Scalable and secure hosting of Cognitive Enterprise Systems using Managed Kubernetes

Autoren: Martin Voigt & Yannic Ahrens

AI4BD

AI4BD delivers Cogntive Enterprise System with Cogntive Business Robotics Coworker as a Software Products. The Method for “Cognitive Business Robotics” (CBR) enables SME and large enterprise for a new class of digitalisation with encapsulating human knowledge. This empowers organisation for new business models and efficient processes. The core of the unique CBR Coworker Product Architecture is the applied generic data structure, combined with AI and Knowledge Graphs. The AI4BD software platform enables to design and implement various AI-driven Coworker as Software products, so-called CBR Coworker, integral over further Data and Processes. The global first Cogntive Enterprise System for two lighthouse projects are currently based on two CBR Coworkers “Cognitive Secretary” (COSY) and “Cognitive Data Scientist”. Already this first step shows significant and measurable value add. AI4BD is in partnering with IBM & SAP to support SAP Partner, SAP Clients to prepare their Data and Processes to enable easier and faster migration and make better use of SAP S/4HANA. Therefore, several AI4BD Partners are already trained and being certified. SW companies and Partners such as Oxaion, powered by CBR Coworker are so called CBR Systems and can enhance their software portfolio with affordable AI capabilities for their clients with measurable added value and new class of digitalisation with high performance also for Big Data. AI for Big Data for Cognitive Enterprise System.

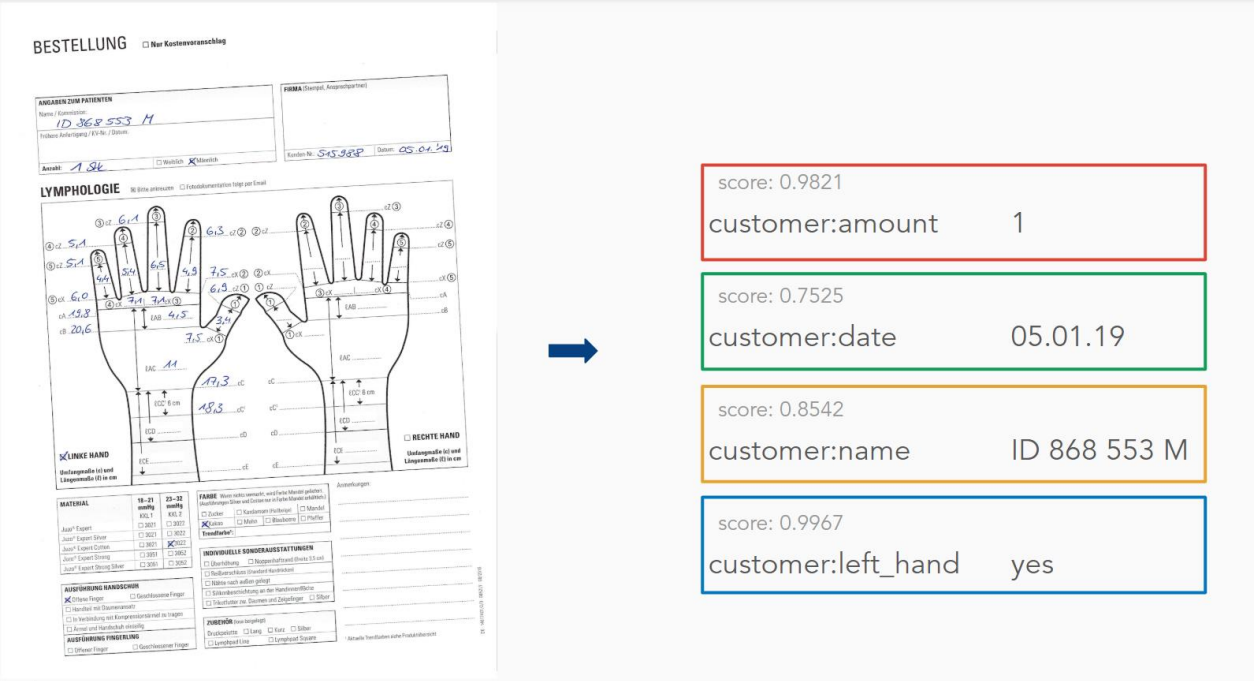

Figure 1: Example of Cognitive Secretary (COSY) for semantic knowledge extraction from various documents

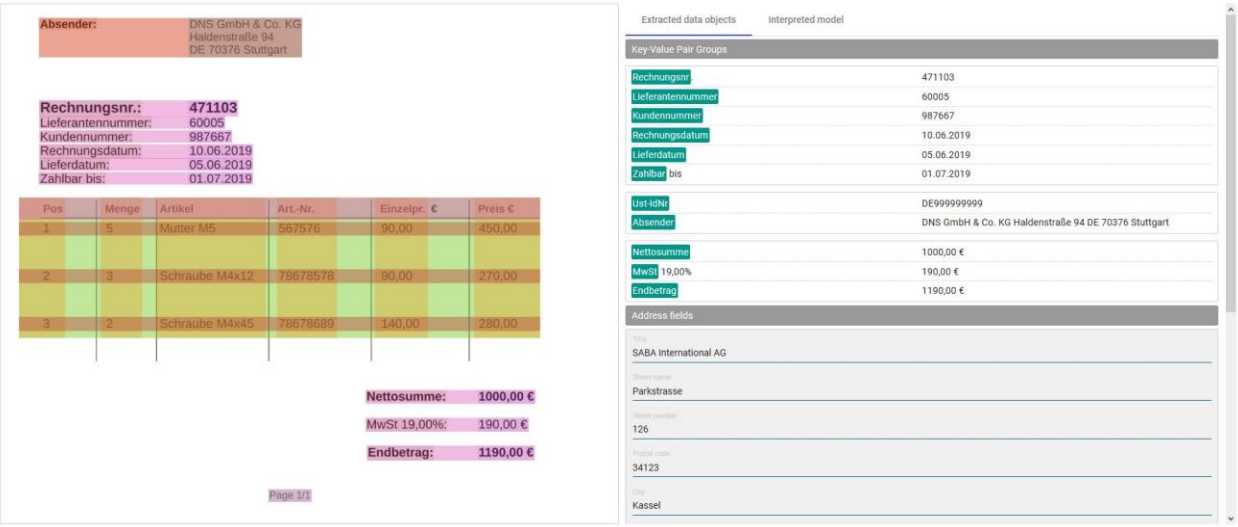

Figure 2: Second Example of Cognitive Secretary (COSY) for semantic knowledge extraction from various documents

From the beginning of the development, the platform is designed to scale – regarding new AI-driven features but also to enable AI4BDs customers to work with small and big data (regarding size and velocity) on demand. To depict the required load: current customers require the analysis of one page within COSY in max. 10s. Therefore, about 400 steps in the ETL pipeline including the inference of about 10 machine learning models is required.

AI4BD's Software-Stack

To achieve this, the software design is driven by a thorough microservice architecture orchestrated using a mature and scalable ETL stack. Rolled-out using a Docker-based containerization, scaling is easy – we thought. As it is possible to setup new instances of a container easily by the DevOps team, it is cumbersome and finally not to handle in a manual way. Thus, AI4BD started to test with Docker Swarm as a native solution but it showed up fast, that this orchestration had different constraints making it impossible for a productive use. Here, initial tests with Kubernetes (K8s) and proper K8s-training by an expert revealed the way to go. However, we asked our-self: is it our value proposition to our customers to setup and manage a Kubernetes stack? No. We need assistance.

Cloud&Heat

Cloud&Heat was and is future-driven cloud provider to jointly define and reach the goals to host our Cognitive Business Robotics platform using K8s. From being an IaaS customer, we turned our head towards Platform-as-a-Service (PaaS). AI4BD tinkered in a self-deployed K8s cluster and came up with three major pains: the cluster should be highly available, protected against attacks from the outside and well-integrated into the block and object storage provided by Cloud&Heat's own OpenStack distribution. Ideally it should also offer a shared filesystem that is needed by, among other things, the distributed MongoDB instances.

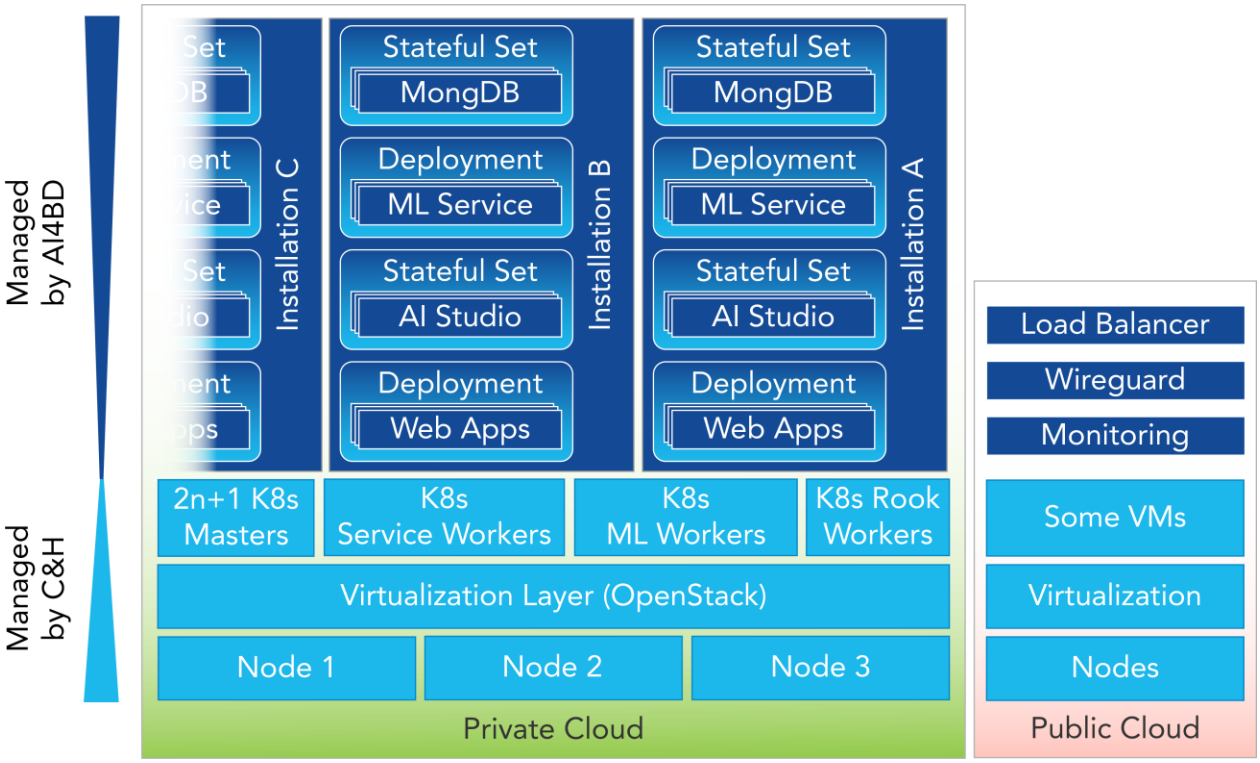

Figure 3: Managed Kubernetes setup provided by Cloud&Heat Technologies for AI4BD applications

Managed K8s

The cluster itself is deployed with kubeadm and only reachable inside a private network. Three master nodes ("multi master") ensure that the cluster is highly available (HA) and can cope with the outage of one node. Another three VMs act as gateways and the endpoint of the cluster. They host HAProxy as a load-balancer and keepalived to handle the virtual public endpoint. The private network can be reached over a wireguard VPN tunnel.

Storage

The cloud controller manager is interfacing with the surrounding OpenStack environment. By using specific storage classes, a block storage can be dynamically allocated to satisfy a Persistent Volume Claim (PVC). It's also responsible for accessing object storage and providing a dynamic load balancer resource (K8s service type). The requirement of a shared file-system was initially addressed by a NFS-provisioner inside the cluster. A big block volume was attached to a pod that creates NFS shares on demand. All NFS shares were hence located on the same volume. Issues like performance, missing quotas and no HA-capability lead to the introduction of rook.io which spawns a Ceph cluster inside the K8s cluster. Effectively it is a Ceph-inside-Ceph cluster because Cloud&Heat's storage uses Ceph with a replica count for 3. To mitigate the storage overhead the replica count in the K8s Ceph was set to 1.

Deployment

The deployment of a fresh cluster happens in four phases. Phase 1 is executed by terraform and creates the necessary VMs, security rules and the network. In phase 2 the gateways are bootstrapped, phase 3 is the deployment of the cluster with kubeadm and phase 4 some rudimentary smoke tests. Phases 2-4 are orchestrated by Ansible. The entire cluster configuration is declarative. kube-prometheus creates a surprisingly sophisticated monitoring environment based on Prometheus, Grafana and Alertmanager that is also used by the operations team.

Security

AI4BD and Cloud&Heat are currently in discussion to increase the security of the platform even further so that the CBR software can be used in a "Kritische Infratrukturen" (KRITIS) environment. In the future, Managed K8s will be deployed on top of SecuStack, a security hardened OpenStack version developed by secustack. secustack is a joint venture owned by Cloud&Heat and the IT security specialist secunet. When rolled out, AI4BD and its customers can profit from a series of additional security features:

- SecuStack is Open Source based and therefore auditable on the source code level.

- The intra-cloud network traffic between servers inside the DC on which the VMs run on is encrypted by default. In upcoming releases more fine-grained levels of tenant separation, chosen by the customer, will be provided. E.g., the customer can opt for dedicated network traffic between all of theirs VM.

- If required, both the customer and the operator can access the cluster with the BSI-certified SINA L3 Box. This feature is encapsulated in "VPN as a Service" (VPNaaS).

- All operating system images are cryptographically validated. This ensures the integrity of the K8s nodes on deployment.

- Particularly sensitive user data can be stored in encrypted OpenStack block storage volumes. SecuStack offers a sophisticated way to deal with key material. This user-specific encryption is on top of the general volume encryption of all devices in Cloud&Heat's Ceph Cluster.

- Use of hardware support like Hardware Security Moduls (HSM), Intel® Software Guard Extensions (Intel® SGX) and Secure Boot ensure the integrity of the complete hardware and software stack. This will be addressed in an upcoming release and offered as a configuration option to the customer.

Current Deployment

AI4BD's k8s cluster run on dedicated hardware in a co-location spot in Cloud&Heat's major data center in Frankfurt a. M. Different virtual machine sizes (flavors) were tested for the worker nodes to make best use of the hardware resources for database and service applications. AI4BD successfully deployed its Cognitive Secretary inside the cluster. The defined and implemented managed K8s environment are currently used for AI4BD's internal development team (staging), for the consultants (demo) and the customers (productive). Especially the latter will grow on demand due to the implementation of new customers over the year.