hystax

Hystax OptScale

MLOps open source platform

Multiply a number of ML/AI experiments with minimal cloud costs

hystax.com | info@hystax.com | 1250 Borregas Avenue Sunnyvale, CA 94089

Hystax

Founded in 2016,
customers in 48 countries

Customers: Airbus, Nutanix, Orange, Nokia, DHL, Burger King

OptScale use cases

Open source under Apache 2.0: https://github.com/hystax/optscale

OptScale schema

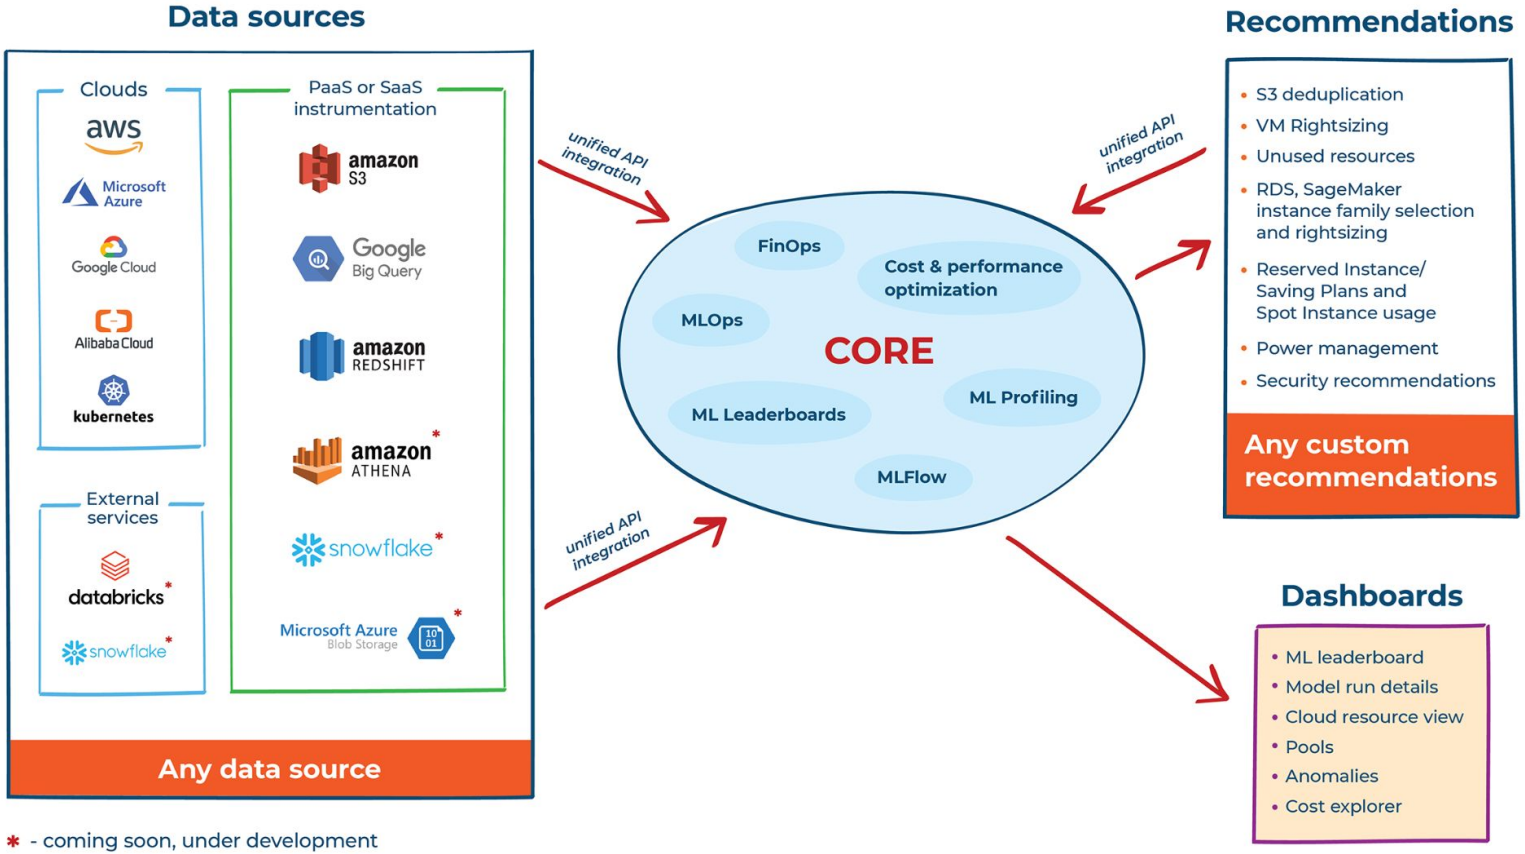

* - coming soon, under development

MLOps

- Team and individual ML engineer progress observability
- ML/AI task profiling, bottleneck identification
- PaaS or any external service instrumentation
- Optimization recommendations
- Runsets to automatically scale a number of experiments

ML R&D status observability

- List of models with goals status and active recommendations
- Tracking the number and quality of experiments run by a team
- Cost of an overall model and individual experiments

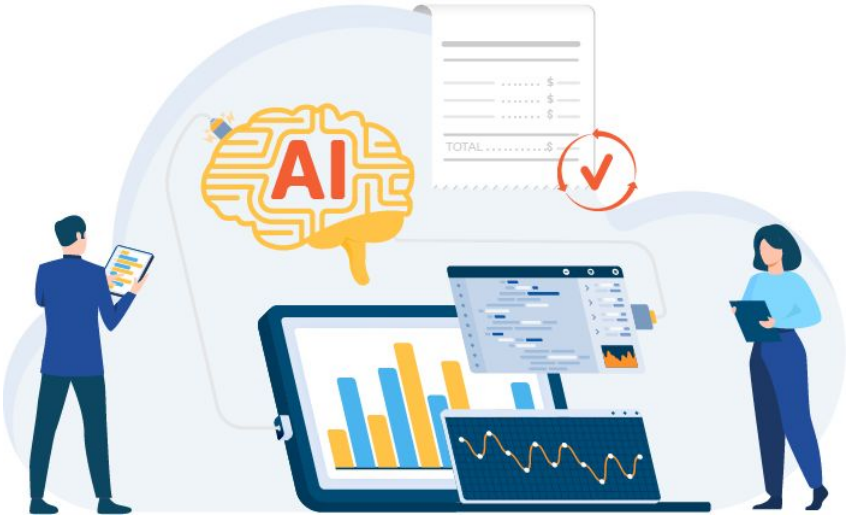

ML R&D status observability

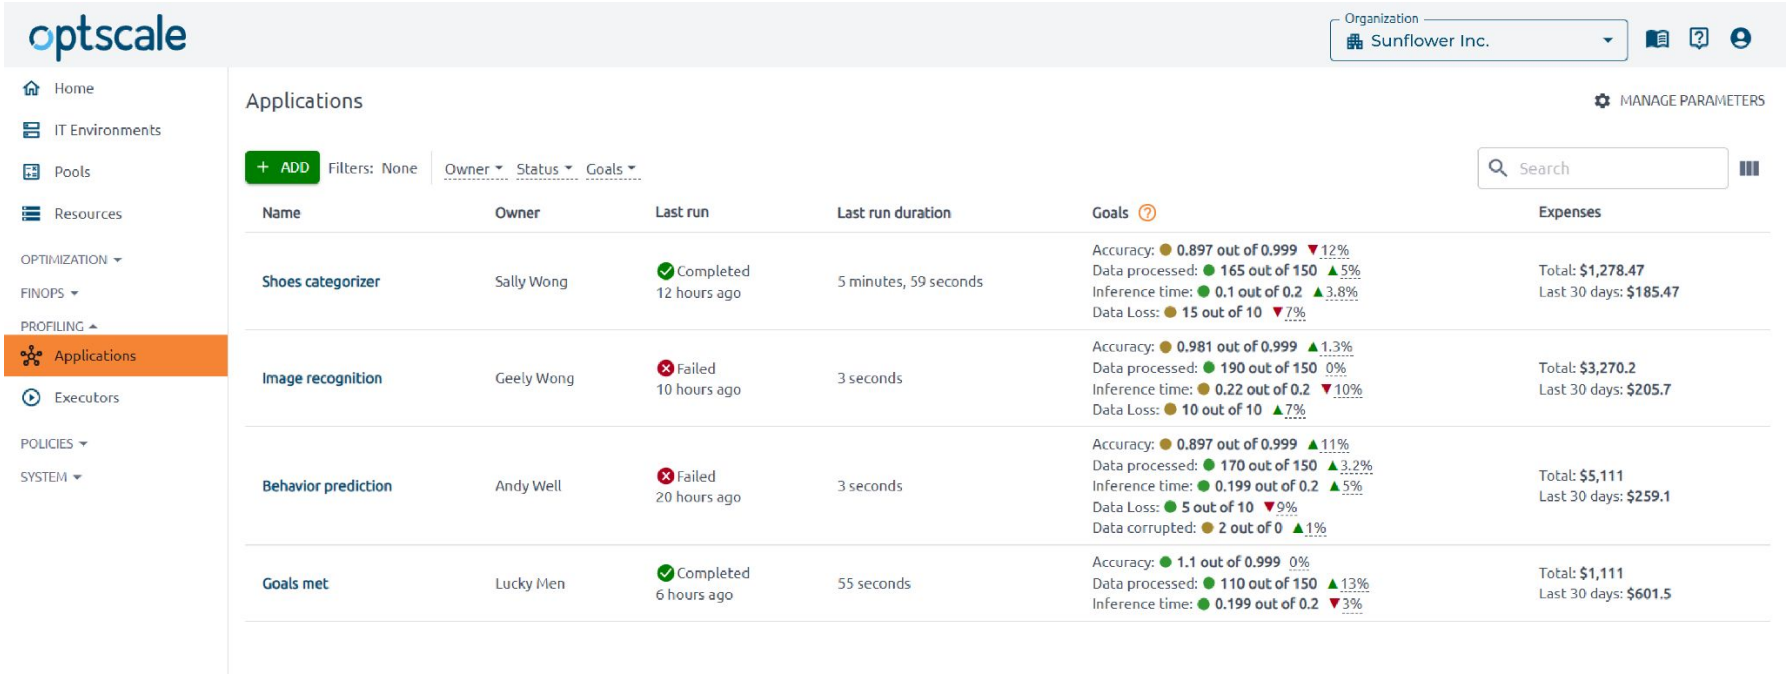

hystax

ML/AI profiling & optimization

- ML/AI model training tracking and profiling, inside and outside metrics collection
- CPU/RAM/GPU/Disk IO correlation tracking
- Minimal cloud cost for ML/AI experiments and development by utilizing Reserved Instances/Savings Plans and dozens of optimization scenarios

Supported platforms:

aws, Microsoft Azure, Google Cloud, Alibaba Cloud, kubernetes, databricks, PyTorch, Kubeflow, Apache Spark, TensorFlow

ML/AI optimization recommendations

- Utilizing Reserved/Spot Instances and Savings Plans
- Rightsizing and instance family migration
- Detecting CPU, GPU, RAM, and IO bottlenecks
- Cross-regional traffic
- Experiment/run comparison

ML/AI profiling & optimization

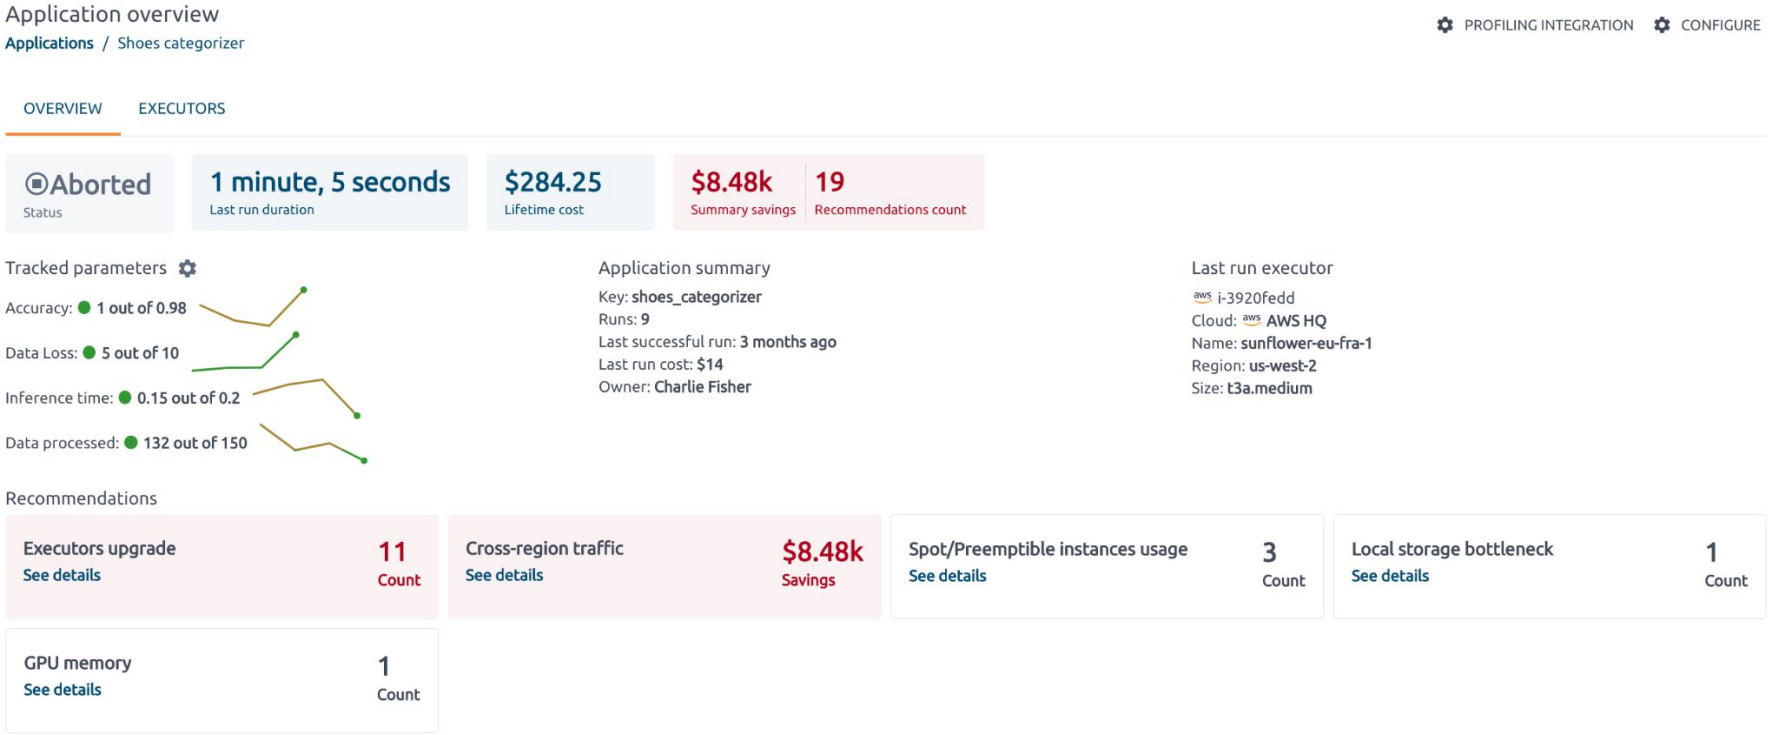

ML/AI profiling & optimization

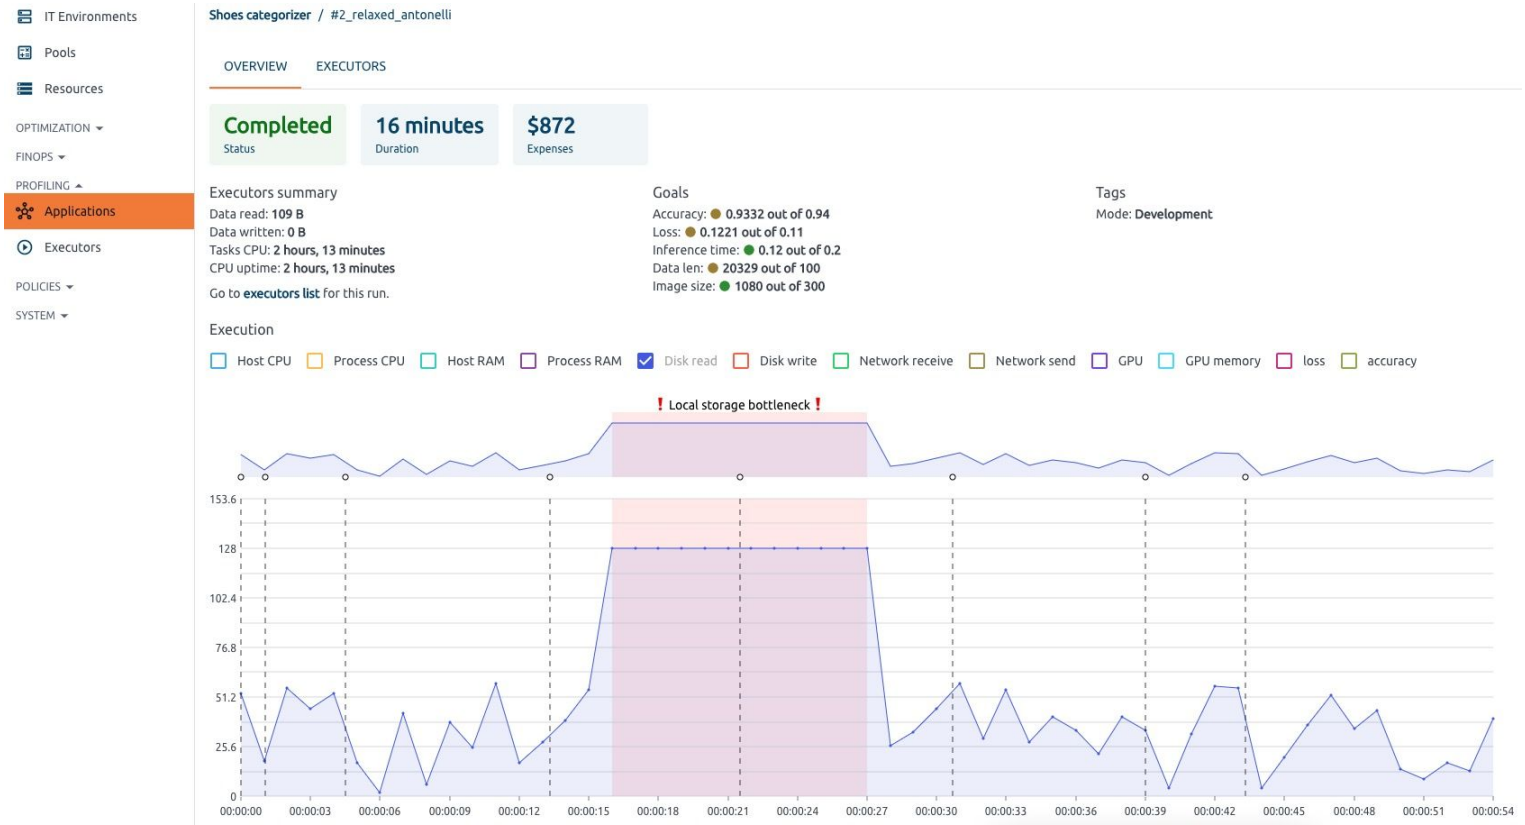

hystax

PaaS or any external service instrumentation

- Cost, performance, and output details of any API call to PaaS or an external service
- Metrics tracking and visualization
- Performance and cost optimization of API calls
- Cross-regional traffic
- S3, Redshift, BigQuery - ready, unified way to add more services

Runsets

- Automated run of a number of experiments with configurable datasets, hyperparameter ranges and model versions
- Optimal hardware with cost-efficient usage of Spot, Reserved Instances / Savings Plans
- Configurable experiment goals and success criteria
- Various complete/abort conditions - take first successful, complete all
- Integrated profiling to identify bottlenecks

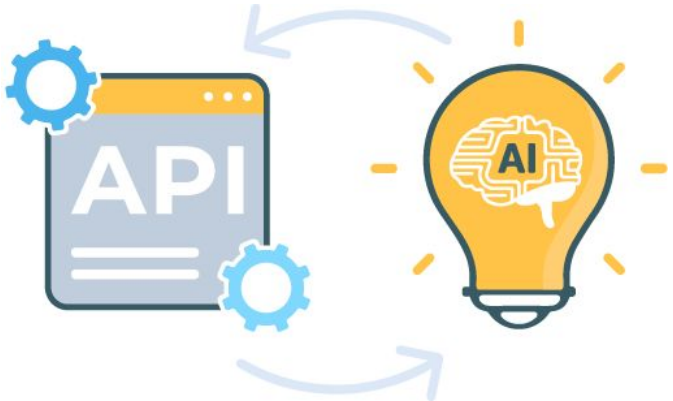

Runsets

Runset overview

AWS GPU Instances / #3_gentle_sky

6	1	$73.2
Configurations tried	Runs met goals	Total expenses

Application: **Shoes categorizer**

Parameters

Data source: aws **AWS HQ**
Region: aws **us-east-1**
Instance type: aws **t3a**
Maximum parallel runs: **14**

Hyperparameters count

Model path: **2**
Dataset path: **3**
Learning rate: **4**

Stop conditions

Stop runset when projected expenses exceeds 20$
Stop individual run if its duration exceeds 3 minutes

Correlations

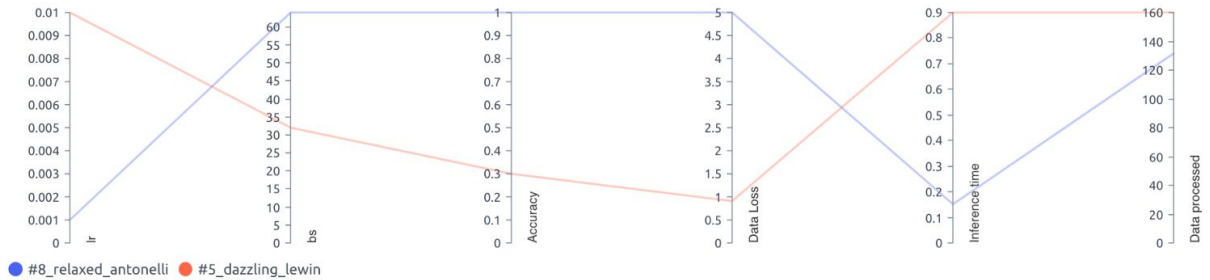

RUNS EXECUTORS

Filters: None Status Goals

Search

#	Status	Hyperparameters	Goals	Started at	Duration	Executors
#9_whispering_fog	Aborted Reached plateau for the **Accuracy** goal with the value of **0.87**	lr: **0.01** bs: **32**	Accuracy: - Data Loss: - Inference time: - Data processed: -	01/26/2023 05:55 AM	1 minute, 5 seconds	aws i-3920fedd Size: **t3a.medium** Expenses: **$14**

Roadmap

- Cost plugin for MLflow, WanDB, and neptune.ai
- Integration with Optuna to optimize Reserved Instance and other hardware parameter usage
- Model versioning
- Better hardware selection recommendations based on usage patterns and algorithms

hystax

FINOPS & CLOUD COST OPTIMIZATION

hystax

FinOps and cost management

- Forecast and monitor an IT infrastructure cost
- Identify wastage and optimize IT expenses
- Bring resource / application / service observability
- IT asset management
- Set TTL and budget constraints
- Establish a long-term FinOps process by engaging engineering teams

Supported platforms: aws Microsoft Azure Google Cloud Alibaba Cloud kubernetes

OptScale vs cloud-native cost explorer

- Cloud resource visibility and filtering across all the clouds, accounts and regions
- Dozens of optimization scenarios not supported by clouds incl. one of the best rightsizing engines
- Cost allocation not just by tags but other properties
- Geo and network traffic map
- TTL rules and budget constraints
- FinOps: OptScale is built for engineers to be responsible for their cloud resources

Contacts

+1 628 251 1280

hystax.com

info@hystax.com

1250 Borregas Avenue, Sunnyvale, CA, 94089