



Covalent Cloud is a cloud agnostic, entirely pythonic, and infinitely scalable computing platform providing on-demand, serverless access to AI infrastructure such as Nvidia H100s and A100s. With Covalent, users can scale to any arbitrary number of compute resources with a single line of Python code.

### Infrastructure as Python

With Covalent, you don't need to write a single line of IaC. No more YAML, scripting or Kubernetes. Infrastructure is entirely managed in Python.

### Serverless GPUs On-Demand

No reservations needed. No minimum commitment required.

- H100 80GB
- A100 80GB
- V100
- L40
- T4
- A10G
- RTX A6000
- & more!

### Simplified AI & High Compute

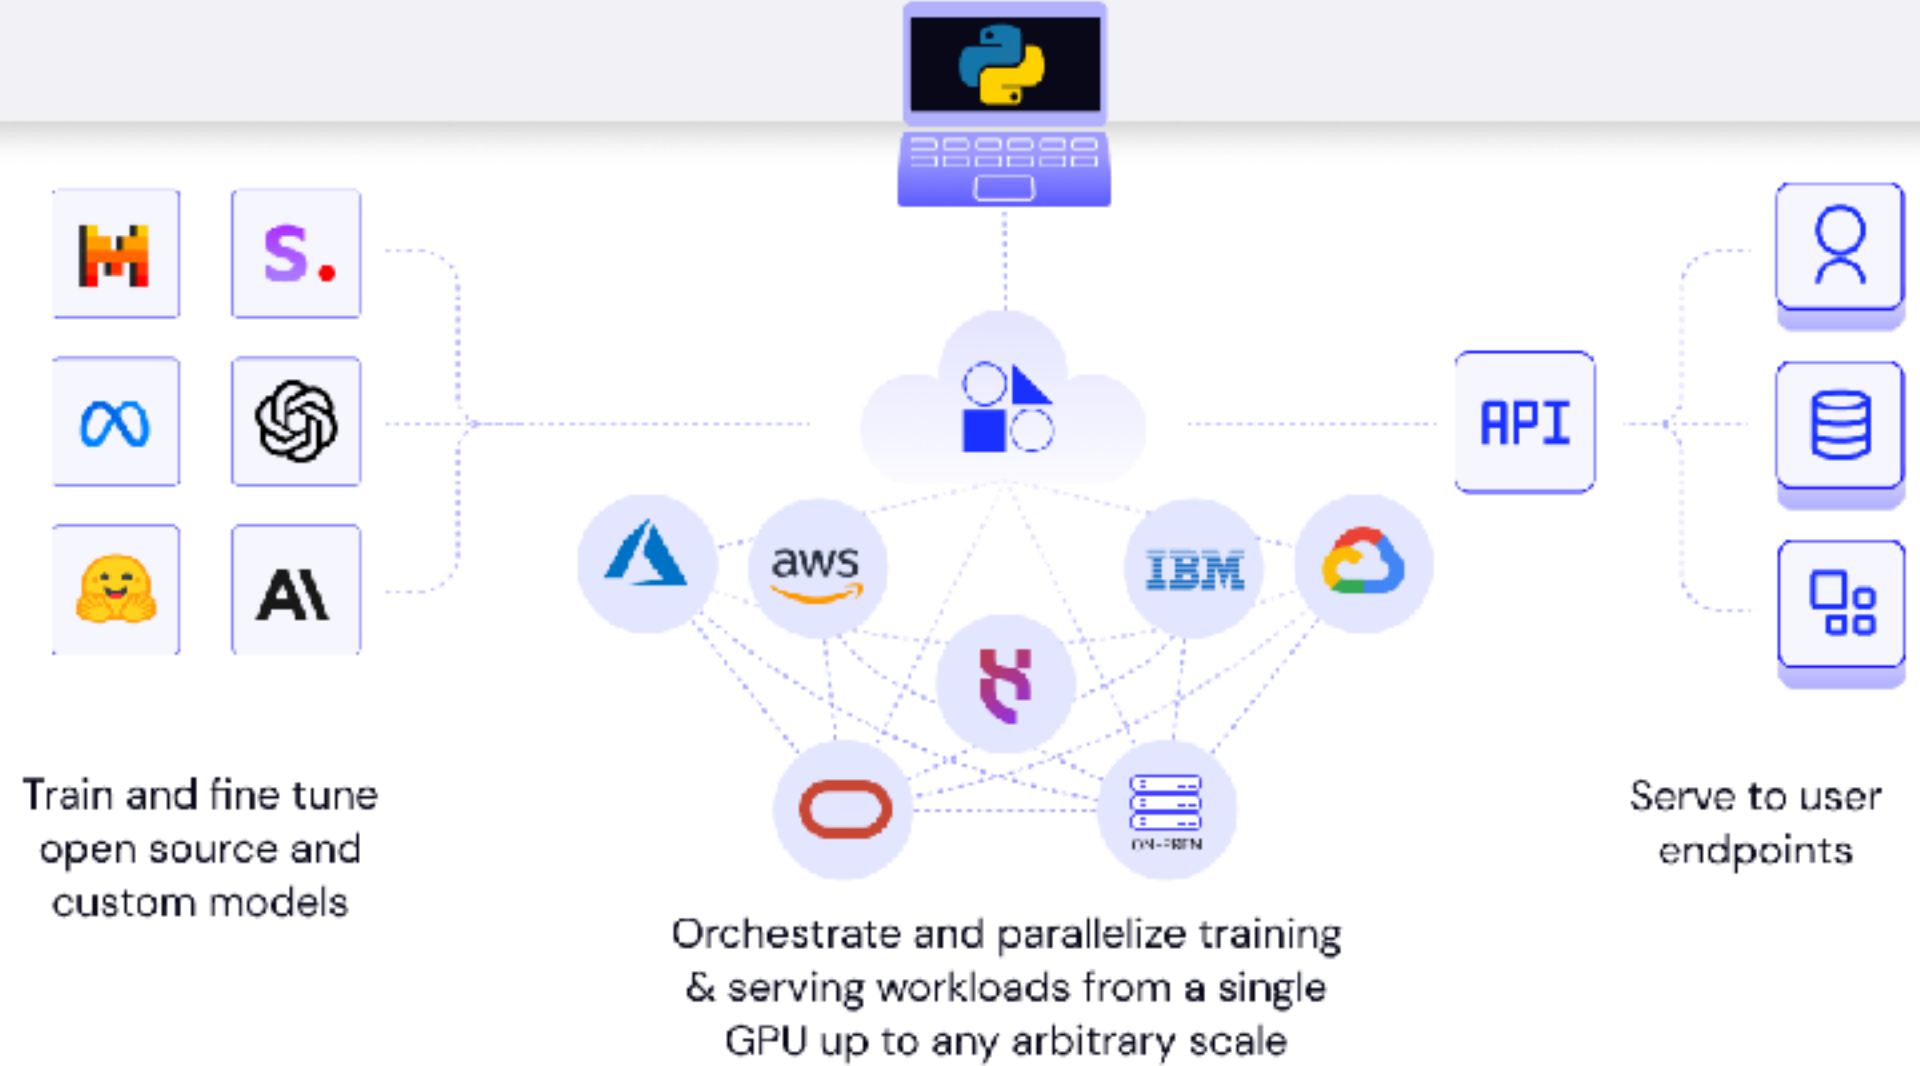
- GenAI / LLM
- Biology, Pharma, & Chemistry
- 3D, video, & audio processing
- Simulation
- Scientific computing

```

○○○
# pip install covalent-cloud
cc.create_env("llm_env", pip=["transformers"], conda=["cuda", "openmp"])

gpu_compute=CloudExecutor(env="llm_env", num_cpus=64, num_gpus=8, gpu_type="H100")

@cc.electron(gpu_compute) # Single line Infrastructure
def finetune_stable_diffusion_image(**prompt):
    ...
    
```



## Accelerate Your AI Journey

### Model training

Train custom AI models directly with Covalent or choose out-of-the-box open source models (e.g. Mixtral, Llama2 and more)

### Fine tuning

Fine-tune any AI model using Covalent without any DevOps. Entirely in Python, with no low level container or infra work

### Serving / Inference

Serve your models with a single-click deployment, in a cost-efficient and highly performant manner

## Become Cloud Agnostic

### Multi-Cloud Mobility

Get the lowest cost and best performance for any given workload with multi-cloud and multi-region support

### Eliminate DevOps/MLOps

You code, and we take care of infrastructure - from containerization, autoscaling, to provisioning & tearing down resources

### Bring Your Own Compute

Orchestrate workloads efficiently across your infrastructure - whether its on-premises, across clouds, or hybrid

### For Teams Large & Small

From dev teams of one to large research organizations sharing precious HPC resources, Covalent simplifies complex compute



### Cloud

- Get started with up to 30 hours of free compute

### Enterprise

- Multi-cloud, on-prem, and hybrid deployments

### Open Source

- pip install covalent