Valohai

ML Platform Infrastructure

This whitepaper is intended for DevOps and infrastructure teams. It will give you an overview of Valohai and how it integrates into your existing IT infrastructure.

INTRODUCTION:

Valohai is a machine learning platform that lets data scientists run ML experiments in the cloud and on-premises without worrying about the hardware setup, workload management or version control. It automatically scales servers on demand, manages workloads between users and shuts servers down to cut costs. Companies save money and data scientists can work autonomously, without the continuous support of a DevOps or ML Engineering department.

On top of machine orchestration, the platform automatically maintains a version control history for each experiment and data lineage for the created models and other data assets. This creates an audit trail that the team and the company can use to work together and solve compliance issues.

LOGICAL TOPOLOGY OF VALOHAI:

Valohai consists of three logical parts:

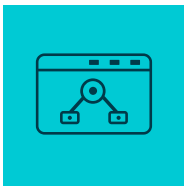

The orchestrator

that launches individual workers and keeps track of metadata about every experiment. This is the web UI and the REST API end-point.

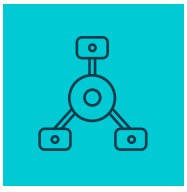

The workers

that the orchestrator launches and that conduct the actual experiments on the selected CPU/GPU/TPU instances, taking care of Docker container setup, data & code fetching and storing the data into selected output folders (e.g. AWS S3, Azure blob, Google cloud, or OpenStack Swift storages).

The integrations

to tools such as Jupyter notebooks or the CLI that integrates with the orchestrator's API. There are several ways to integrate Valohai into your existing CI/CD pipelines or other machine learning tools like external hyperparameter optimizers or Airflow.

CONFIGURATION ALTERNATIVES:

Valohai can be installed in five configurations:

1

Cloud orchestrator & Cloud workers

The default installation is to run Valohai orchestrator at app.valohai.com and the cloud workers under Valohai's account. No external agreement with AWS, GCP or Azure is required and usage will be billed according to the publicly available price list of each cloud provider. This setup usually requires less than an hour to set up.

2

Cloud orchestrator & Private workers

The second alternative is to hook up Valohai with the customer's own cloud account and have the orchestrator running at app.valohai.com. This way the customer will be using their own cloud agreement with AWS, GCP or Azure and will pay cloud usage directly to the provider. As all of the actual computation happens in the private cloud; none of the training data or other artifacts ever leave outside the customer cloud account. This setup usually requires a few hours.

3

Private orchestrator & Private workers

Valohai can also be installed fully into the customer's cloud environment, where both the orchestrator as well as the workers are running in their own cloud environment. The orchestrator will be running in the customer's own cloud/intranet environment and usually requires two days on-site to set up and train the DevOps team. Updates through VPN or other measures will be agreed upon separately.

4

Full on-premises installation

A full on-premises installation is the same as a private orchestration & private workers setup, with the difference that instead of using a public cloud environment both are set up on the customer's physical servers. This setup usually requires 2 to 4 days to install and train. The environment needs to be running Linux with capabilities to run Docker, Python 3.4+, and the wanted Nvidia drivers, depending on your use-case.

5

Hybrid solutions

On top of the above, various hybrid combinations are possible. Want to run the orchestrator in the cloud but workers on-premises and the cloud? No problem! Everything is possible and easy to set up. We will do the heavy llifting and guide your DevOps team to manage the possible worker combinations per project and access rights in the on-boarding sessions.

In all configurations your code can be hosted from a public/private/enterprise Git repository (e.g. GitHub, GitLab or Bitbucket), a Jupyter notebook or from the local machines of the data scientists. The data can be hosted in various cloud storages (AWS S3, Azure Blob, Google Cloud, OpenStack Swift, behind an HTTP end-point or local mount in on-premises installations). Your Docker images can be hosted from any accessible public or private Docker registry.

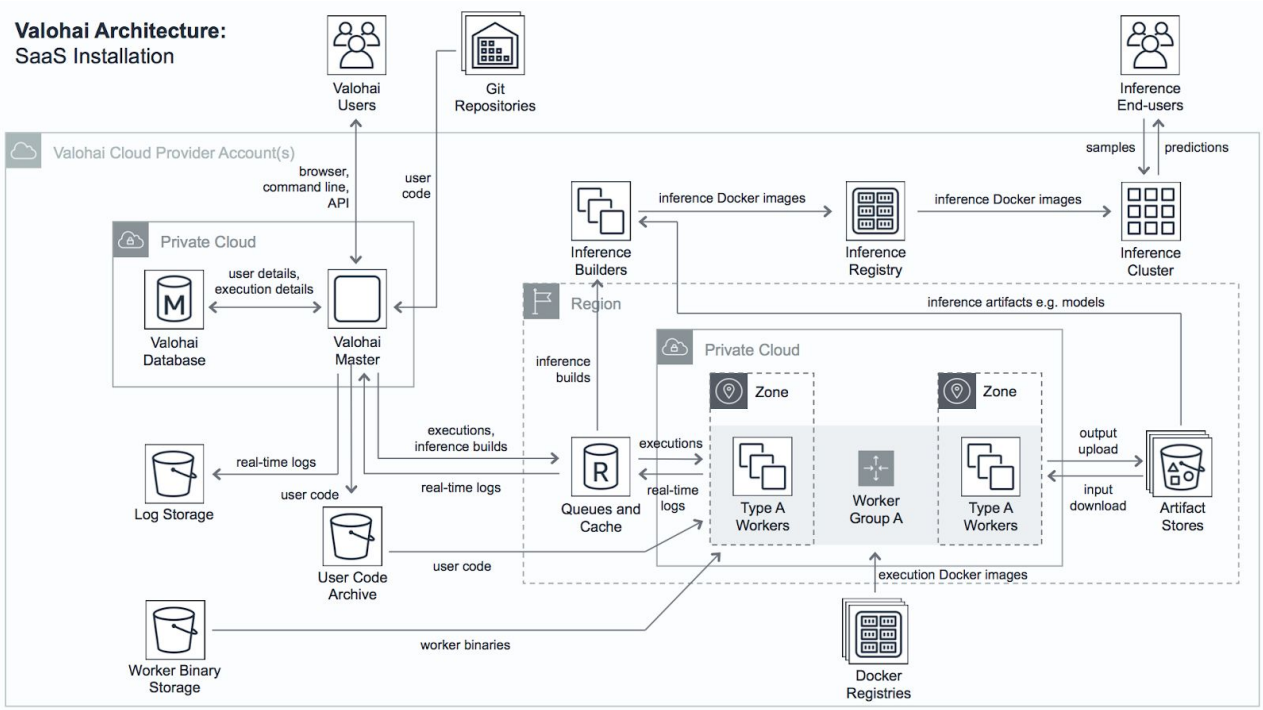

Example: Valohai running in a full cloud installation

INSTALLATION PROCESS:

Installation of Valohai is included in the platform's onboarding process. Depending on the solution you choose, the actual installation is from <1 hour (cloud) up to 4 days (exotic on-prem installations). Valohai will fly onsite and do the installation for you in conjunction with a training session to the data science teams. We will install the software and teach you to maintain it, after which we will concentrate on the data science teams so that they can run their experiments autonomously.

MAINTENANCE & SUPPORT:

Valohai is a managed service, meaning that support and updates are included in the license fee. Valohai also actively maintains the service either directly using provided privileges, through a jointly set-up VPN tunnel or physically on-site when required. Updates are released roughly twice per month including new features and bug fixes to both the web platform as well as integrations (e.g. Jupyter notebooks, CLI, API).

The DevOps and ML Engineering team at Valohai are reachable over Slack, email, or phone, and a separate SLA can be agreed upon in the agreement.

QUESTIONS?

If you have any questions, in regards to the installation of Valohai contact us at sales@valohai.com or check out architecture diagrams in the documentation.

The managed platform for building production-scale machine learning.