

Columbo: A Reasoning Framework for Kubernetes' Configuration Space

Matthijs Jansen, Sacheendra Talluri, Krijn Doekemeijer,
Nick Tehrany, Alexandru Iosup, Animesh Trivedi



@Large Research
Massivizing Computer Systems



m.s.jansen@vu.nl



<https://atlarge-research.com/mjansen/>



VRIJE
UNIVERSITEIT
AMSTERDAM



Columbo
Open-source Code


Containerization

Container benefits:

- Virtualization with **little overhead**
- **Portable** and **reproducible** packaging and execution


In data centers:

- Many containers across many machines
- **Resource managers** govern life cycle



CNCF 2024 Annual Survey

How are containers used in your organization



The new normal

Over 90% uses containers

The production standard


80% uses Kubernetes

Does your organization use Kubernetes

Yes, in production

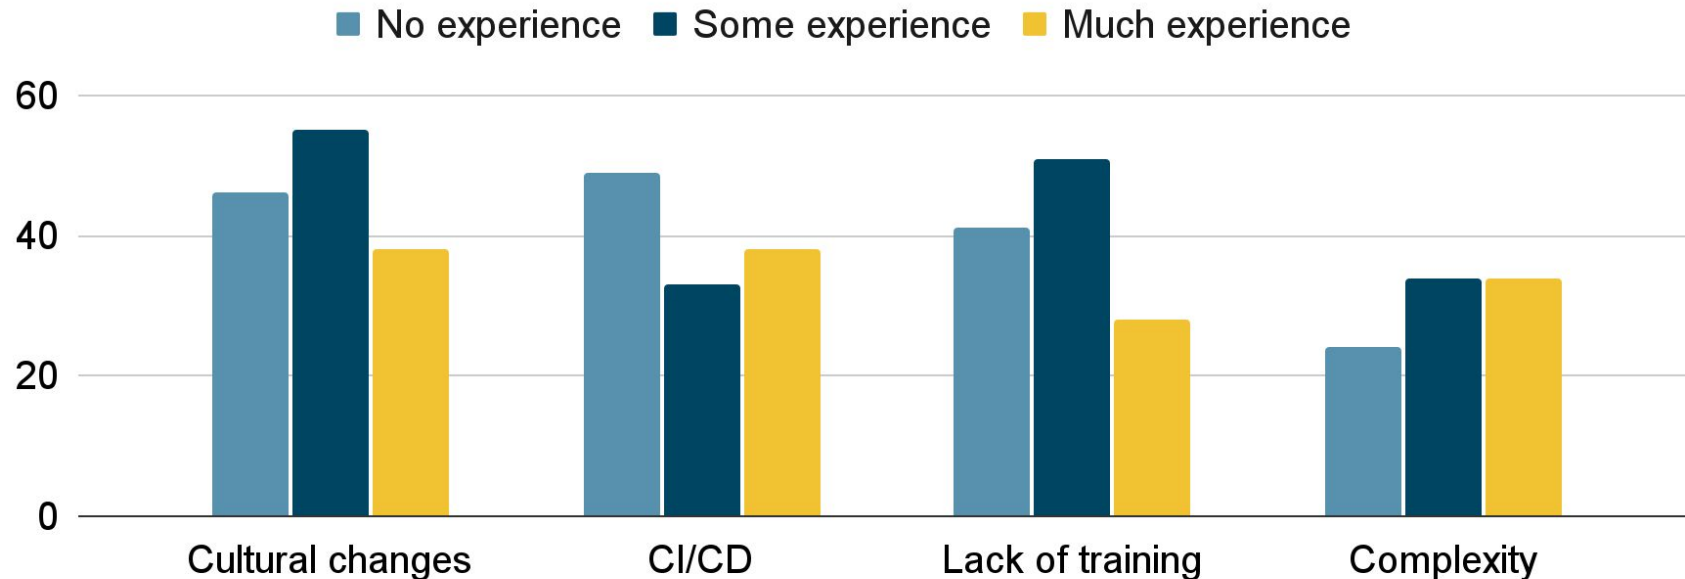
Yes, piloting

No

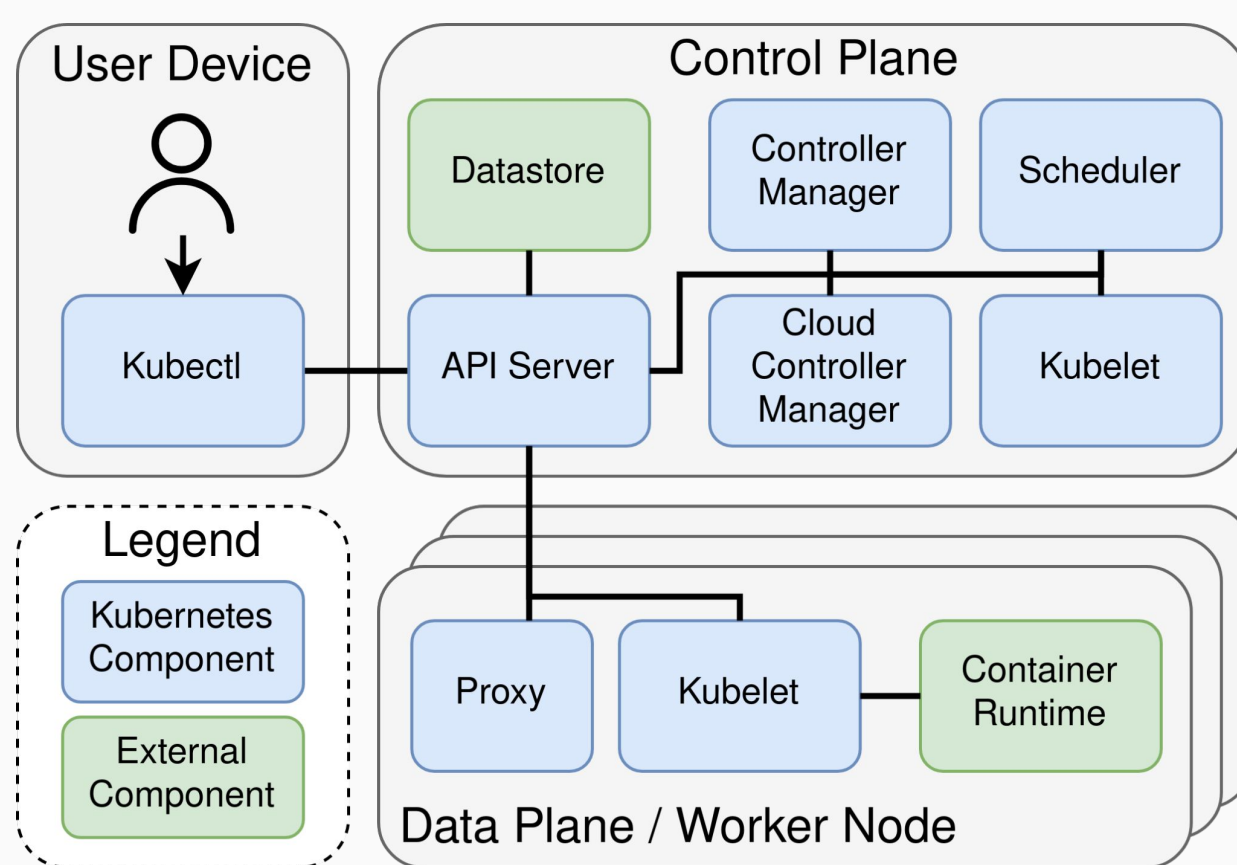


But: Containers are **difficult to use**

Container use challenges by container maturity level



Kubernetes Configuration Complexity



Architecture with 9 distributed components

- **Control plane:**
Decision making
- **Data plane:**
Decision execution


Configuration API:

- **234 resources**
- **1598 parameters**

Configurations Are Important

Goal: Tune configurations to optimize container deployment latency


- **Google Kubernetes Engine** deploys only **7% faster**
- Configuration tuning with **our approach**: Deploys **37% faster**



Configuration Tuning Strategy

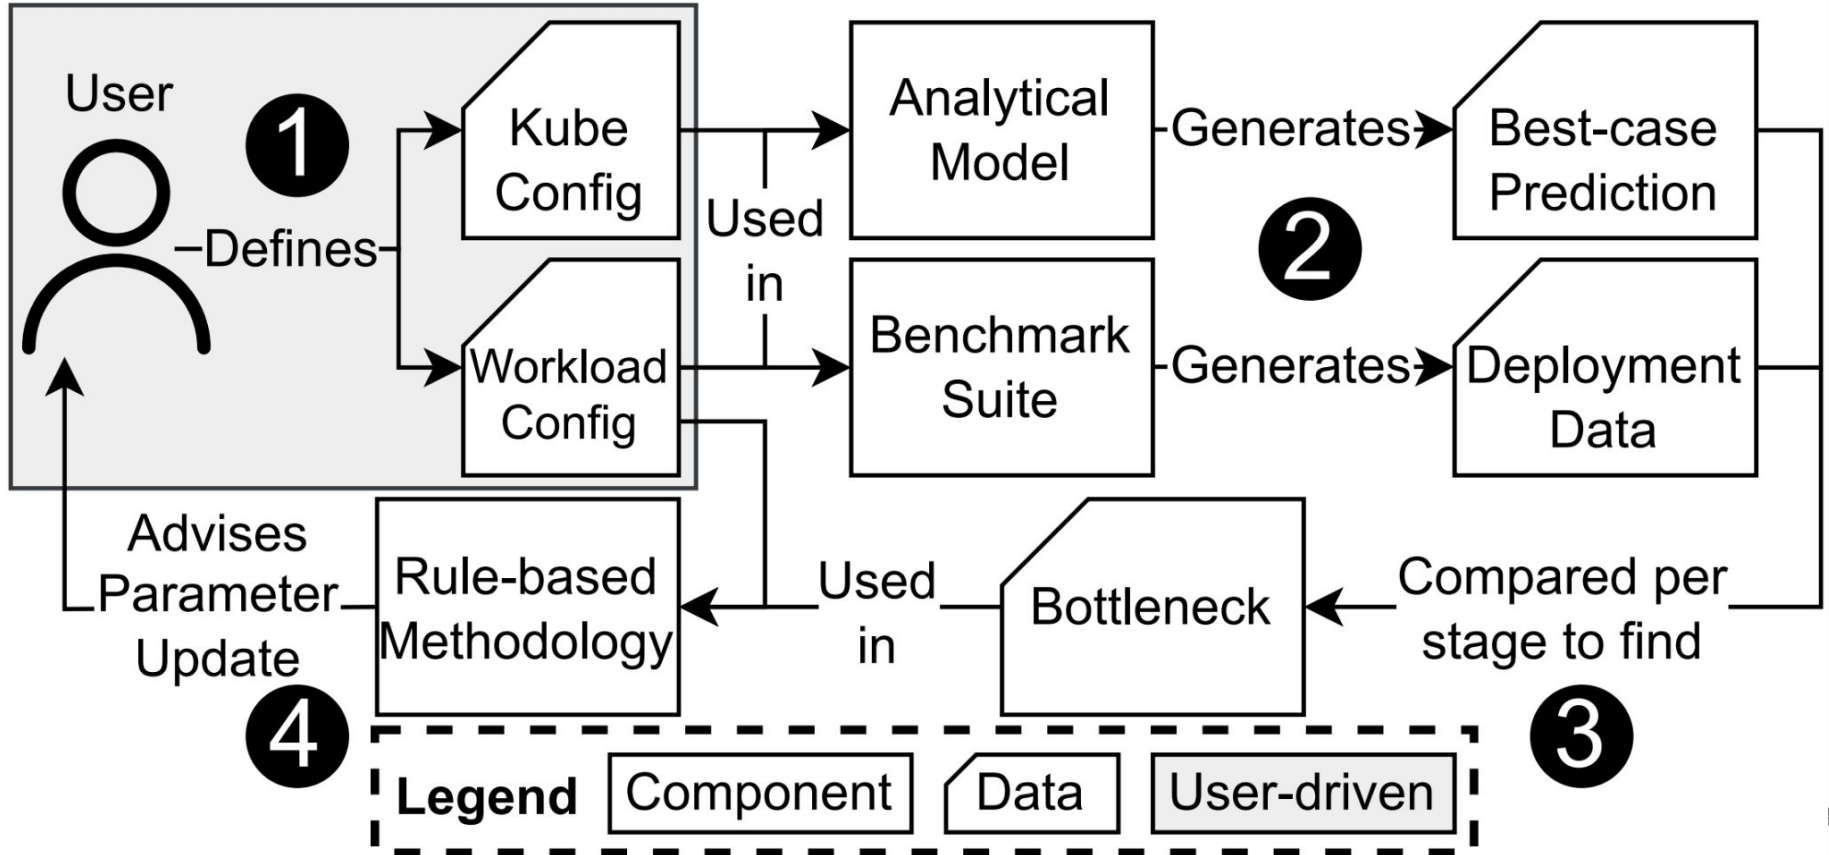
Approaches

- **Brute force? No** → Too many parameters, slow updates
- **Users to restrict parameter space? No** → Requires expertise
- **Root-cause analysis? Yes** → Fast and user-friendly









Pipeline Stage	Parameter	Object	Current value	Demanded
3	<i>Batch Size</i>	Pods	5	100

Columbo's Configuration Tuning



Workload Deployment Pipeline

6 stages: Unique data object / components control loop
26 steps: Move data between components




Best case: Objects are **independent** and are processed **in parallel**
So: Benchmark execution of **1 container**, **extrapolate** to demand

$$T_{s1} = T_{start-kubectl} + T_{translate} + Latency + T_{write}$$

$$T_{total} = T_{s1} \times [jobs/CPU s]$$

Stage 1: Create Workload Object

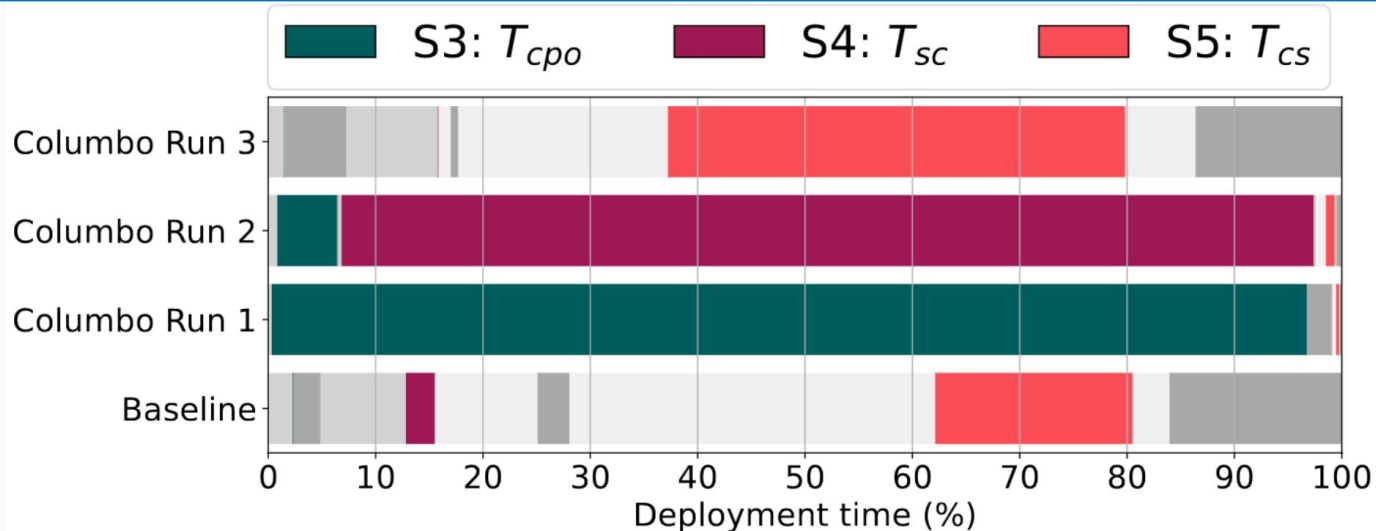


Columbo Day-to-day Operation

Use case: 1600 containers over 16 worker, 1 control plane node

Bottleneck detection:

- Benchmark finds **97% execution time in phase 3, step 9**
- Analytical baseline predicts **<1% of execution time in this step**



Resolve Bottleneck with Parameter Rules

1. Map configuration parameter to pipeline stages / steps
 - a. Automatic analysis of documentation and source code
 - b. Filter for performance-sensitive parameters
2. Define parameters' current value and matching object
 - a. Extract from documentation / source code and pipeline

One-time effort → Expert can augment this automated approach

Stage	Step	Parameter	Current Value	Scales with
3	8	<u><i>InitialBatchSize</i></u>	<u>1</u>	#pods
4	14	<i>KubeApiQPS</i>	50	#pods

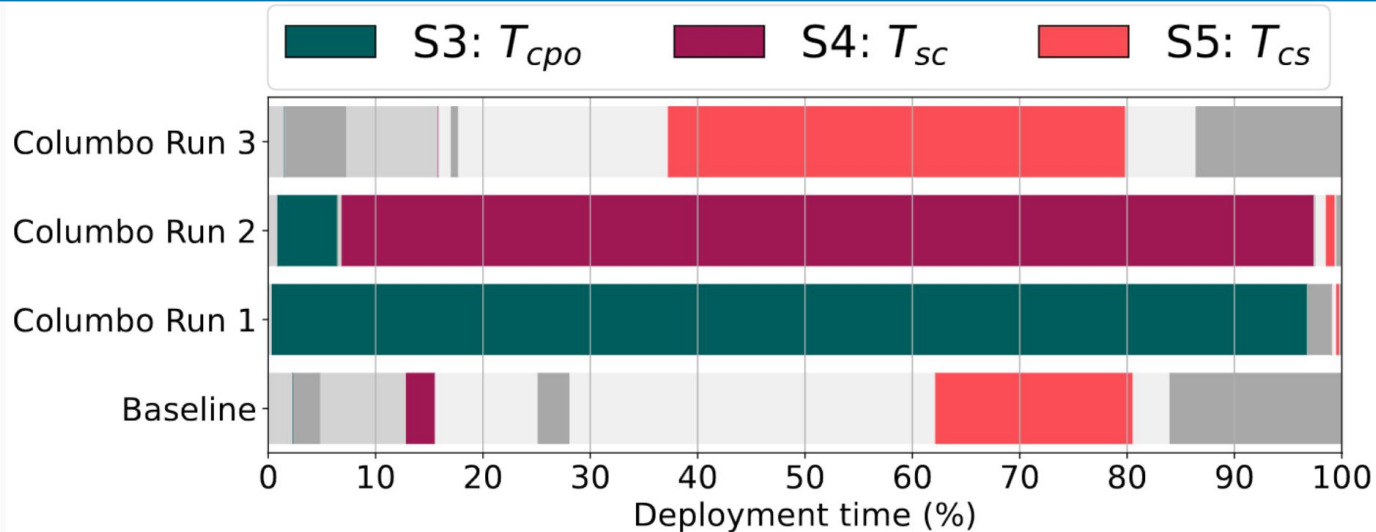
**We have
1600 pods**

Columbo Day-to-day Operation

Update *InitialBatchSize* to 1600 to resolve bottleneck

Columbo Run 2: 80.9 seconds to 31.1 seconds

- Benchmark finds 97% execution time in phase 4, step 14
- Analytical baseline predicts ~6% of execution time in this step




Columbo Day-to-day Operation

Update *KubeApiQPS* to 1600 to resolve bottleneck

Columbo Run 3: 80.9 seconds to 16.9 seconds → 79.1% reduction

- CPU utilization at 100%
- No further optimization possible unless resources are added




Optimize Deployment Method

100 containers over:

- 100 kubectl calls
- 1 call, 100 jobs
- 1 job, 100 pods
- 1 pod, 100 cont.

Results:

- Cont.: **-16%**
- Pod: **+1%**
- Job: **+5%**
- Call: **-37%**



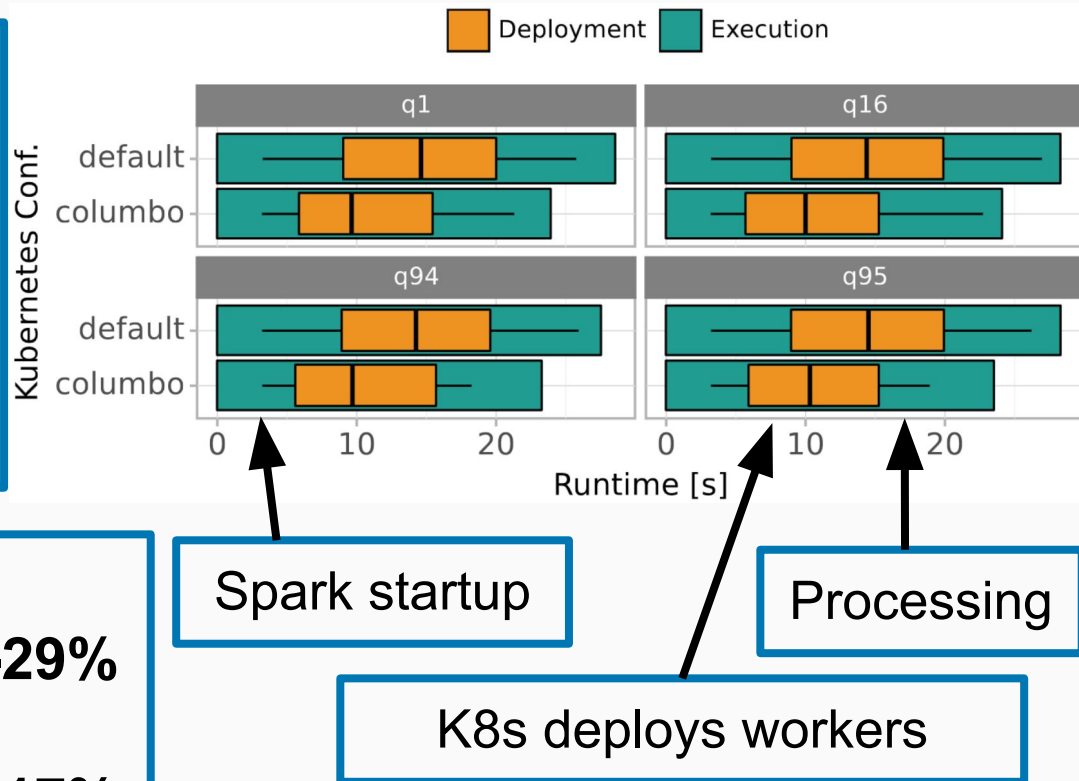
End-to-end Impact of Columbo

Setup:

- 100GB TPC-DS on Spark
- Spark control plane and workers on Kubernetes
- 400 Spark workers run 4 queries with cold starts

Findings:

- **Avg. deployment time: -29%**
 - Similar to microbench.
- **Avg. end-to-end time: -17%**

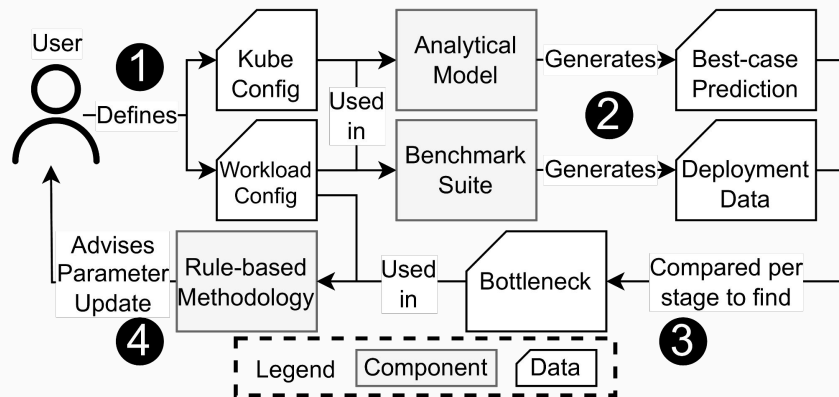


Conclusion

We present Columbo:

- Automatically detect and resolve bottlenecks for Kubernetes configurations
- Columbo reduces deployment time (avg. 28%) and total execution time (17% for Spark)

Open-source and free-to-use



m.s.jansen@vu.nl



atlarge.science/mjansen



Columbo
Open-source Code