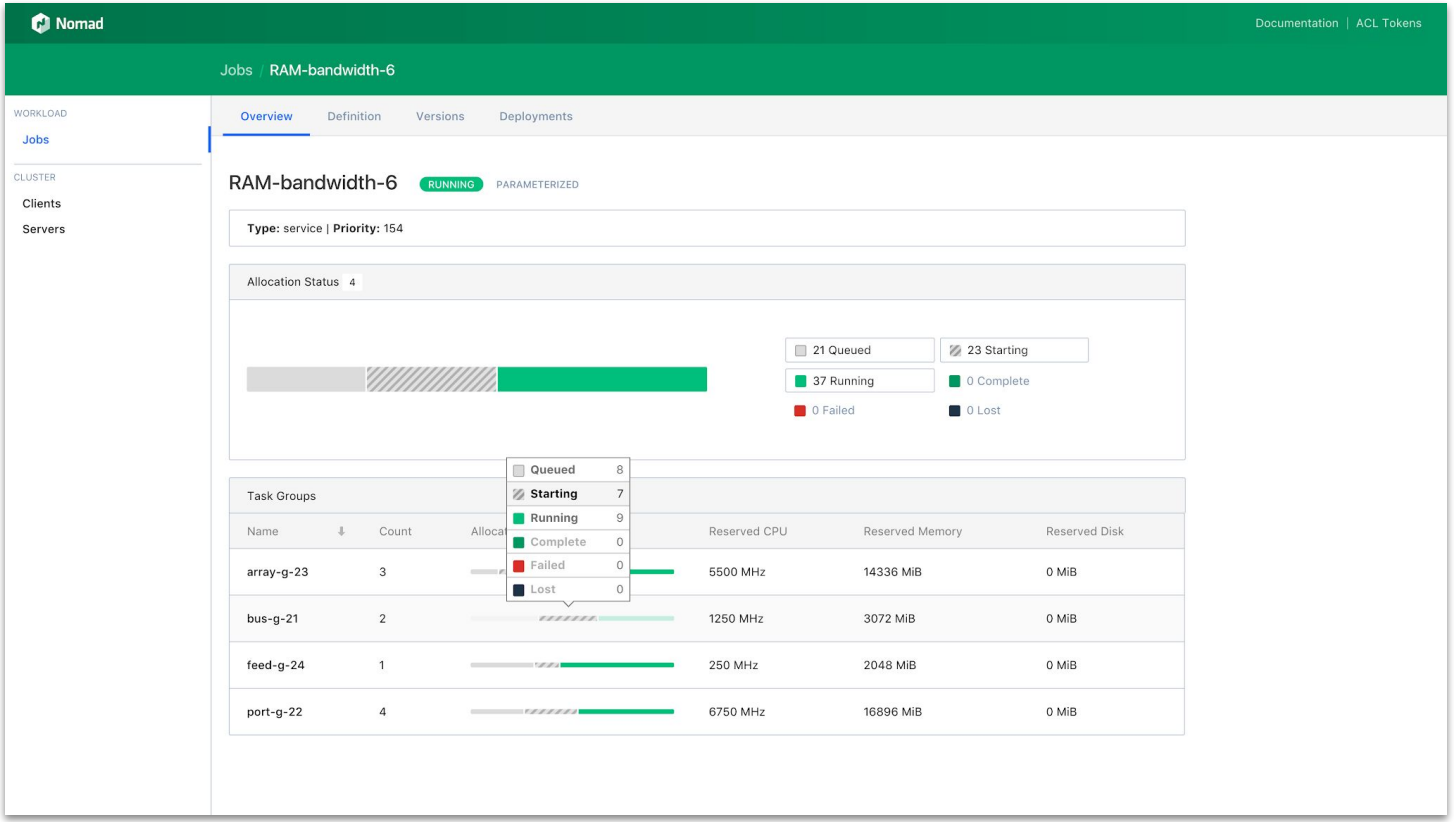

REWE digital

Nomad

Benefits - Cluster Level

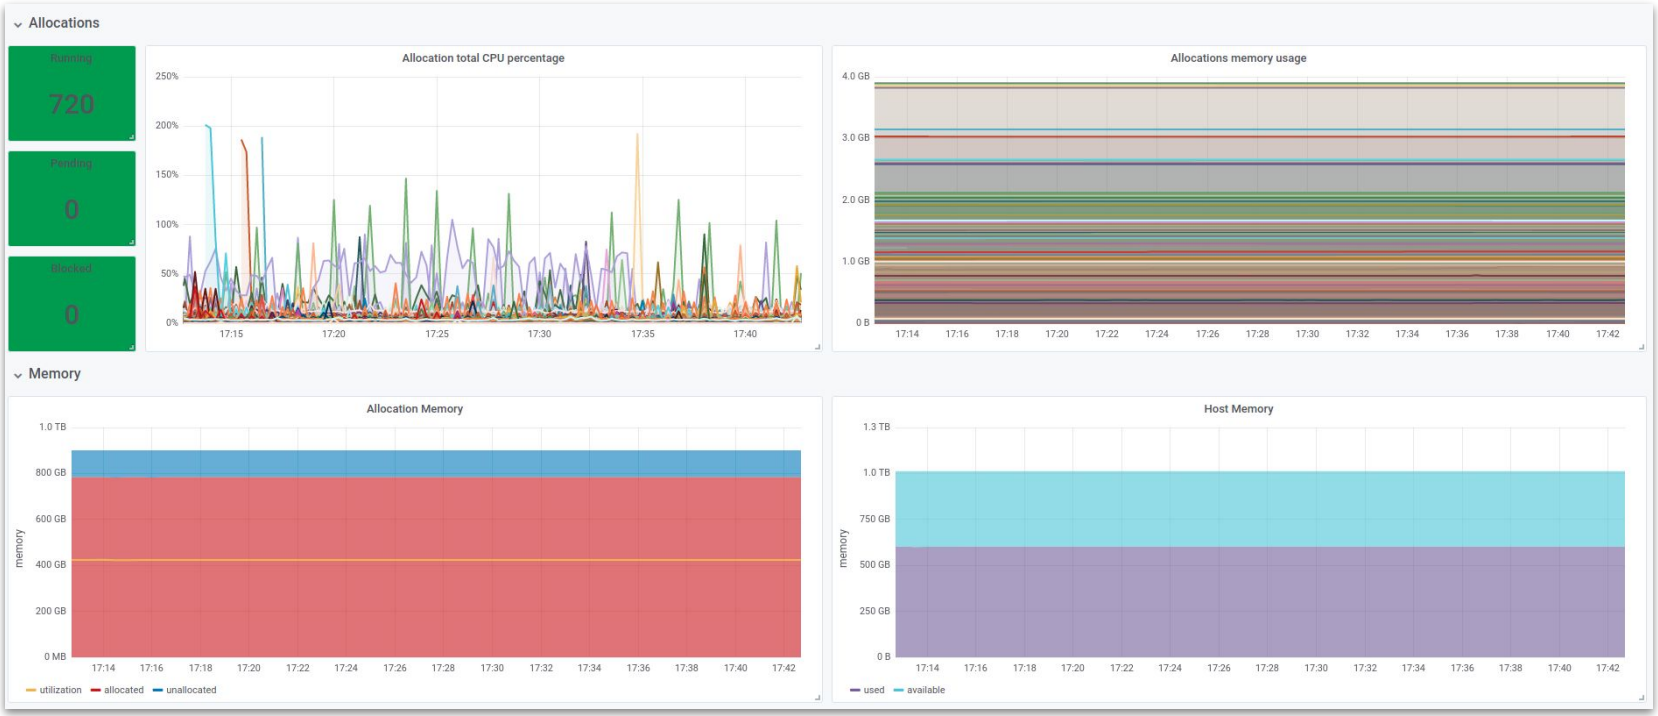

Nomad

Benefits - Cluster Level

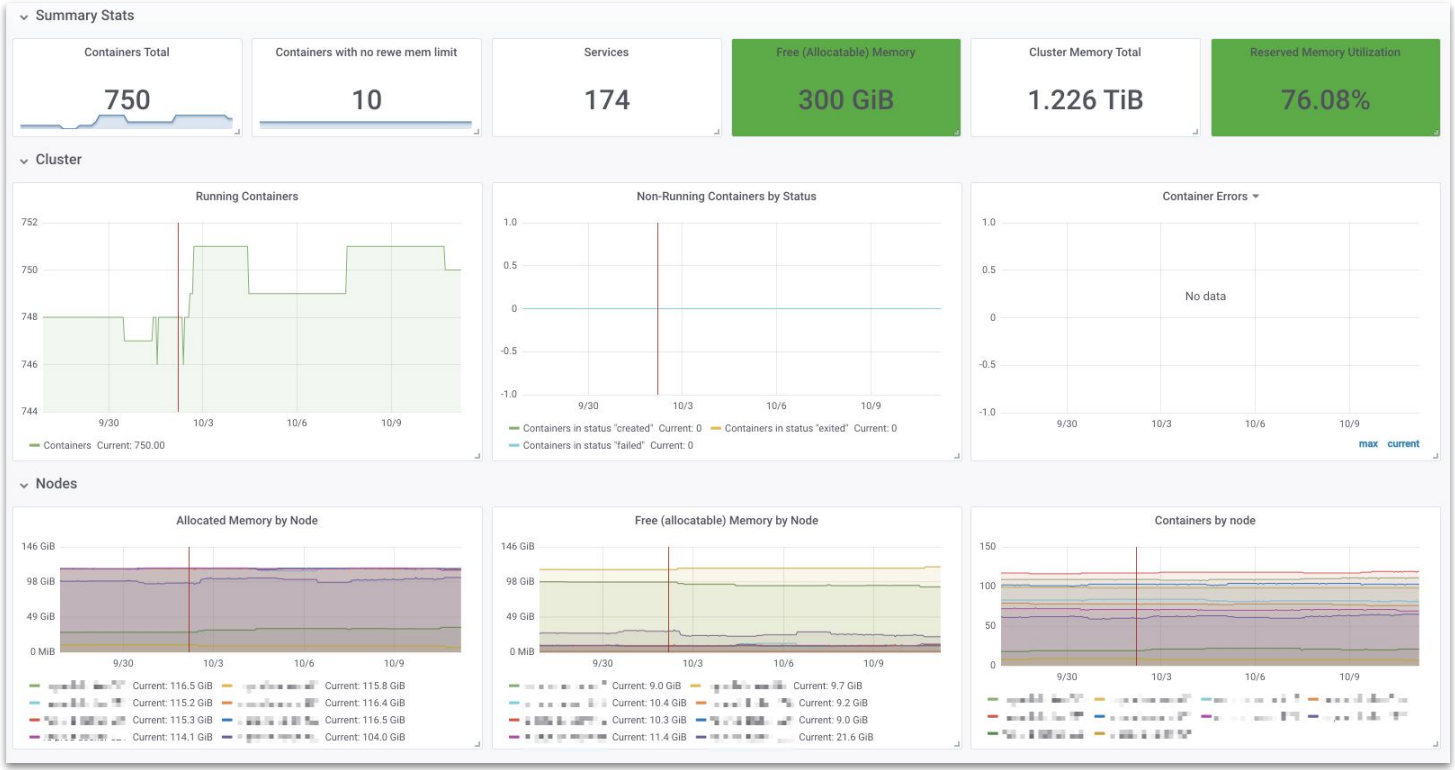

Nomad

Benefits - Service Level

REWE digital

Nomad

Benefits

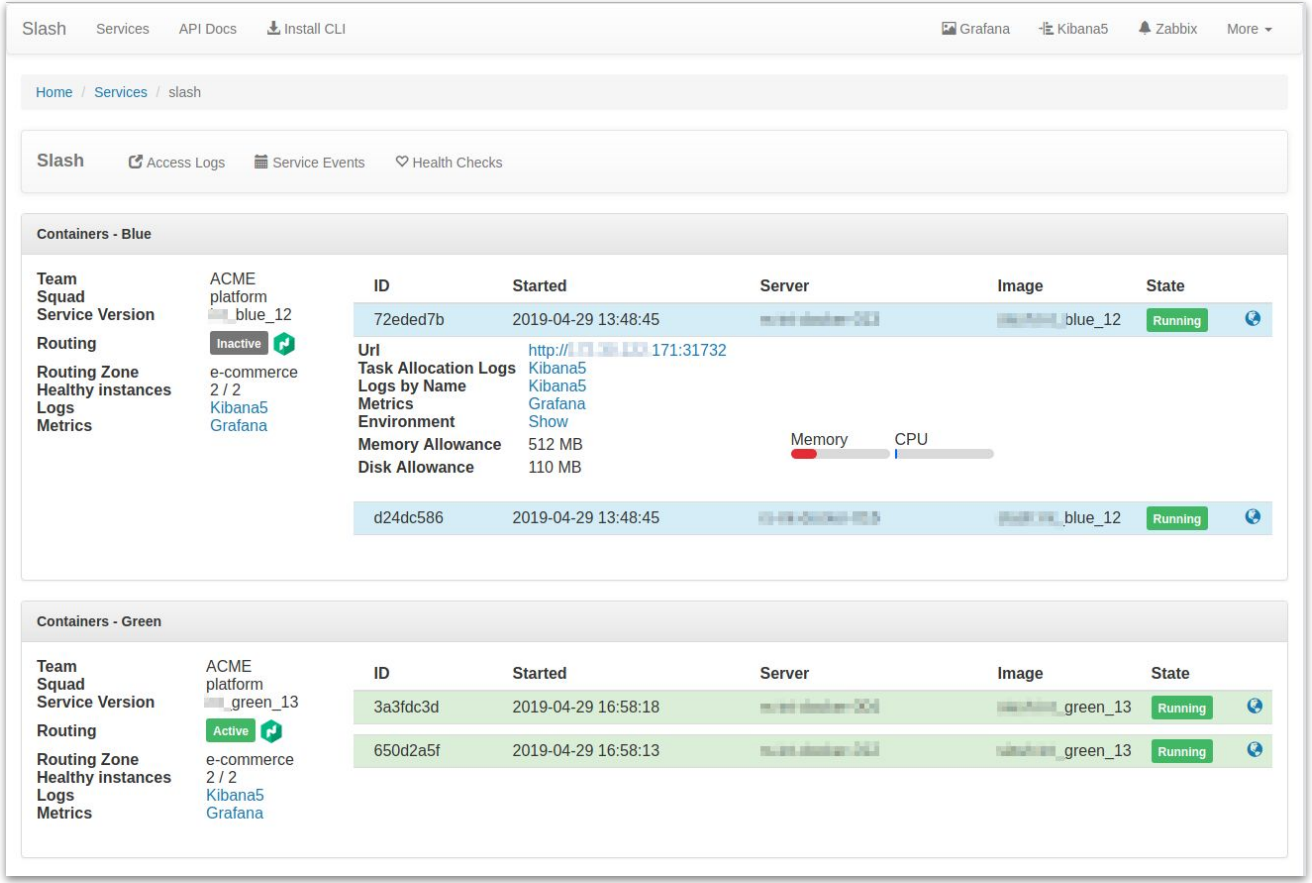

State of 2019

We’re operating a custom Docker-Environment consisting of:

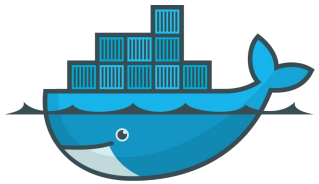

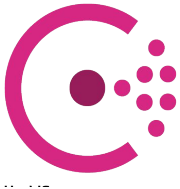



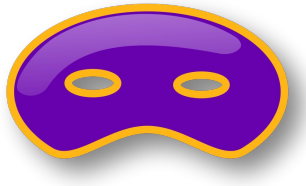


State of 2019

And we're also using

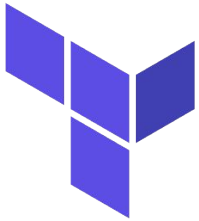

HashiCorp Terraform


What we Learned

What helped us most?

- Having a centralised deployment-toolset
  - perform all changes for all teams / developers at the same time
- Do Canary-like changes on our infrastructure
  - fully interoperable changes
  - Nginx <-> Traefik

What did we learn?

- Distributed systems can be hard
- Keeping your architecture pluggable helps a lot
- Computing resources can be finite
    - Enforcing resource limits can be interesting
- You might not need Kubernetes...

Evolution of a Microservice Infrastructure

OSAD 2019, Munich

Thank You!

Paul Puschmann @ppuschmann

www.rewe-digital.com @rewedigitaltech

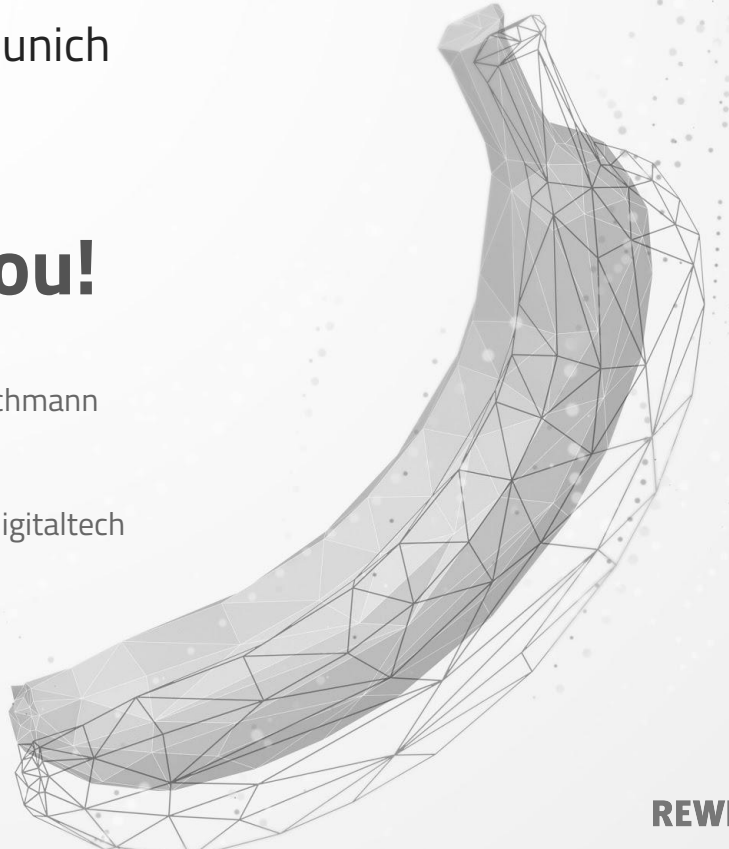

REWE digital