

# Predictable elasticity of Docker applications

Manuel Ramírez López

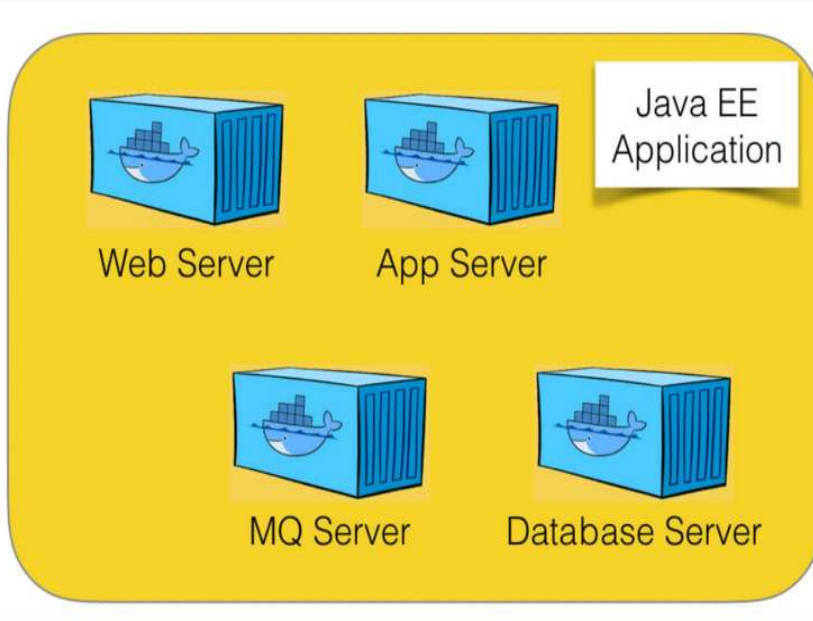
Service Prototyping Lab. ZHAW

ramz@zhaw.ch

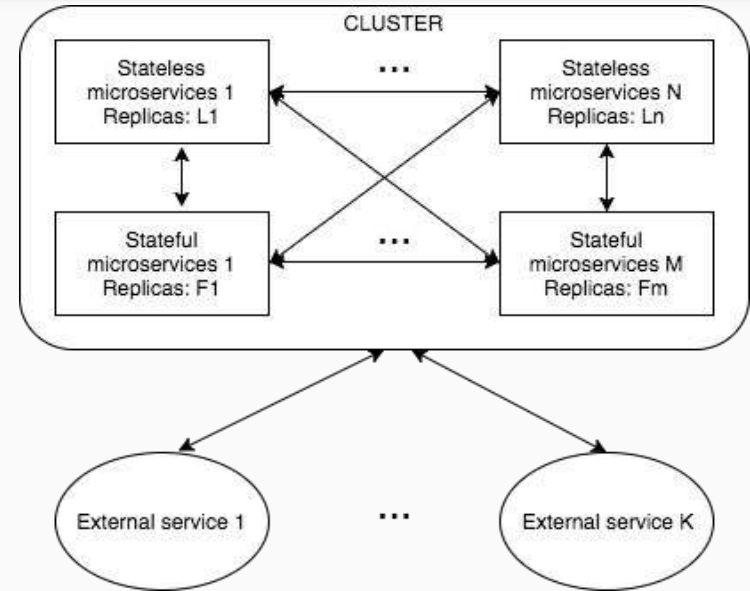
31/05/2017

14th Docker Switzerland User Group Meetup

# Application composed of multiple Docker containers

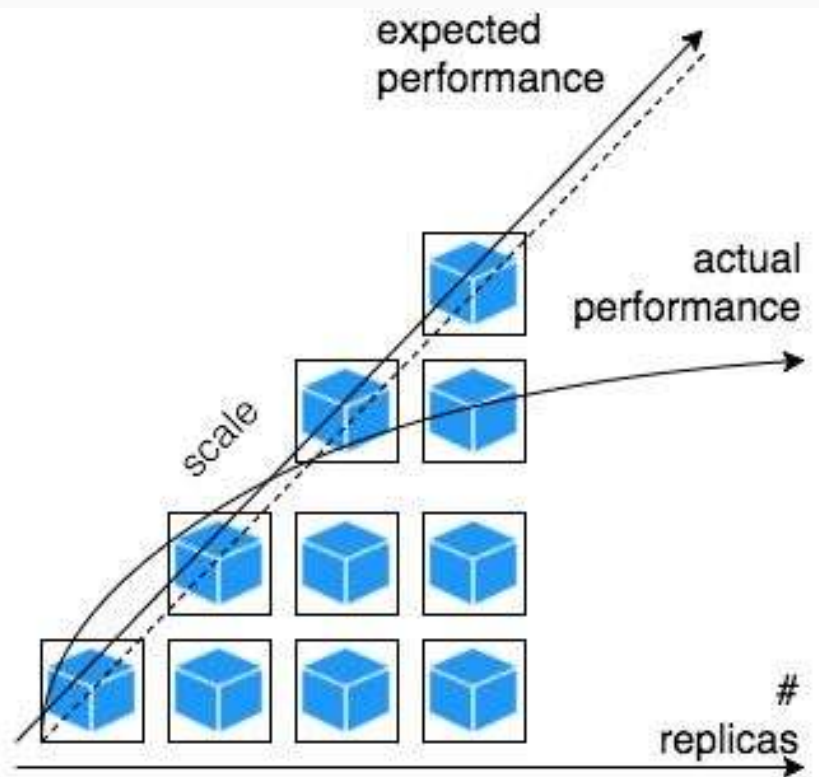


Example of docker app



Microservices architecture

# Motivation



## Why?

- The replicas are stealing resources from other microservices which are on the critical performance path
- The ineffective scaling containers which are not the bottleneck of the application

# How do Docker containers scale?

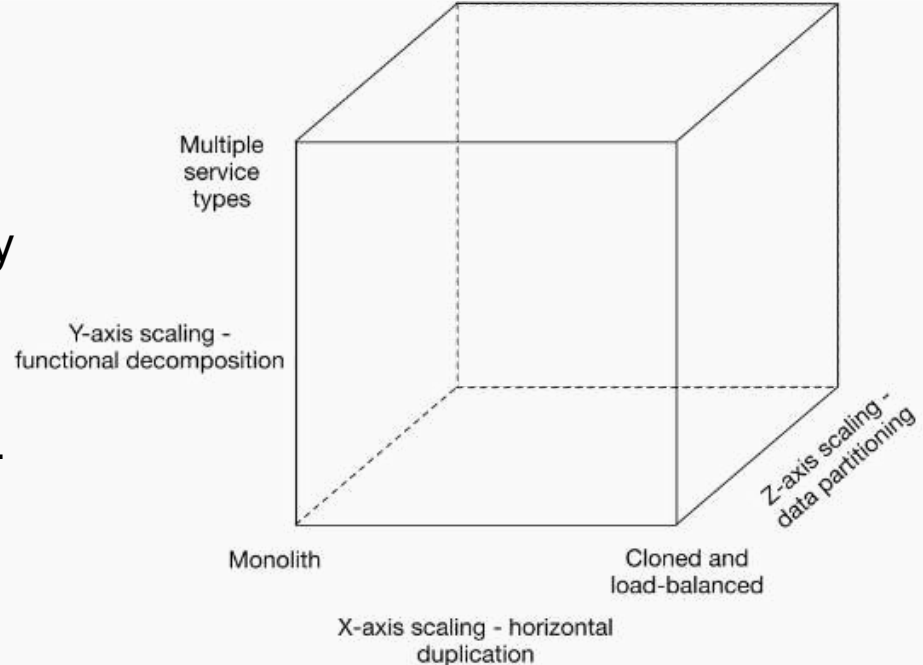
**Stateful docker container:** A container which maintains state locally across invocations. So must to handle the persistence of the state. Examples: databases and message queues.

**Stateless docker container:** Do not keep any state and therefore do not persist data except through other services.

# How do Docker containers scale?

## Scale-cube

- **X-axis**: horizontal duplication. Scale by cloning.
- **Z-axis**: data partitioning. Scale by splitting similar data structures.
- **Y-axis**: functional decomposition. Scale by splitting different functionality.



# Autoscaling

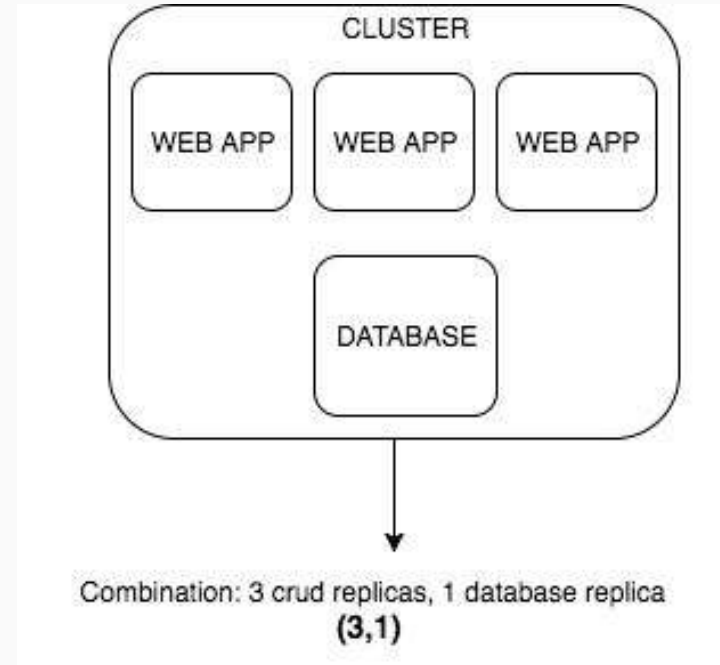
1. Some microservices do not scale out or in as fast as is needed. A prediction of when they need to scale is necessary to achieve elasticity.
2. Not all microservice implementations can be auto-scaled by instantiation alone. Stateful services are often not recognised automatically.
3. Auto-scaling over the top of a bottleneck is in vain.

## What is the best combination?

- Combination =  $(n_1, n_2, \dots, n_k)$

The 3 factors:

- Use case (fix)
- Cost
- Performance



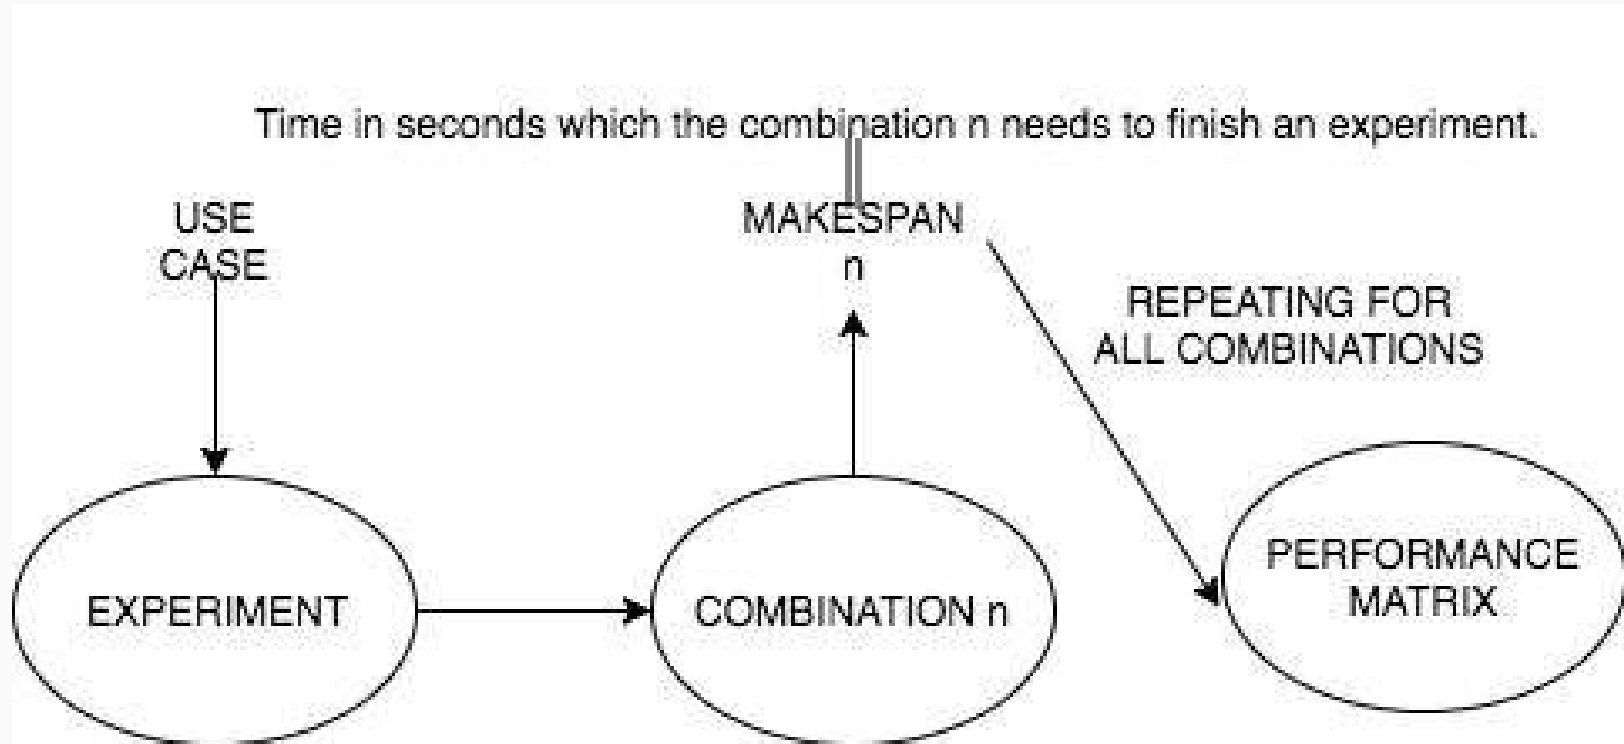
Example combination

# Scalability formula

- What is the most economical combination satisfying minimum performance constraints?
- What is the fastest combination satisfying maximum price constraints?



# Step 1: Create a performance matrix



And now some math to show how that works

# Step 2: Obtain a solution

$$\begin{aligned} & \text{fastest}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa) \\ & = i \mid \min_{\forall i \in I} \{m_i \in M_e \mid m_i < \text{max}_\mu, \text{cost}(i, \text{prices}) < \text{max}_\kappa\} \quad (1) \end{aligned}$$

$$\begin{aligned} & \text{cheapest}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa) \\ & = i \mid \min_{\forall i \in I} \{\text{cost}(i, \text{prices}) \mid M_e \ni m_i < \text{max}_\mu, \text{cost}(i, \text{prices}) < \text{max}_\kappa\} \quad (2) \end{aligned}$$

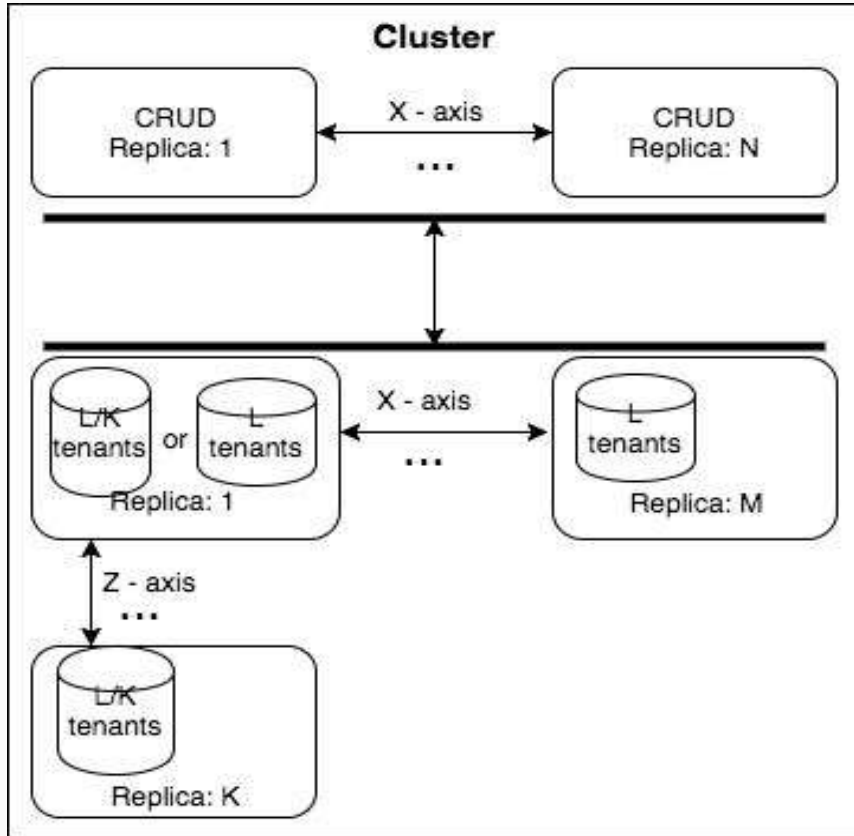
$$\begin{aligned} & \text{fastest\_rate}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa, \text{rate}) \\ & = i \mid \min_{\forall i \in I} (i \mid \frac{m_i}{m_k} \leq \text{rate}, \prec_{\text{cost}}) \end{aligned}$$

$$\text{where } k = \text{fastest}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa) \quad (3)$$

$$\begin{aligned} & \text{cheapest\_rate}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa, \text{rate}) \\ & = i \mid \min_{\forall i \in I} (i \mid \frac{\text{cost}(i, \text{prices})}{\text{cost}(k, \text{prices})} \leq \text{rate}, \prec_{\text{perf}}) \end{aligned}$$

$$\text{where } k = \text{cheapest}(M_e, \text{prices}, \text{max}_\mu, \text{max}_\kappa) \quad (4)$$

# Practical example



## USE CASE (Experiment)

- Search using word:
  - *documents/search/tenant/D/replica/word*
- Return the last “number” documents:
  - *documents/tenant/D/tenant/replica/lim/number*
- Return the last documents:
  - *documents/tenant/D/replica/last*
- Return the document with the id 4:
  - *documents/tenant/D/replica/4*

Docker images:

- mongo
- chumbo/arkiscrud:1.6.1

# Practical example

$$M_2 = \begin{pmatrix} 89.16 * 45.53 * 43.79 * 41.88 * 42.05 * 40.45 \\ 71.70 * 48.11 * 40.07 * 35.92 * 36.05 * 36.35 \end{pmatrix}$$

Title	Policy	$max_{\mu}$	$max_{\kappa}$	Rate	#S-ful	#S-less	Cost	Makespan
fastest	fastest	X	X	X	2	7	0.83	35.92
cheapest	cheapest	X	X	X	1	1	0.33	89.16
fastest with C	fastest	45.0	0.8	X	2	5	0.75	40.07
cheapest with C	cheapest	45.0	0.8	X	1	5	0.5	43.79
fastest(C & rate)	fastest	45.0	0.8	1.06	1	7	0.58	41.88
cheapest(C & rate)	cheapest	45.0	0.8	1.2	1	7	0.58	41.88

Legend: S-less = stateless, S-ful = stateful, C = constraints

# Demo

DOCKER APP

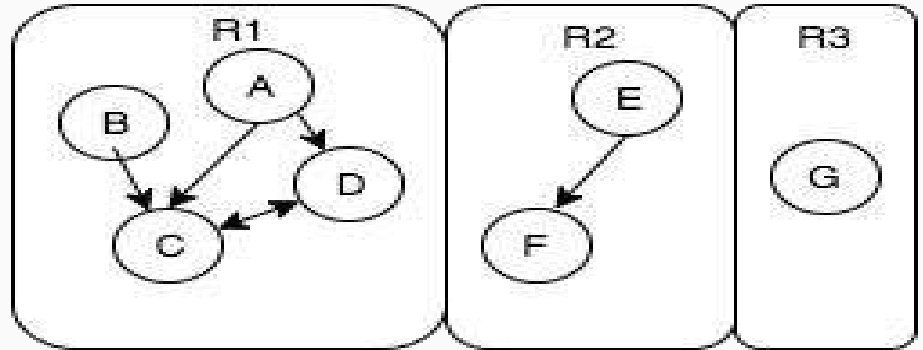


$$M_2 = \begin{pmatrix} 89.16 * 45.53 * 43.79 * 41.88 * 42.05 * 40.45 \\ 71.70 * 48.11 * 40.07 * 35.92 * 36.05 * 36.35 \end{pmatrix}$$

Formula script  
(Python)

Title	Policy	$max_{\mu}$	$max_{\kappa}$	Rate	#S-ful	#S-less	Cost	Makespan
fastest	fastest	X	X	X	2	7	0.83	35.92
cheapest	cheapest	X	X	X	1	1	0.33	89.16
fastest with C	fastest	45.0	0.8	X	2	5	0.75	40.07
cheapest with C	cheapest	45.0	0.8	X	1	5	0.5	43.79
fastest(C & rate)	fastest	45.0	0.8	1.06	1	7	0.58	41.88
cheapest(C & rate)	cheapest	45.0	0.8	1.2	1	7	0.58	41.88

Legend: S-less = stateless, S-ful = stateful, C = constraints

# Next steps



Until now:

- Specific workload

Future work:

- Formula to find the relation for each of the bi-directed connected graphs which compose the microservice architecture

# Our research on Cloud-Native Applications

- One of the research initiatives of the Service Prototyping Lab at Zurich University of Applied Sciences
- Successful transformation of legacy software into cloud-native apps
- Proven track record with CRM, DMS and other Swiss business apps

We co-innovate with software SMEs - contact us for more information!

## Cloud-Native Applications



# Links

<b>CNA research initiative</b>	<a href="https://blog.zhaw.ch/icclab/category/research-approach/themes/cloud-native-applications/">blog.zhaw.ch/icclab/category/research-approach/themes/cloud-native-applications/</a>
<b>Docker blog posts</b>	<a href="https://blog.zhaw.ch/icclab/tag/docker/">https://blog.zhaw.ch/icclab/tag/docker/</a>
<b>Application composed of multiple Docker containers</b>	<a href="https://github.com/serviceprototypinglab/scaling-containers">github.com/serviceprototypinglab/scaling-containers</a>
<b>Experiments - test Results</b>	<a href="https://github.com/serviceprototypinglab/scalability-experiments">github.com/serviceprototypinglab/scalability-experiments</a>
<b>Open Science Notebook</b>	
<b>Formula</b>	

