

A Dynamic Kubernetes Load Generation Solution Mimicking Human Traffic Patterns

CLOUDSTARS secondment experience

Ranjan Ojha
June 19


WOSCx3



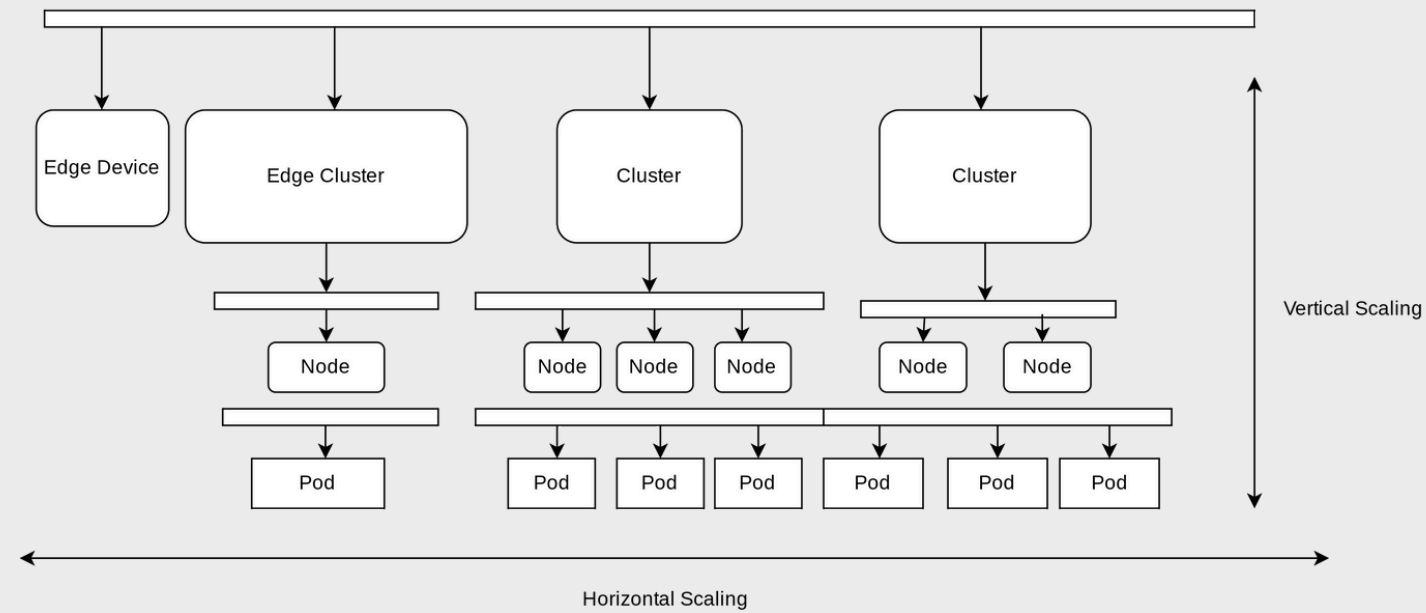


Kubernetes


The Operating system of Cloud




Autoscalers




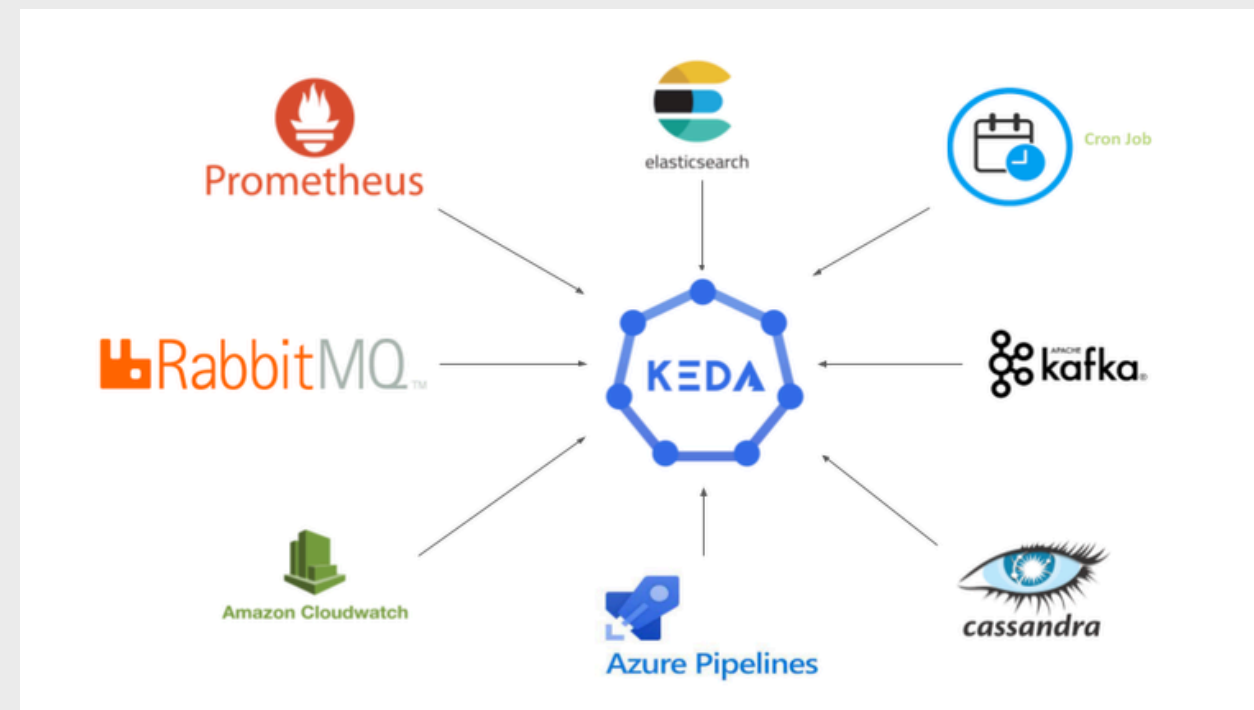
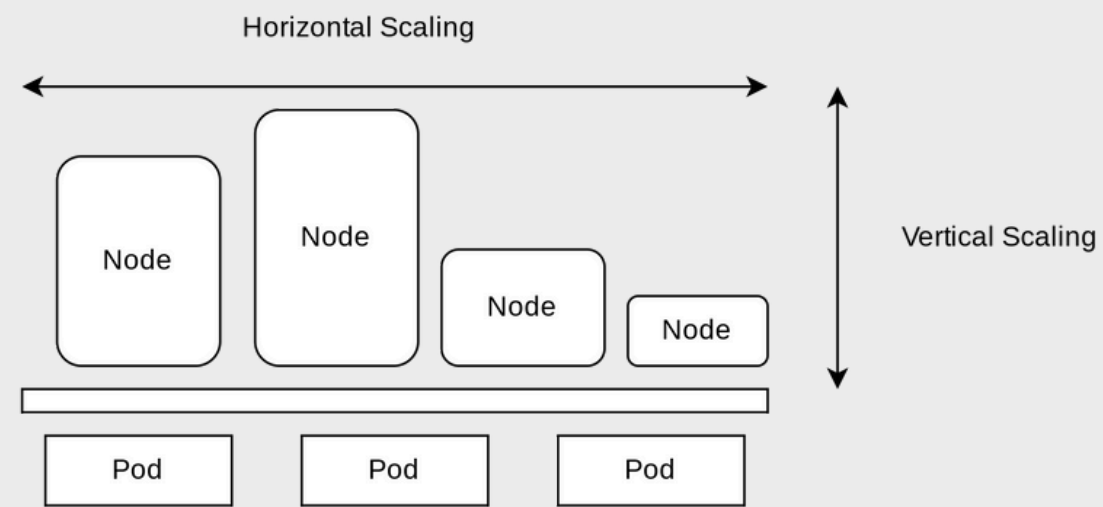
HPA & VPA



Traditional Autoscaling solution



Cluster Autoscaler



Alternatives?

Intelligent Autoscaling of Microservices in the Cloud for Real-Time Applications

ABEER ABDEL KHALEQ, (Member, IEEE), AND ILKYEUN RA^{ORCID}, (Member, IEEE)

Department of Computer Science and Engineering, University of Colorado at Denver, Denver, CO 80204, USA

Corresponding author: Ilkyeun Ra (ilkyeun.ra@ucdenver.edu)

Proactive Random-Forest Autoscaler for Microservice Resource Allocation

LAMEES M. AL QASSEM, THANOS STOURAITIS, ERNESTO DAMIANI, AND IBRAHIM (ABE) M. ELFADEL

Center for Cyber Physical Systems (C2PS) and Department of Electrical Engineering and Computer Science

Khalifa University, P.O. Box 127788, Abu Dhabi, UAE

E-mail: {lamees.alqassem, thanos.stouraitis, ernesto.damiani, ibrahim.elfadel}@ku.ac.ae

Corresponding author: Ibrahim (Abe) M. Elfadel (e-mail: ibrahim.elfadel@ku.ac.ae).

Proactive Autoscaling for Edge Computing Systems with Kubernetes

Li Ju

li.ju@it.uu.se

Dept. of Information Technology
Uppsala University
Uppsala, Sweden

Prashant Singh

prashant.singh@it.uu.se

Science for Life Laboratory
Dept. of Information Technology
Uppsala University
Uppsala, Sweden

Salman Toor

salman.toor@it.uu.se

Dept. of Information Technology
Uppsala University
Uppsala, Sweden

Most of the algorithms, are an improvement over HPA so suffer from being reactive rather than proactive.

They don't have the reproducibility to compare with existing solutions.

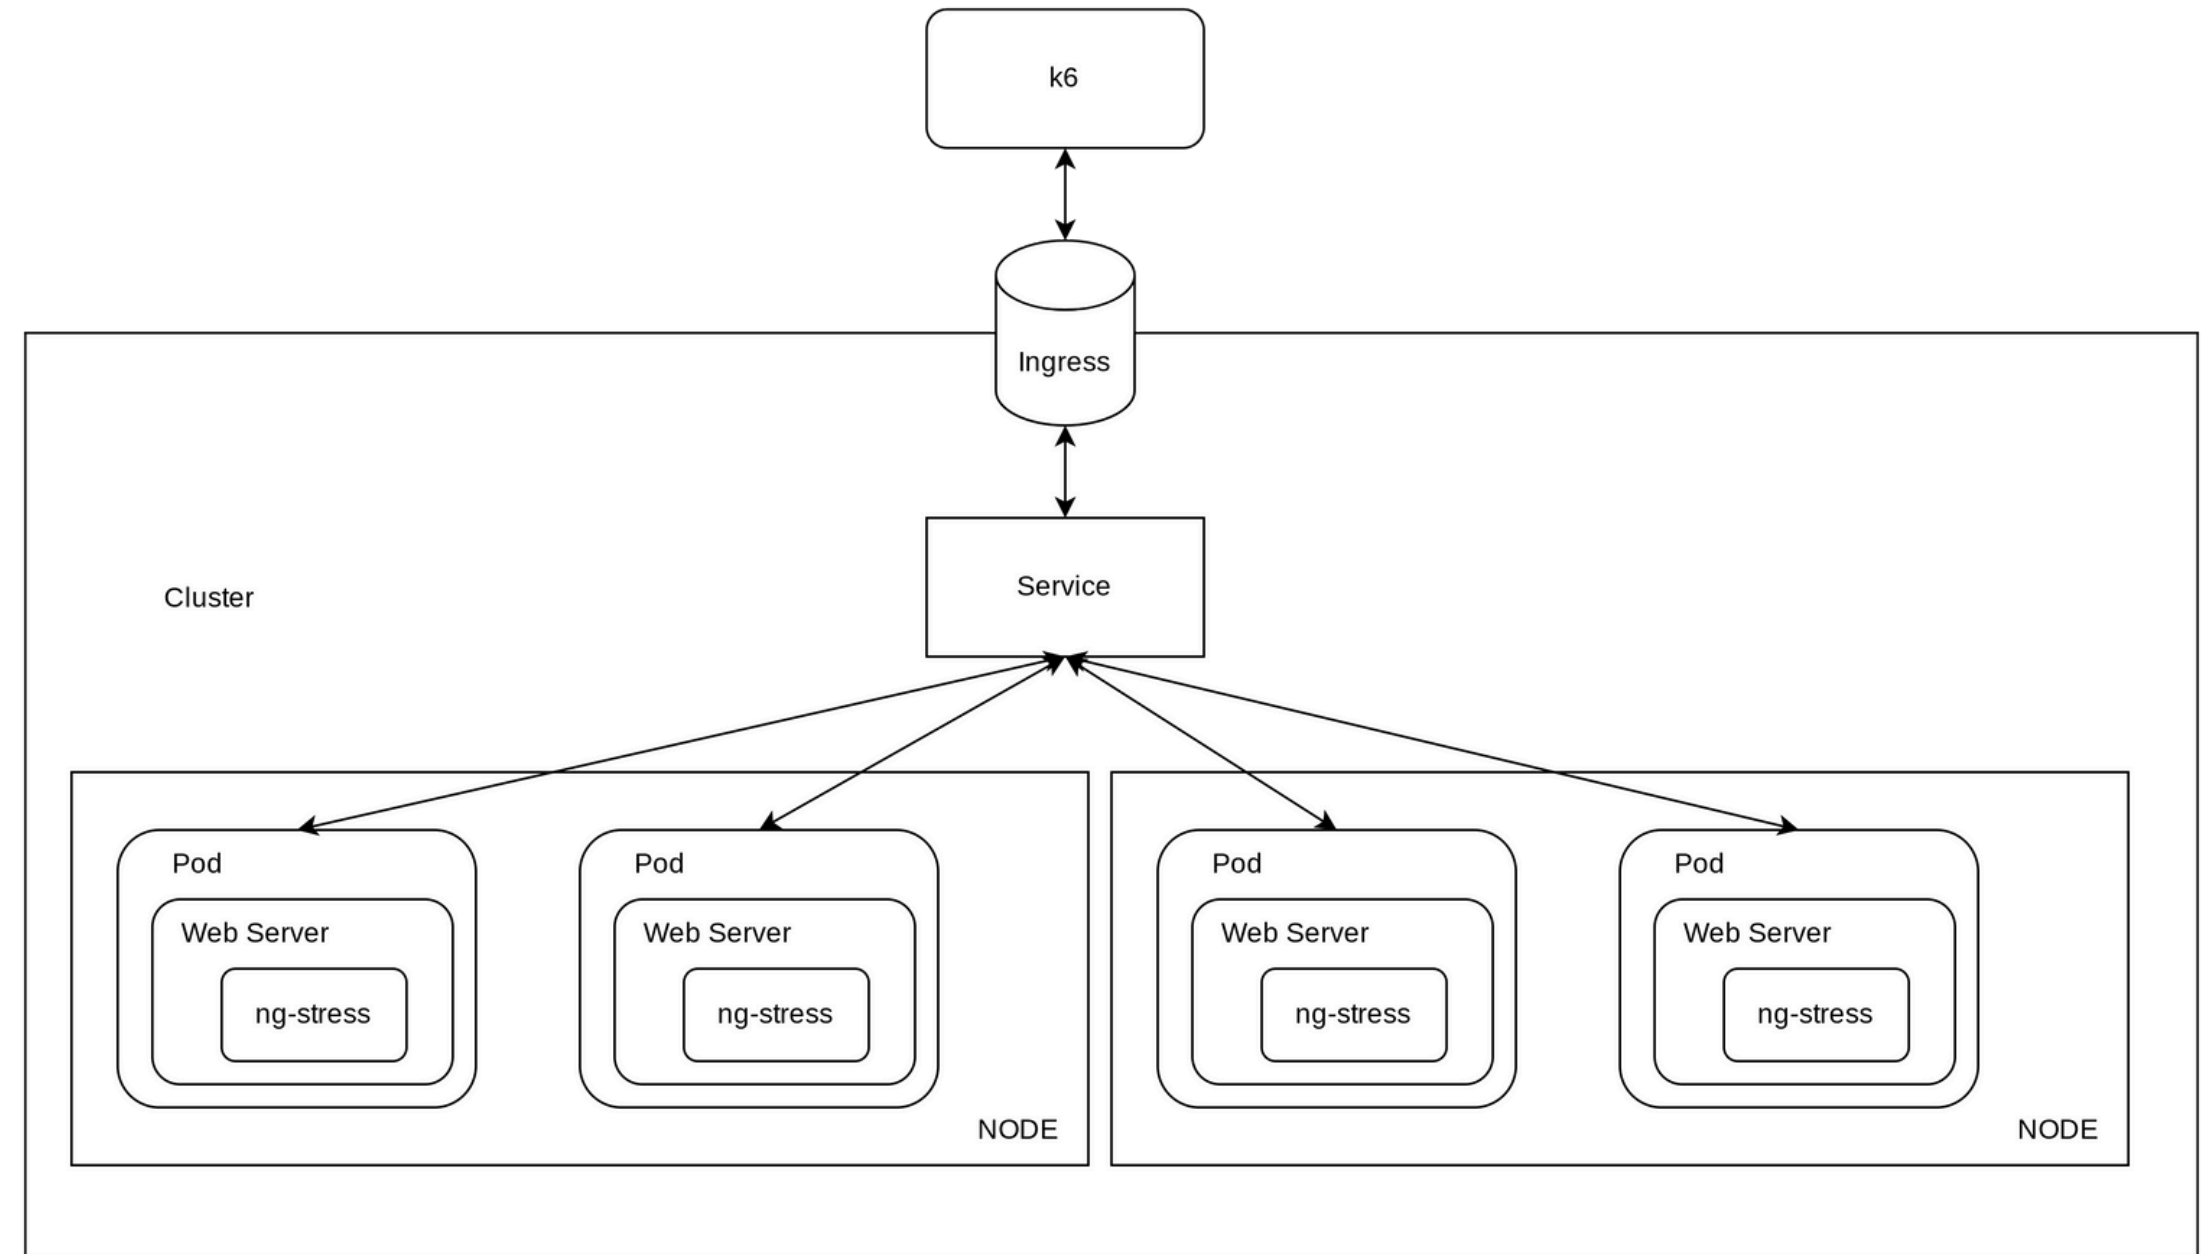
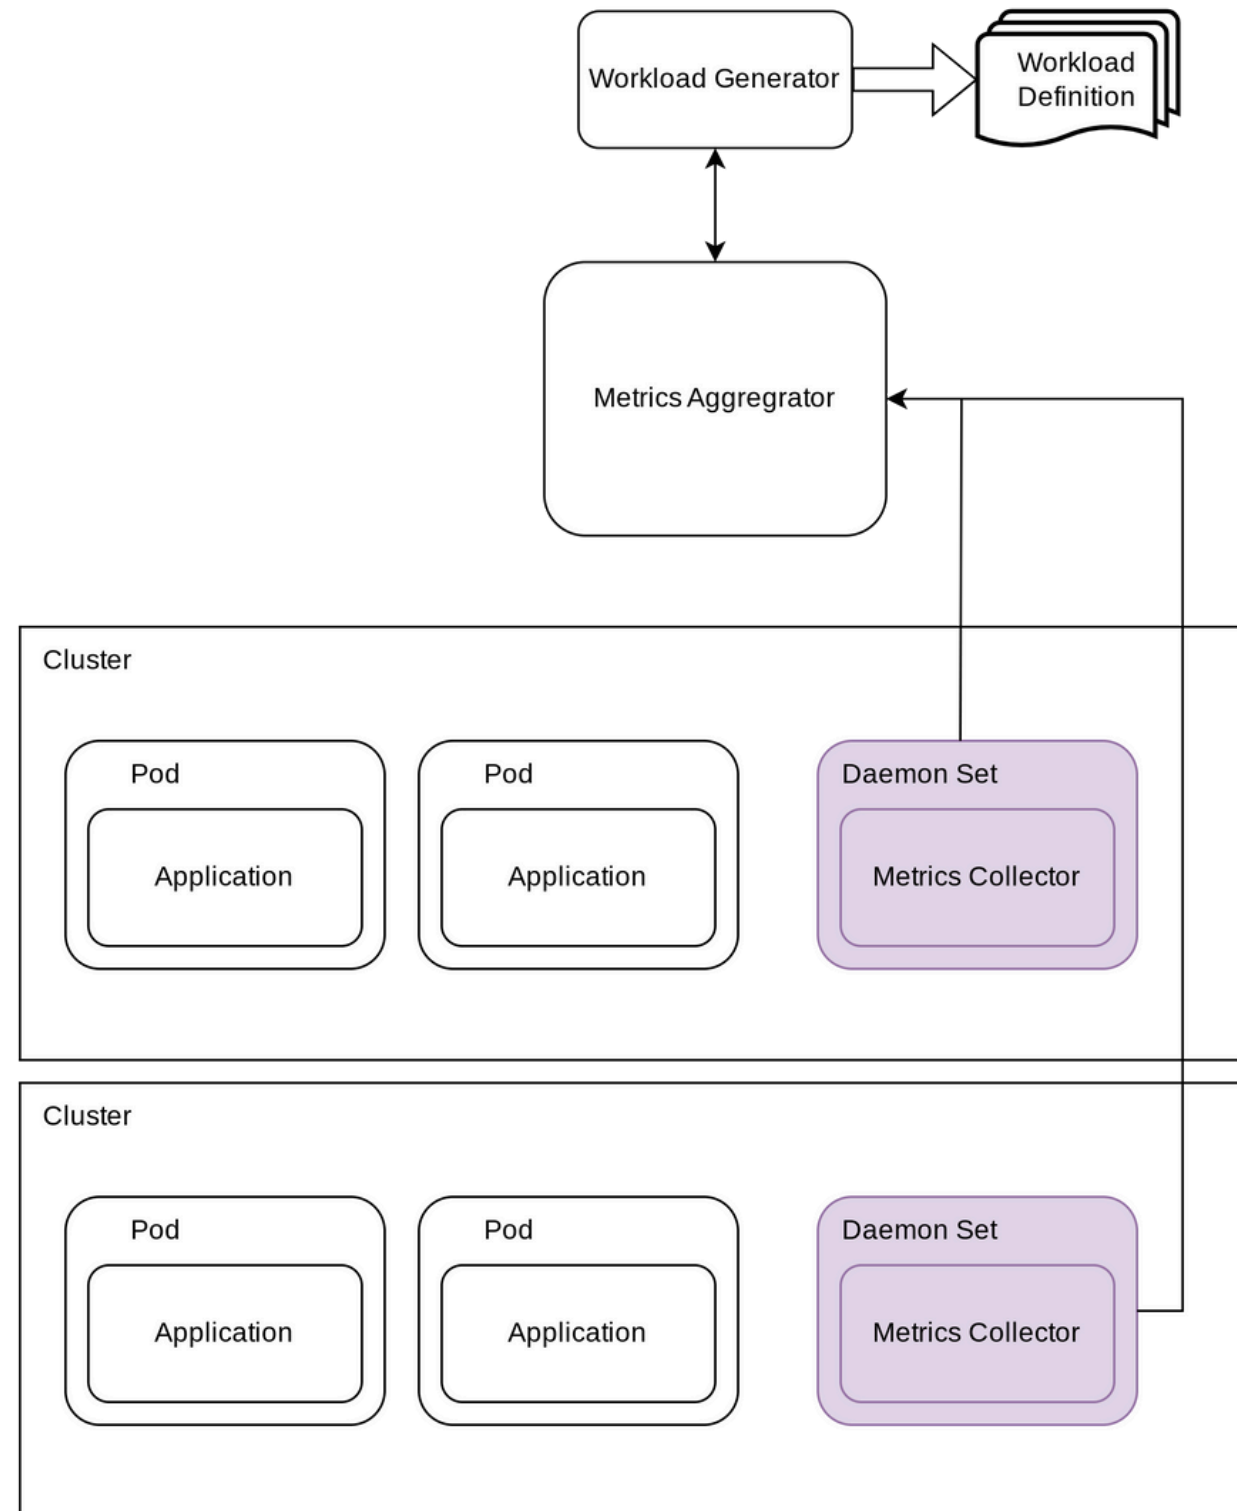
Autoscalers don't adapt to the application, and require lots of manual tuning, the time and resource not available for PaaS providers.

Problem statement

- Most of the current autoscaling algorithms are reactive. They recommend having overhead of resources in the system. For the system to grow into while the surge of traffic causes autoscaling to kick in start adding in resources for the system to utilize.
- Autoscaling solutions of today require that operators of the cluster are well aware about the behavior of application. This however is not always possible. PaaS providers for instance have an SLA to meet and aren't aware of the application behavior or are aware about the general trends. This means they can't prepare in advance. All the scaling decisions must be made in the moment. How can we guarantee that a particular system will allow providers to meet the SLA ?

First Step

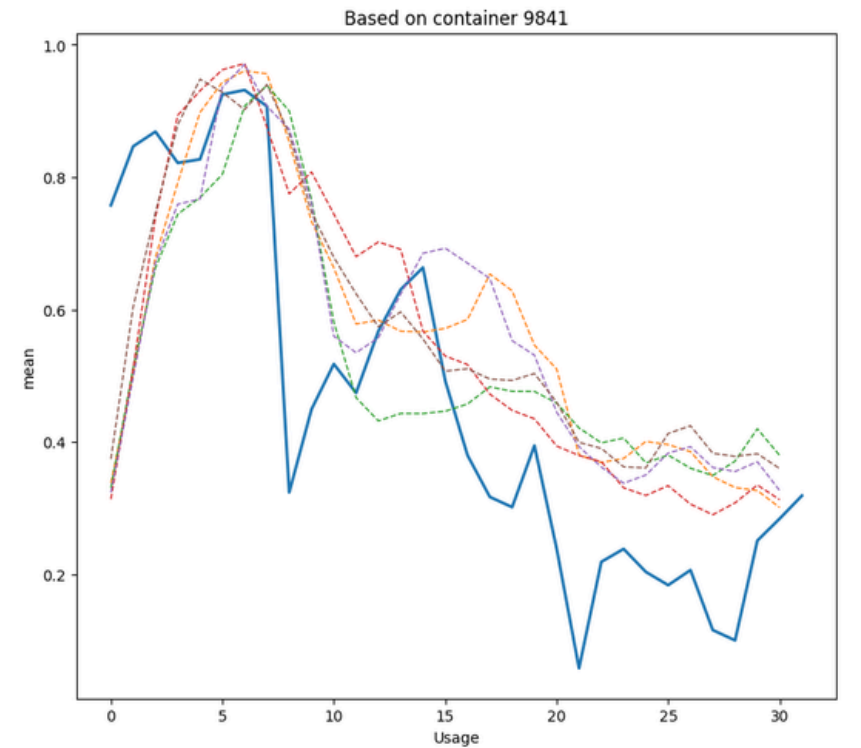





New Tool Design for mimicking human interaction with Kubernetes hosted applications



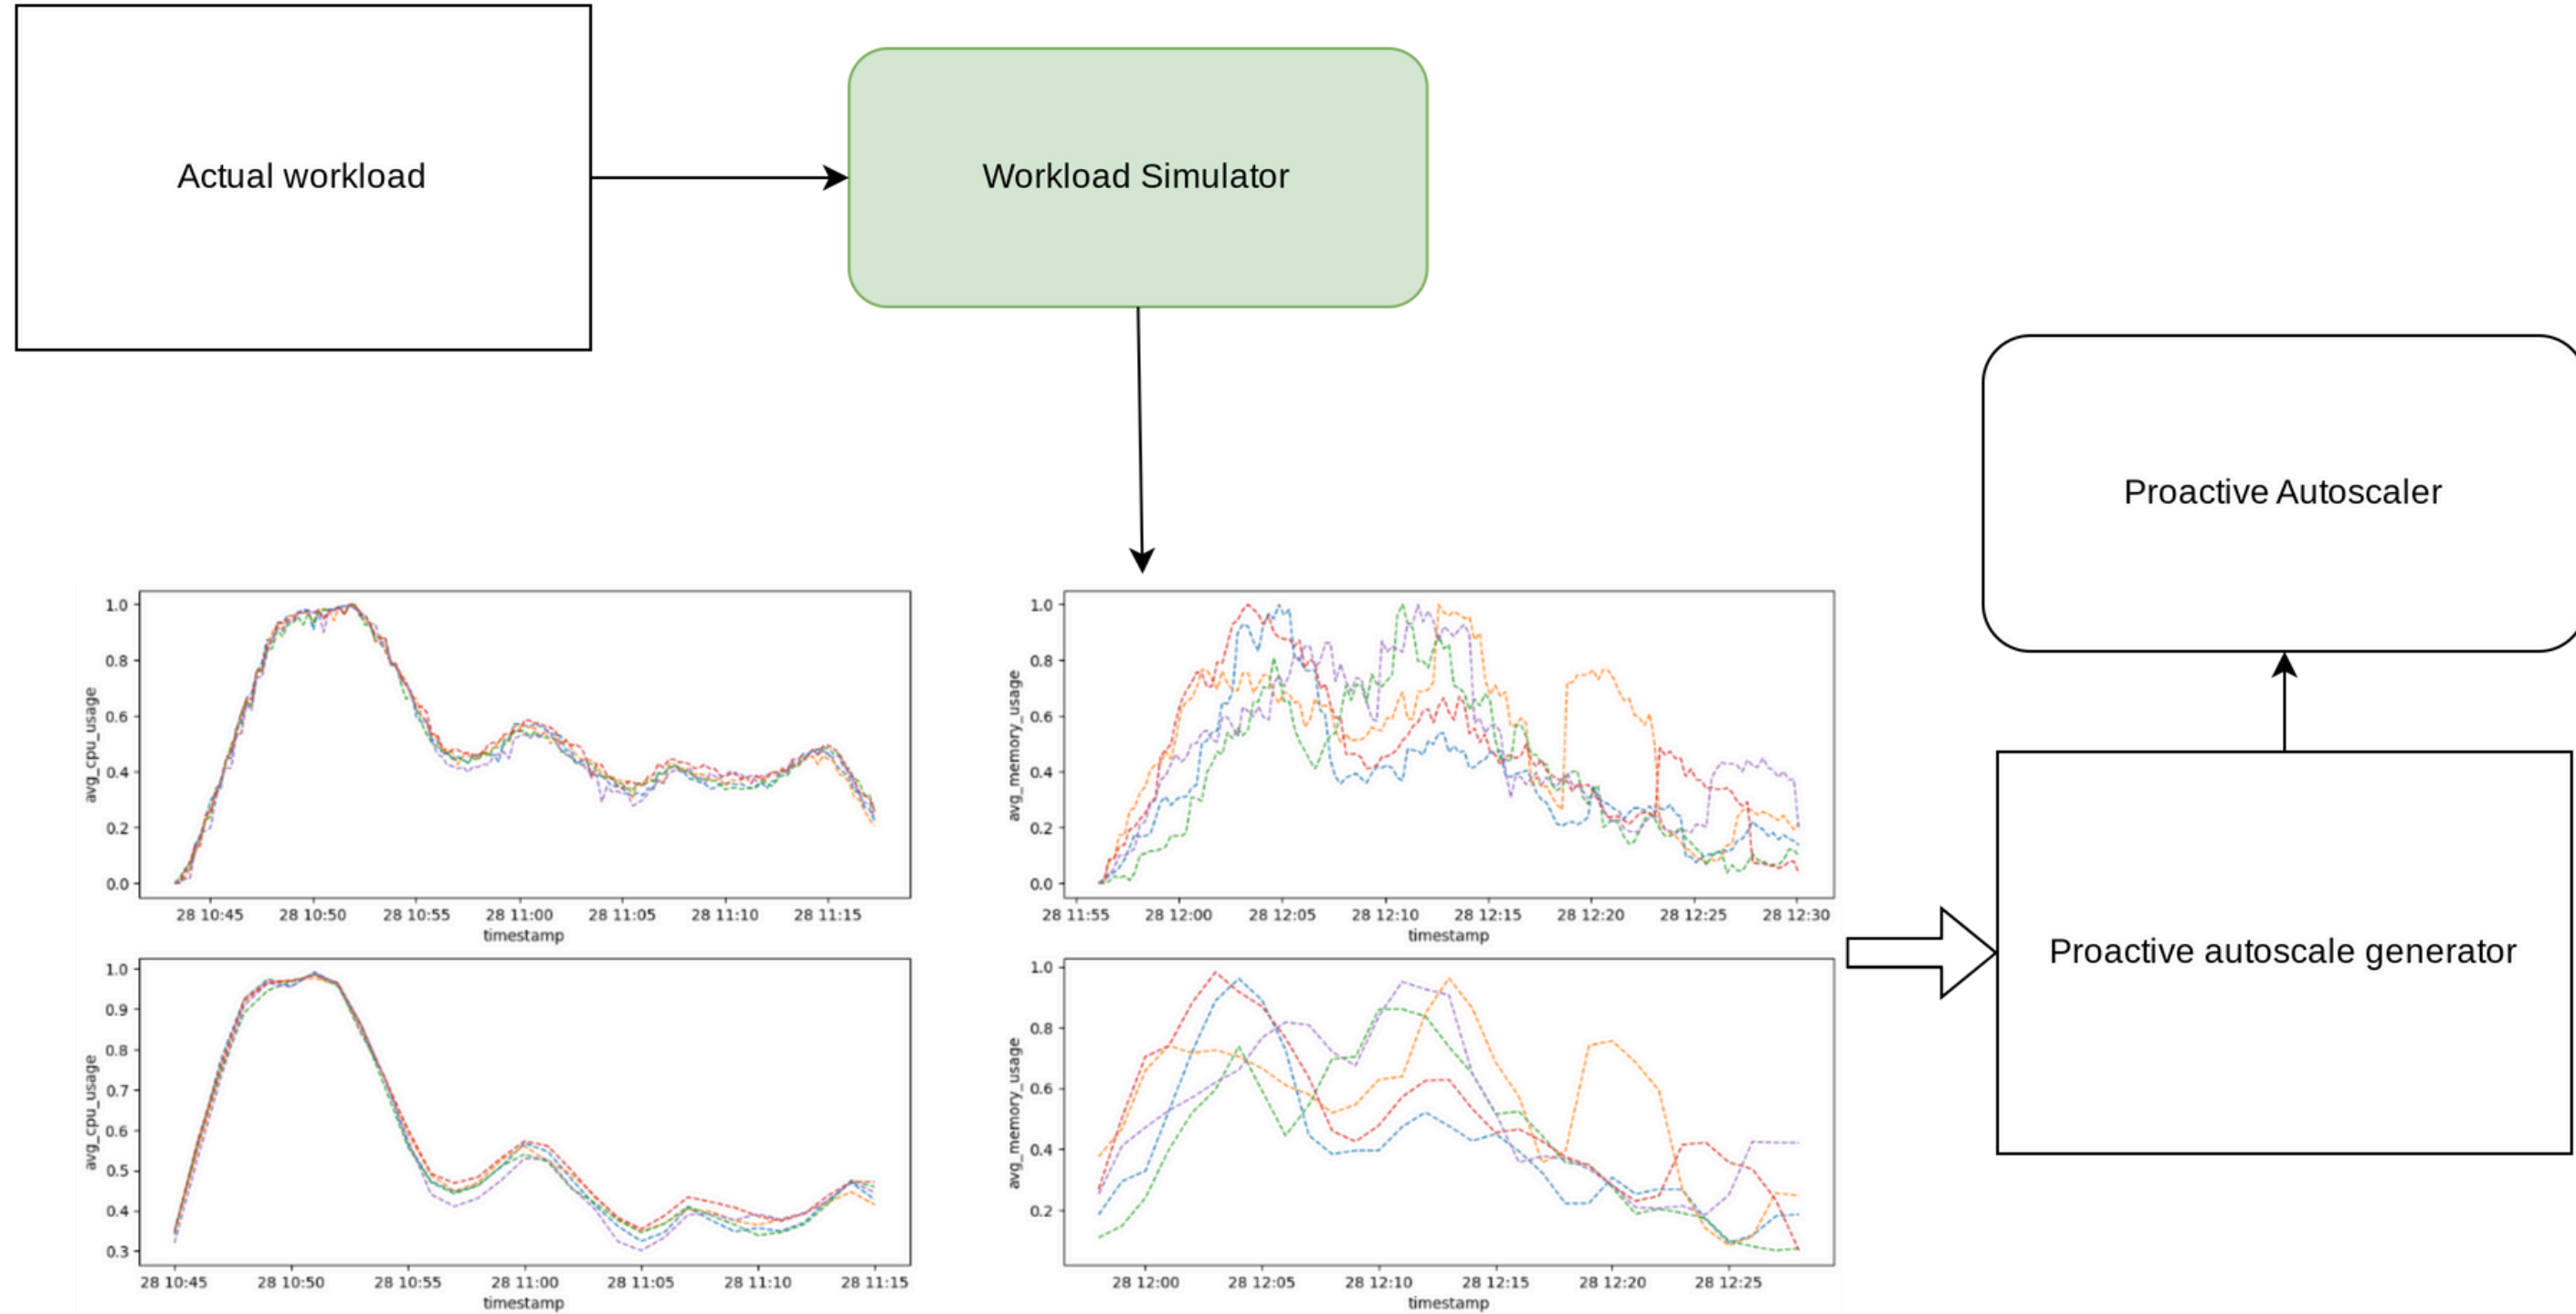
Performance

collection_id	avg_cpu_usage	max_cpu_usage	avg_memory_usage	max_memory_usage	avg_cpi	avg_mai
9453240160	9.58442687988281E-05	0.000104904174804688	0.000388622283935547	0.000426292419433594	2.13405412435532	0.0138982711359859
9453240160	9.48905944824219E-05	0.000102043151855469	0.00038814544677344	0.000425338745117188	2.2089404463768	0.0143667533993721
9453240160	9.58442687988281E-05	0.000103950500488281	0.00038909912109375	0.000426292419433594	2.50274127721787	0.0162047441117466
9453240160	9.67979431152344E-05	0.000107765197753906	0.000389575958251953	0.000426292419433594	2.50274127721787	0.0162047441117466
9453240160	9.67979431152344E-05	0.000105857849121094	0.000388622283935547	0.000424385070800781	2.50274127721787	0.0162047441117466
9453240160	9.5367431640625E-05	0.000104904174804688	0.000387668609619141	0.000422477722167969	2.53739476203918	0.01738870077591
9453240160	9.34600830078125E-05	0.000102996826171875	0.000387191772460938	0.000423431396484375	2.70294910669327	0.0190583779823
9453240160	9.39369201660156E-05	0.000103950500488281	0.00038814544677344	0.000423431396484375	2.70294910669327	0.0190583779823
9453240160	9.67979431152344E-05	0.000105857849121094	0.000386714935302734	0.000421524047851563	2.70294910669327	0.0190583779823
9453240160	9.44137573242188E-05	0.000102043151855469	0.000385284423828125	0.0004205703777156	2.70294910669327	0.0190583779823
9453240160	9.44137573242188E-05	0.000101089477539063	0.00038623809814531	0.00041961669921875	2.70294910669327	0.0190583779823
9453240160	9.5367431640625E-05	0.000101089477539063	0.00038623809814531	0.00041961669921875	2.70294910669327	0.0190583779823
9453240160	9.77516174316406E-05	0.000103950500488281	0.0003870758619922	0.00041961669921875	2.71121311187744	0.0164402015507221
9453240160	9.87052917480469E-05	0.00010681152344E-05	0.00038623809814531	0.000418663024902344	2.71121311187744	0.0164402015507221
9453240160	9.5367431640625E-05	0.000103950500488281	0.000387191772460938	0.000421524047851563	2.5532229862213	0.0168146705254912
9453240160	9.63211152344E-05	0.000103950500488281	0.000387191772460938	0.000422477722167969	2.52237248420715	0.0168877900578082
9453240160	9.918212890625E-05	0.000103950500488281	0.00038909912109375	0.000421524047851563	2.52237248420715	0.0168877900578082
9453240160	9.58442687988281E-05	0.000103950500488281	0.000386714935302734	0.00041961669921875	2.55864500999451	0.0172922308556736
9453240160	9.77516174316406E-05	0.000111579895019531	0.000386714935302734	0.000422477722167969	2.90861594676971	0.0185499745421112
9453240160	9.5367431640625E-05	0.000107765197753906	0.00038814544677344	0.00041961669921875	3.06957566738129	0.0188797526061535
9453240160	9.44137573242188E-05	0.000103950500488281	0.000387191772460938	0.00041961669921875	3.06957566738129	0.0188797526061535
9453240160	9.918212890625E-05	0.000109672546386719	0.000385761260986328	0.000417709350585938	3.06957566738129	0.0188797526061535
9453240160	9.48905944824219E-05	0.000102996826171875	0.000386714935302734	0.000418663024902344	3.0213440656662	0.018806538078934
9453240160	9.39369201660156E-05	0.000102996826171875	0.00038623809814531	0.00041961669921875	2.95778286457062	0.0187772121280432
9453240160	9.20295715332031E-05	0.00010013580322656	0.000386714935302734	0.000418663024902344	2.88512849807739	0.0189613737165928
9453240160	9.34600830078125E-05	0.000102996826171875	0.000385284423828125	0.000418663024902344	2.88512849807739	0.0189613737165928
9453240160	9.25064086914063E-05	0.000102996826171875	0.00038623809814531	0.00041961669921875	2.93886077404022	0.0190484225749969
9453240160	0.00177658928765191	0.00811767578125	0.000308036804199219	0.000432014465332031	3.02831137180328	0.019193337764591

Input Dataset from Google



Moving Onwards



Thank You

WOSCx3

