



Kubernetes - Build or Buy

Martin Schmidt – Beck et al.

Christoph Stoettner – panagenda



Cloud



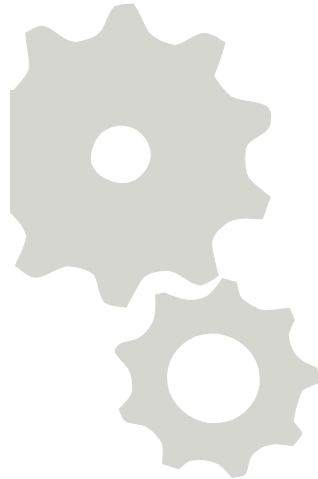
On Premises



Agenda

- Short Intro on containers and orchestration
- Disclaimer (actual support statement)
- Building a Kubernetes cluster
 - Planning
 - Network (Overlay, DNS)*
 - Storage (HA)*
 - Backup / Restore
 - On Premises caveats
 - Using a Cloud provider
- Sizing / Costs

* On Premises only



Cloud



On Premises

Martin Schmidt



+49 89 538863 878

Martin.Schmidt@becketal.com

linkedin.com/in/
martin-schmidt-526267153

@MSSputnic

- Senior IT Architect at [Beck et al.](#)
 - IBM / HCL Connections since 2007
- Experience in
 - Deployment, Migration, Operation and Customization
 - Kubernetes EKS and AWS, Container
- Focusing on
 - enhancing with own containers
 - make the most out of customizer
- More and more
 - DevOps, Automation



Christoph Stoettner



+49 173 8588719

christoph.stoettner@panagenda.com

linkedin.com/in/christophstoettner

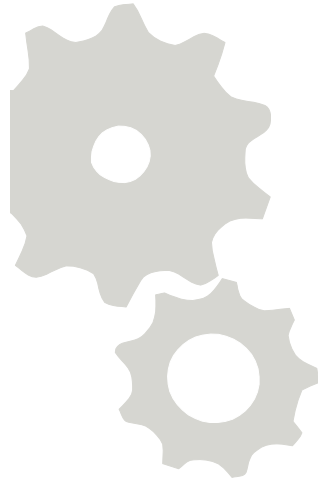
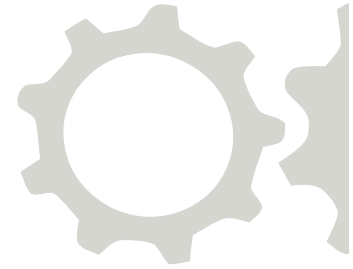
stoeps.de

christophstoettner

@stoeps



- Senior Consultant at [panagenda](https://panagenda.com)
 - IBM Connections since 2009
- Experience in
 - Migrations, Deployments
 - Kubernetes, Container
 - Performance Analysis, Security
- Focusing in
 - Security, Deployments
- More and more
 - DevOps, Automation, Ansible



Containerization

engage^{gears}

Container?

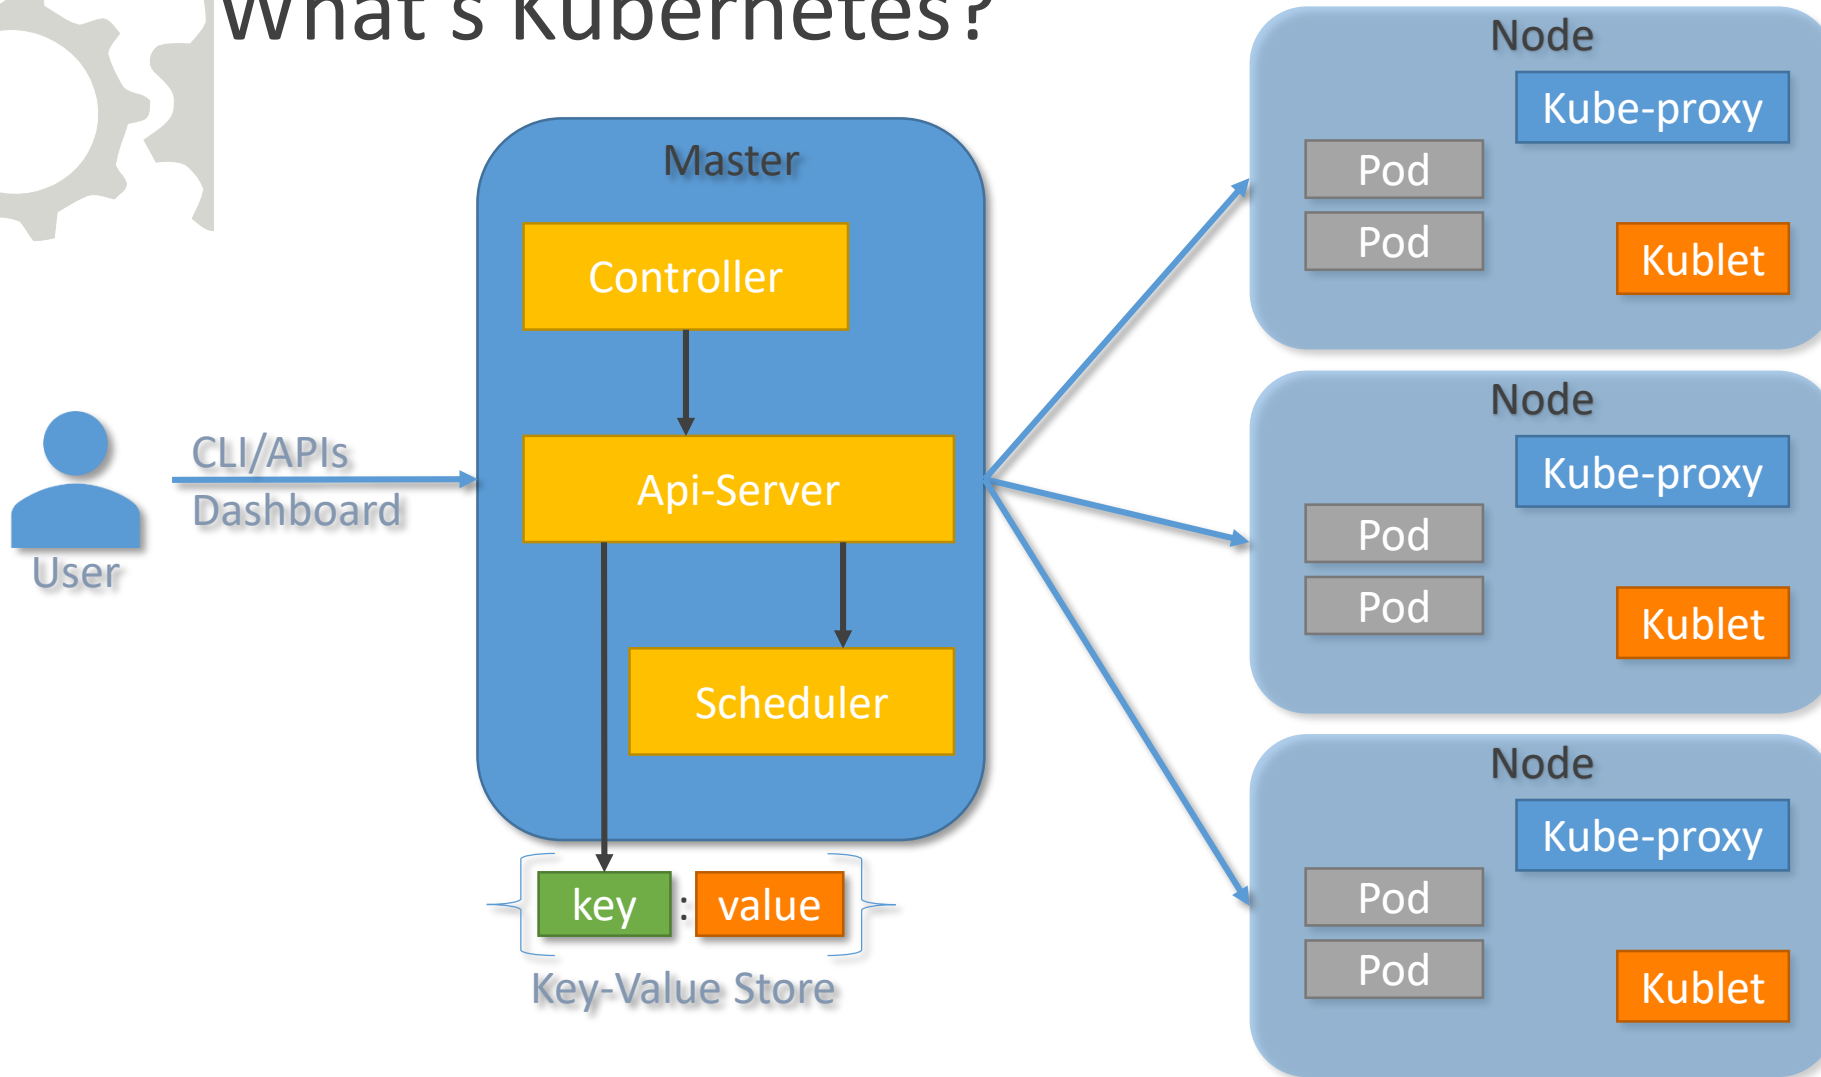




Kubernetes is containers on steroids



What's Kubernetes?

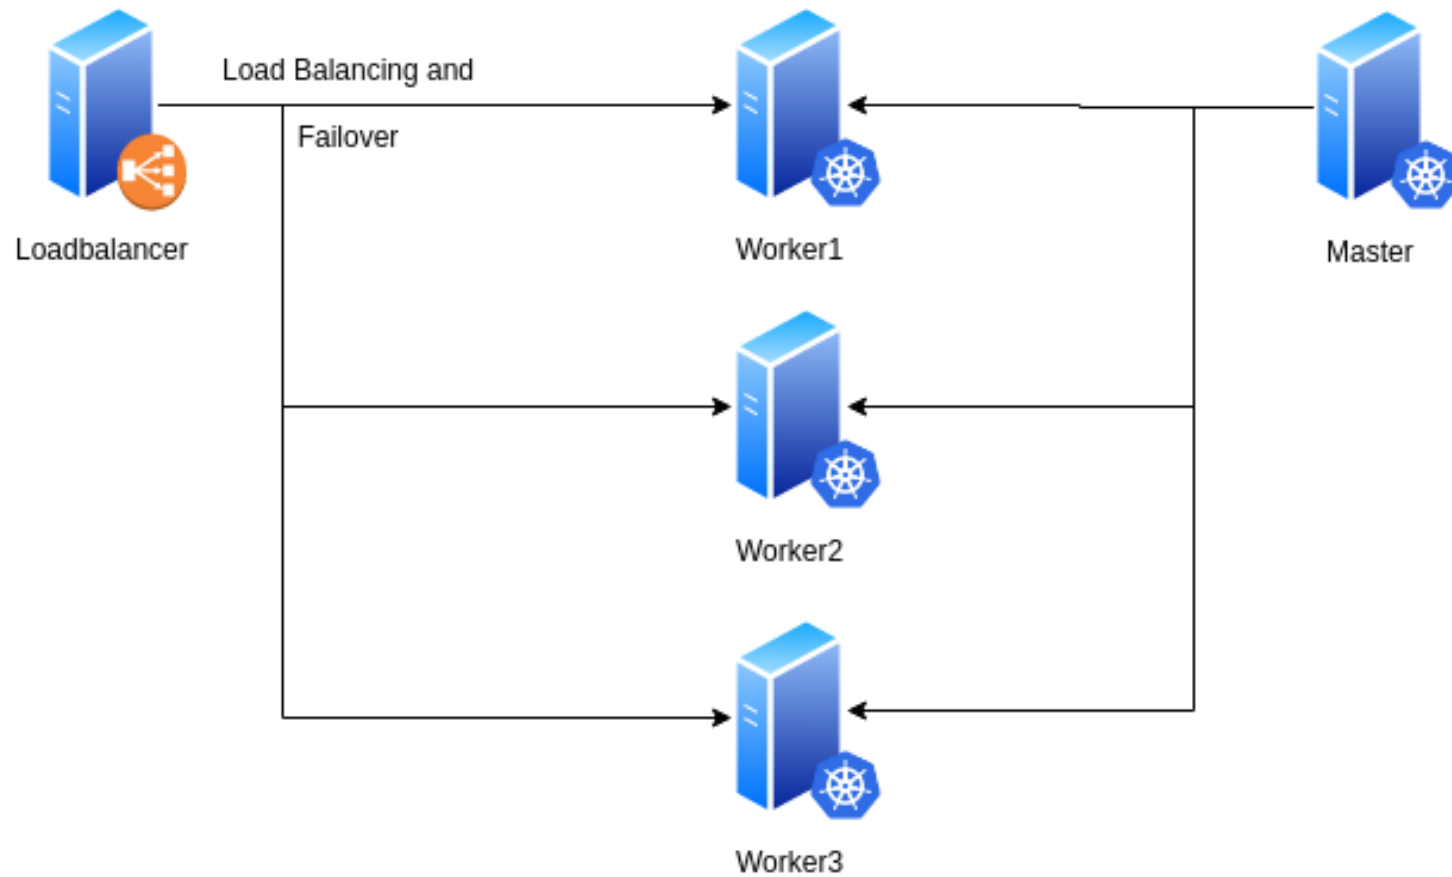


Cloud



On Premises

Infrastructure example



Cloud

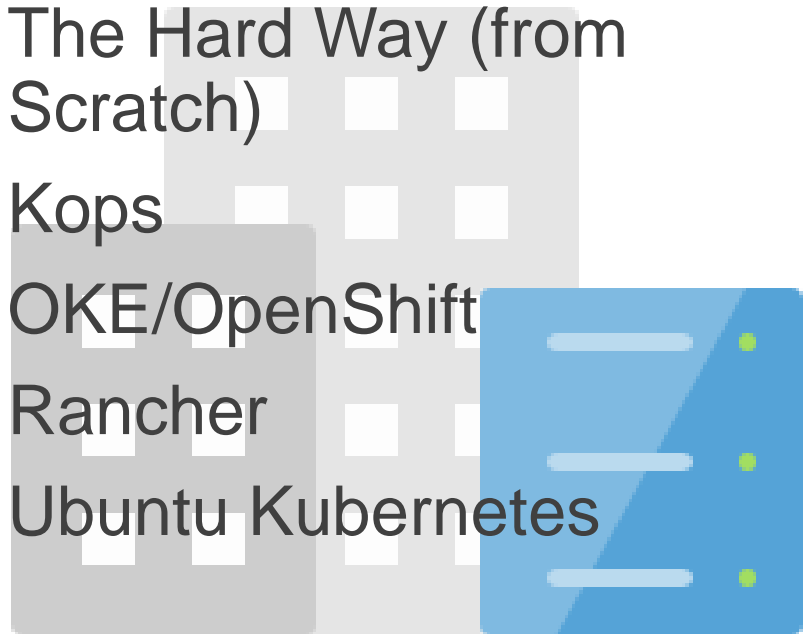


On Premises



Distributions / Cloud provider

- AWS EKS
 - Google GCE
 - Azure
 - Digital Ocean
- 

- Kubeadm
 - The Hard Way (from Scratch)
 - Kops
 - OKE/OpenShift
 - Rancher
 - Ubuntu Kubernetes
- 



There are over 90 Certified
Kubernetes offerings.



<https://www.cncf.io/certification/software-conformance>



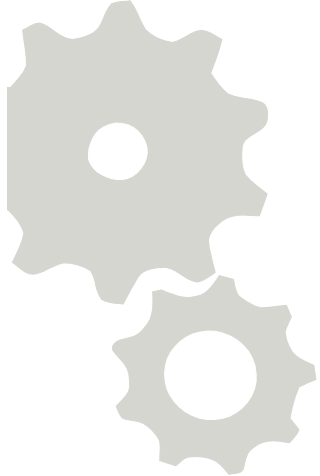
Cloud



On Premises 10



Disclaimer

- HCL currently supports the documented reference installation only.
 - Other installations are not officially supported (but they will not refuse just you are not on the reference platform)
 - In the future this will change....
- 

Supported Kubernetes for ComponentPack 6.5

- Red Hat 7.6 and CentOS 7.6
- Docker 17.03 (EE/CE + devicemapper storage)
Docker 18.06.2+ (CE + devicemapper storage)
- Kubernetes version [1.11.9](#)
- Calico v3.3 used as the network add-on
- Helm v2.11.0

The Component Pack was validated on a Kubernetes v1.11.6 cluster that was set up on virtual machines using the kubeadm tool.

This limits us a little bit and of 90 compliant options only few are left






Supported Cloud provider with version 7


Improved compatibility
with different offerings of Self and Provider Managed Kubernetes

Announcing Support for with Connections v7




Amazon
EKS


Amazon Elastic
Kubernetes Service



AKS: MS Azure
Kubernetes Service



GKE: Google
Kubernetes Engine

 **HCL SOFTWARE**

Copyright © 2020 HCL Technologies Limited | www.hcltechaw.com



Deployment Options – Testserver

- All services on **one** server (master and worker):
 - 16 CPU, 2.x GHZ, 64GB memory, 100GB disk.
 - 50GB+ for Device Mapper block storage

16 CPU cores
64 GB memory
150 GB disk

Additional requirements for Customizer

1 - Reverse proxy: for Customizer: 4 CPU, 2.x GHZ, 4GB memory - 100GB disk



Deployment Options – Standard

- 1 - Master: 4 CPU, 2.x GHZ, 16GB memory, 100GB+ disk
 - 50GB+ for Device Mapper block storage.
- 1 – Generic Worker
 - 8 CPU, 2.x GHZ, 24GB memory, 100GB disk.
 - 50GB+ for Device Mapper block storage
- 1 – Infrastructure Worker
 - 8 CPU, 2.x GHZ, 24GB memory, 100GB disk.
 - 50GB+ for Device Mapper block storage.
- Storage: Persistent volumes for Elasticsearch, Customizer, MongoDB, Zookeeper, and Solr indexes, 100GB disk.

Additional requirements for Customizer


- 1 - Reverse proxy: for Customizer: 4 CPU, 2.x GHZ, 4GB memory - 100GB disk



Deployment Options – high available

- 3 - Masters
 - 4 CPU, 2.x GHZ, 16GB memory, and at least 100GB disk.
 - Add 50GB+ per master for Device Mapper block storage.
- 3 - Generic Workers
 - 6 CPU, 2.x GHZ, 24GB memory, and 100GB disk.
 - Add 50GB+ per master for Device Mapper block storage.
- 3 - Infrastructure Workers
 - 6 CPU, 2.x GHZ, 24GB memory, and 100GB disk.
 - Add 50GB+ per master for Device Mapper block storage.
- Storage (available on all nodes)
 - Elasticsearch, Customizer, MongoDB, Zookeeper, and Solr indexes - 150GB disk.
- Loadbalancer (haproxy, i5, nginx)

48 CPU cores
156 GB memory
2.4 TB disk



Additional requirements for Customizer

1 - Reverse proxy: for Customizer: 4 CPU, 2.x GHZ, 4GB memory - 100GB disk

On Premises



What do we want to achieve? How to start?

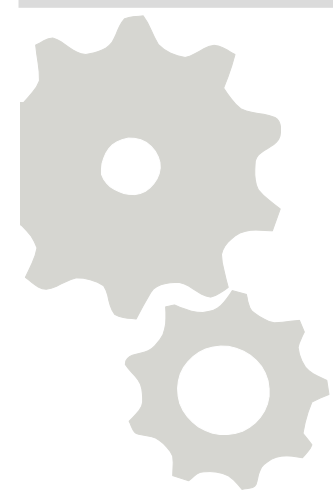
- Get a cloud like experience
 - Automate the deployment process
 - Choose a networking solution
 - Choose a storage solution
 - Handle security and authentication
- Ansible can help automating the installation
 - Prepare VM
 - Disable Swap
 - Install Prerequisites
 - ...
 - Deploy Kubernetes




ANSIBLE



Build machines

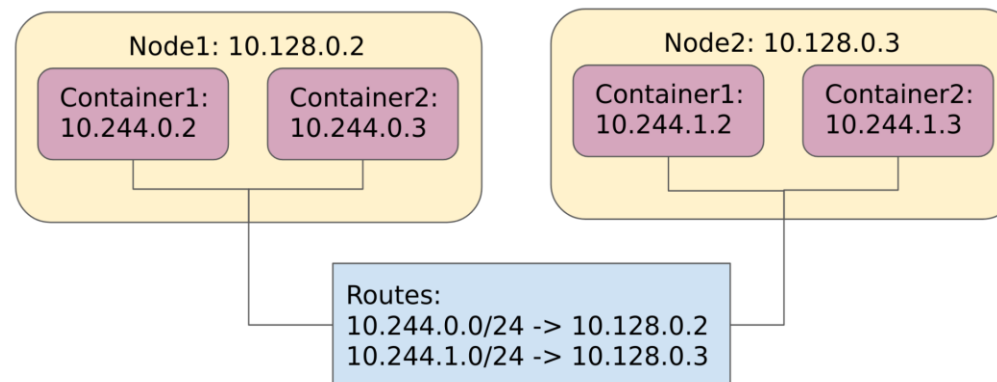
- Bare Metal or Virtual Machines
 - If you plan high availability workers and nodes
 - Cluster must be distributed across multiple hypervisors
 - Hardware errors can happen
 - Think about dumb machines which doesn't hold data
 - Just running containers (pods)
- 

Operating System

- **Supported: Red Hat 7.6 or CentOS 7.6 with devicemapper**
- Better optimized
 - Fedora CoreOS.
 - Automatically updating Linux OS for containerized workloads.
 - RancherOS
 - lightweight, secure Linux, built from containers to run containers
- Evolution from Infrastructure as Code
 - No separate Packer -> Terraform -> Ansible -> Kubernetes
 - Provide a pxeboot.cfg and ignition file on a webserver
 - VM grabs config from pxe and ignition file
 - To update or repair, just reboot into new image

Networking

- **Supported: Calico**
- CNI – Container Network Interface
 - [Calico - a layer 3 virtual network](#)
 - [Weave - a multi-host Docker network](#)
 - [VMware NSX – enables automated NSX L2/L3 networking and L4/L7 LB](#)
 - Many more





High available network storage

- Replication
 - Ceph
 - GlusterFS
 - Rancher Longhorn
- Amount of disk space is needed multiple times
- Hardware storage
 - NAS / SAN
 - High available





Kubeadm

- Script/Binary to prepare Kubernetes environment
- Saves from copying “hundreds” of certificates
- Manually deploy and update masters and nodes



Alternatives

- Kubespray
 - Ansible roles to build Kubernetes Cluster
- Rancher
 - Run it free or buy support
- OpenShift
 - Special
 - Secure, but special
- Not for ComponentPack,
but k3s.io even runs on Raspberry Pi





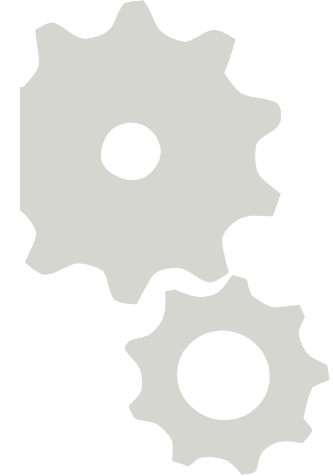
Updates

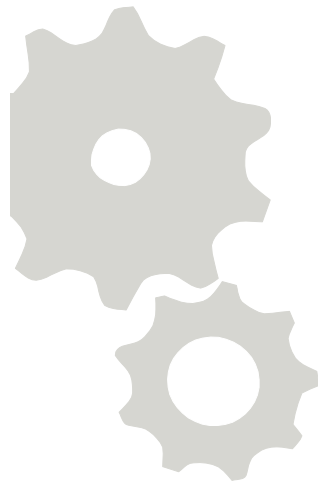
- You only can upgrade
 - from one MINOR version to the next MINOR version 1.16 → 1.17
 - between PATCH versions of the same MINOR 1.16.1 → 1.16.5
 - So going from 1.11.6 to 1.17.3 means
 - 1.11.6 → 1.12.x → 1.13.x → 1.14.x → 1.15.x → 1.16.x → 1.17.3
- The upgrade workflow at high level is the following:
 - Upgrade the primary control plane node
 - Upgrade additional control plane nodes
 - Upgrade worker nodes
 - One at a time
 - Zero downtime
 - In place or add temporary additional nodes





Our goal

- We should end with
 - Supported Kubernetes Deployment
 - Supported Connections / ComponentPack
 - Kubernetes supports
 - Last 3 minor versions
 - Today: 1.15, 1.16, 1.17
 - Several companies sell support contracts for Kubernetes
 - Or build the skills inhouse
- 



Cloud



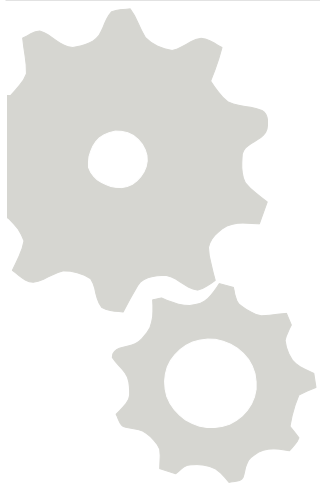
Buy Kubernetes from a Cloud Provider

- Kubernetes Cluster can be purchased from different provider
 - AWS *
 - Azure *
 - Google *
 - Digital Ocean
 - IBM Cloud
 - ...
- They are all a little bit different and you need to check if the products fit your needs.
- **None of them matches the “Reference Implementation” => currently not supported -> * v7 will change this**



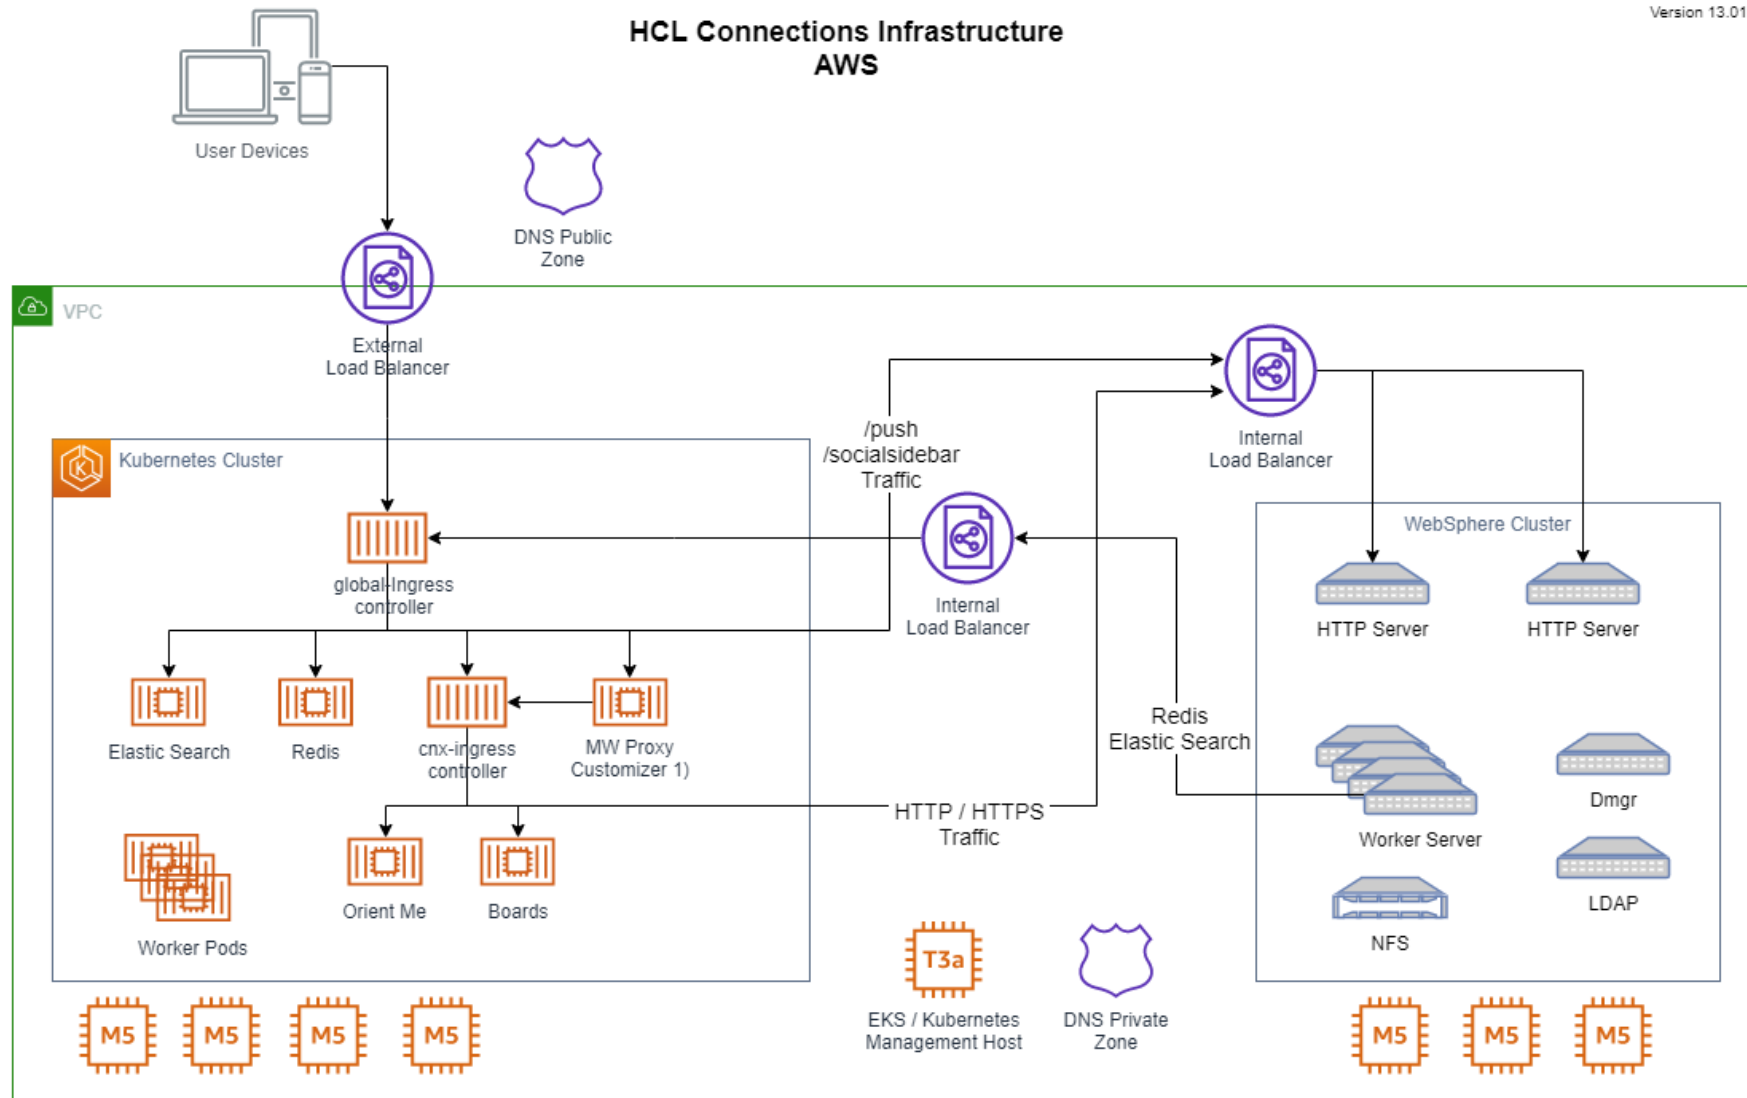


Component Pack on AWS

- AWS has the most complete offer (as far as I know)
 - Azure was tested as well (Component Pack 6.0.0.7)
 - We have tested Component Pack 6.0.0.7 – 6.0.0.9 and 6.5.0.0 on EKS.
 - Other services can also be used
 - Elasticsearch Service (for metrics and type ahead search of classic infrastructure. For OrientMe it does not work – open PMR)
- 

HCL Connections Infrastructure AWS

Version 13.01.2020



EFS
File Store

ECR
Container Registry

Kubernetes
Master

DNS
Services

1) The mw-proxy service (customizer component) forwards the traffic to the cnx-ingress-controller and not to the backend connections http servers.
Advantages: Everything can be customized, No proxy configuration on existing HTTP Server, No load balancer for the cnx-ingress-controller necessary.
Disadvantage: The mw-proxy pod needs to be customized to support http as protocol to the cnx-ingress-controller service





Installation

Pro

- With a single command you get an operational Kubernetes Cluster. (eksctl)
- With a single command you get an operational Elasticsearch cluster.

Con

- You get only dedicated Kubernetes / Elasticsearch versions.
- Limited individual setup possible.
- Minimum Infrastructure Requirements must be available.





Operations

Pro

- AWS is managing the Master / Cluster for you.
- Restart / Recreate / Add / Remove Kubernetes node / Elasticsearch node is just 1 command.

Con

- In case of problems, you have no access possibilities to the Kubernetes Master other than API.
- In case of problems, you can not install patches or fixes on Kubernetes Nodes.
- In case of problems, you have no direct access to the Elasticsearch Nodes.





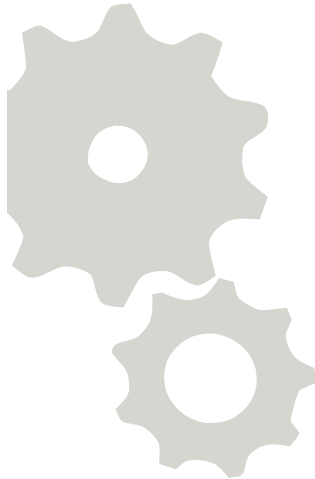
Support

Pro

- AWS can help you with the Kubernetes Cluster / Elasticsearch Cluster (Professional Services)
- The infrastructure can be build reproducible without effort (e.g. Cloudformation, eksctl)
- Can be easily integrated with other products (e.g. Lambda, CodePipeline, CodeBuild, CloudWatch, ...)

Con

- Officially not (yet) supported by HCL





Update

Pro

- Update to a later Kubernetes version is just one command.
- Update to a later Elasticsearch version is just one command.
- AWS makes sure these updates work.

Con

- Updates are only possible to dedicated versions.
- AWS removes support for older versions of Kubernetes and you must update to a newer version.
- AWS removes support for older versions of Elasticsearch and you must update to a newer version.





Used Services

Pro

- Kubernetes (EKS)
- Elasticsearch*
- Load Balancer (Classic Load Balancer)
- NFS Server (efs) for persistent storage
- Docker Registry (ECR)
- DNS (Route53)

Con

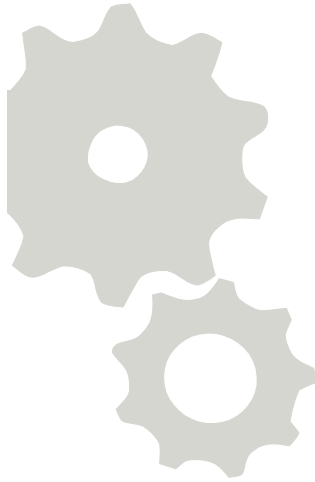
- Every service has a price tag.
- 





AWS cost estimation

- Based on
 - 1year reserved instances, no upfront
 - Instances hosted in Ireland (zone eu-west-1)
 - HA deployment
 - Elastic Search Server on Kubernetes (Not recommended due to storage problems but currently AWS Elasticsearch does not work for OrientMe)
 - EC2 m5.xlarge are smaller than recommended and m5.2xlarge are larger than recommended
- Price calculation from 27. February 2020
 - Prices for network traffic, backup storage and AWS support is not included
 - No free tier included
 - Classic infrastructure is not included



Rough Price Estimation *

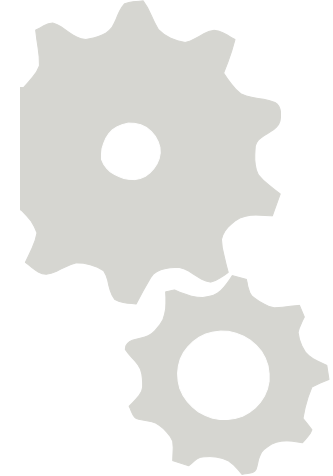
AWS resource	Size	Number required	Price per unit / per resource	Price per month
EKS Master		1	\$ 0.10	\$ 73.0
EKS Worker Node	m5.xlarge	6	\$ 0.1500	\$ 657.00
HDD	100 GB	6	\$ 0.11	\$ 66.00
EFS Share	200 GB	1	\$ 0.33	\$ 66.00
Classic LB		2	\$ 20.50	\$ 20.50
Admin Host 20%	t3.medium	1	\$ 0.0408	\$ 6.71
Admin Host HDD	30GB	1	\$ 0.11	\$ 3.30
ECR	30GB	1	\$ 0.10	\$ 3.00
Sum				\$ 895.51

* no responsibility is taken for the correctness of this information.





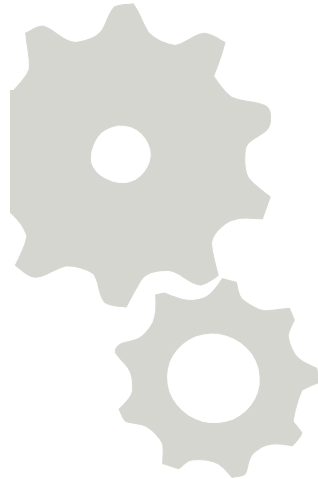
Summary

- In case your company already has AWS or Azure know how but lacks Kubernetes know how EKS or AKS is a good start.
 - Due to the lack of support we can't recommend it for production use (yet) but for test, development or prove of concept this infrastructure works fine.
- 



Links and references

- Setup Component Pack in the cloud:
<https://becketalservices.github.io/beas-cnx-cloud/>
- Kubernetes reference implementation:
https://help.hcltechsw.com/connections/v65/admin/install/cp_prereqs_kubernetes.html
- Kubernetes: <https://kubernetes.io>
- K8s documentation (Tutorials, Basics, Start):
<https://kubernetes.io/docs/home/>
- Cloud Native Computing Foundation (CNCF):
<https://cncf.io>



Cloud



On Premises 39



<https://www.flickr.com/photos/wwwworks/4759535950>