



**OST**

Eastern Switzerland  
University of Applied Sciences

# Distributed Systems (DSy)

## Deployment

Thomas Bocek

17.04.2023

# Learning Goals


- Lecture 8 (Deployment)
  - Different ways to deploy your service
    - High-level overview
  - Cloud Infrastructure [[link](#)], Cloud Operations [[link](#)] - Laurent Metzger
  - Cloud Solutions [[link](#)] - Mirko Stocker

# Back in the old days...

- **OTS**: apt-get / yum / pacman install package, e.g., Apache – configure – run
- Custom SW: Java: **war**, provide custom /etc/init.d script with binary or script
- Problem:
  - It runs on my machine, who installs Java in the right version?
  - What happens on crashes?
  - Scaling?
  - HW defect?
  - Misconfiguration - access to complete PC?
- VMs / Containers help a lot
  - No access to complete PC, can scale, move to another machine, pre-install the right Java version
- So, how to deploy your app?
  - **Ansible** (**Progress Chef**, **Puppet**) - and **more**
    - Playbooks with ssh host list – your host needs to run the same OS (apt/yum)
  - Docker Swarm
    - Works with docker-compose.yml – with docker you package your application the same way on any platform - **simple**
      - Which to use? [[link](#)]
  - Kubernetes
    - Widespread

# Docker Swarm

- Use `docker --context` to run/maintain containers on other machines
- Docker Swarm
  - Deploy with `docker-compose.yml` ([deploy](#))
  - Built into docker
    - `docker swarm` – manage swarm
    - `docker node` – manage nodes
  - Scheduler is responsible for placement of containers to nodes
  - Can use the same files, [easy to setup?](#)
    - [Azure](#), [Google cloud](#), [Amazon](#)




docker swarm

- [Kubernetes vs. Docker Swarm](#)
- “Docker Swarm has already lost the battle against Kubernetes for supremacy in the container orchestration space” [[link](#)]
- “Kubernetes supports higher demands with more complexity while Docker Swarm offers a simple solution that is quick to get started with.” [[link](#)]

# Docker Swarm

- 3 “Machines”
  - KVM instances, alpine running
    - Workers: 192.168.1.238, 192.168.1.103, 192.168.1.173
    - Manager: 192.168.1.166
- Run on manager
  - `docker swarm init --advertise-addr 192.168.1.166`
- To add a worker to this swarm, run the following command:
  - `docker swarm join --token .. 192.168.1.166`

- `docker info`
- `docker node ls`
- **Manager are setup**
- Join nodes
  - Run the `docker swarm join` command
  - `docker node ls`
- **Workers are setup**



# Docker Swarm

- Create service
  - docker service create --name registry --publish 5000:5000 registry:2
  - Where to find the docker image
- Check service
  - docker service ls
- Many options in docker-compose
  - docker stack deploy --compose-file docker-compose.yml

```
worker:
  image: gaiadocker/example-voting-app-worker:latest
  networks:
    voteapp:
      aliases:
        - workers
  depends_on:
    - db
    - redis
  # service deployment
  deploy:
    mode: replicated
    replicas: 2
    labels: [APP=VOTING]
  # service resource management
  resources:
    # Hard limit - Docker does not allow to allocate more
    limits:
      cpus: '0.25'
      memory: 512M
    # Soft limit - Docker makes best effort to return to it
    reservations:
      cpus: '0.25'
      memory: 256M
  # service restart policy
  restart_policy:
    condition: on-failure
    delay: 5s
    max_attempts: 3
    window: 120s
  # service update configuration
  update_config:
    parallelism: 1
    delay: 10s
    failure_action: continue
    monitor: 60s
    max_failure_ratio: 0.3
  # placement constraint - in this case on 'worker' nodes only
  placement:
    constraints: [node.role == worker]
```




# Kubernetes

- What is [Kubernetes](#) (K8s)
  - Container orchestration
    - Automates deployment, scaling, and management of containerized applications
  - Started by Google in 2014, now with [CNCF](#)
    - Widely adopted in the industry for managing complex applications
- Kubernetes-based PaaS
  - [Google](#), [Amazon](#), [Azure](#) (book), [Digital Ocean](#), ...
    - [Difficult pricing schemes](#)
- Why Kubernetes?
  - Simplifies application deployment and management
    - Development: run on one machine, deployment how and where to distribute?
  - Ensures high availability and fault tolerance
    - Containers can crash, machine that runs container can crash (e.g., out of memory)
  - Supports auto-scaling based on demand
  - Facilitates rolling updates and rollbacks
    - Rollbacks are hard, especially with state, stateless rollback is easier
  - Provides a powerful ecosystem of tools and services
    - Package manager [Helm](#) released in 2016 ([convert docker-compose](#))

# Kubernetes

- Design principles
  - Configuration is declarative – declare state with YAML/JSON
  - Immutable containers
    - Don't store state in a container. If a health check fails, Kubernetes removes the container and starts a new one
    - Rollback applications, use older version of container – may need to change schema



- Architecture
  - Master Node: Controls the overall state of the cluster
    - API Server: Manages communication within the cluster
    - etcd: Stores configuration data for the cluster
    - Controller Manager: Ensures the desired state of the cluster
    - Scheduler: Assigns workloads to worker nodes
  - Worker Node: Runs application containers
    - kubelet: Communicates with the master node and manages containers
    - kube-proxy: Handles network routing and load balancing
    - Container runtime: Executes containers (Docker, containerd, etc.)



# Kubernetes

- Key Concepts [[link](#)]

- **Pod**: Smallest deployable unit, contains one or more containers
- **Service**: Stable network endpoint to expose a set of Pods
- **Deployment**: Manages the desired state of an application, define scale, HW limits
- **ConfigMap**: Stores non-sensitive configuration data for an application
- **Secret**: Stores sensitive configuration data, like passwords and API keys

- **Volume**: Persistent storage for data generated by a container
- **Namespaces** – run multiple projects on one cluster, separate with namespaces

## ▼ Concepts

- ▶ Overview
- ▶ Cluster Architecture
- ▶ Containers
- ▶ Windows in Kubernetes
- ▶ Workloads
- ▶ Services, Load Balancing, and Networking
- ▶ Storage
- ▶ Configuration
- ▶ Security
- ▶ Policies
- ▶ Scheduling, Preemption and Eviction
- ▶ Cluster Administration
- ▶ Extending Kubernetes

# Kubernetes

- Getting Started with Kubernetes: [Minikube](#), [k3s](#)
  - Minikube: Run a single-node Kubernetes cluster locally
  - kubectl: Command-line tool for managing a Kubernetes cluster
  - Kubernetes Dashboard: Web-based user interface for managing a cluster
- Deploy any containerized application
  - Use health endpoints
    - [Liveness/Readiness](#)
- Official documentation: <https://kubernetes.io/docs>
- Kubernetes tutorials: <https://kubernetes.io/training>
- [Youtube course](#)

