Wordpress en la era de Kubernetes

Fernando Perales

<me>

Ingeniero de Software Senior @ michelada.io

Promotor @ FLOSS

I ❤️ 🍺 & 🤘

Anfitrión @ RubyGDL

Co-organizador @ RailsBridge México

Estudiante de bajo eléctrico @ TEM UdeG

Aprendiz Krav Maga @ Alianza de Krav Maga México

</me>

Bertrand Russell

Filósofo, matemático, lógico y escritor británico ganador del Premio Nobel de Literatura y conocido por su influencia en la filosofía analítica, sus trabajos matemáticos y su activismo social.

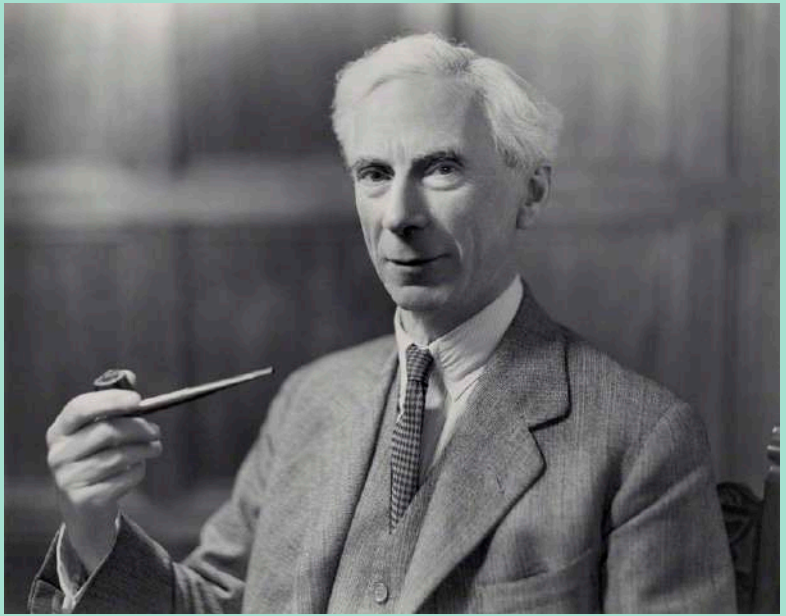

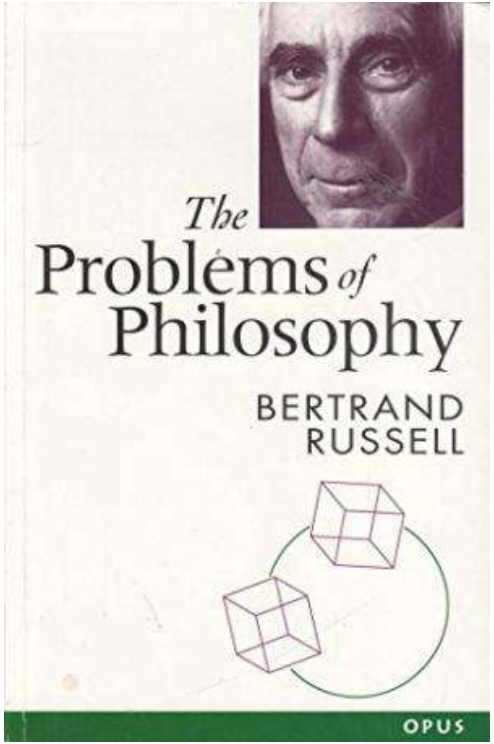

Los problemas de la filosofía

Los problemas del software*

Los problemas del software*

*Spoiler: todo

Recolección de requerimientos

An Apple gathering - Frederick Morgan

Estimación y planeación

http://www.texasholdem.com/wp-content/uploads/2014/09/These-Poker-Hands-Made-Me-Start-Winning-More.jpg

Colaboración

https://www.pamwoldow.com/files/2016/08/collaboration-rowers.jpg

Desarrollo

https://eusbf10p4zh1kzrai2rg2qyf-wpengine.netdna-ssl.com/wp-content/uploads/2018/02/FEATURED_An-Overview-of-the-Top-7-Software-Development-Methodologies-1080x675.jpeg

Pruebas

techbeacon.scdn7.secure.raxcdn.com/sites/default/files/styles/article_hero_image/public/field/image/mobile-testing-new-qa-frontier.jpg

Lanzamiento

https://i6.govx.net/images/189785_deployment_checklist_main.jpg?v=leB21RPmBcqKLBRe8uGfNw==

Seguridad

http://matryxconsulting.com.au/wp-content/uploads/2013/07/Physical-Security.jpg

Escalabilidad

http://blogs.cuit.columbia.edu/scaling/files/2013/05/cropped-COLOURBOX263277411.jpg

¿Por qué la escalabilidad es un problema especial?

Nunca se está listo

https://www.smileballard.com/wp-content/uploads/Dr.-Curelli-NYC-Marathon.jpg

Sobreingeniería

https://horasyminutos.com/wp-content/uploads/2015/09/Vacheron-Constantin-Reference-57260-Tivoli-anverso-horas.jpg

4 soluciones generales

Hacer menos

https://ei.marketwatch.com/Multimedia/2017/01/09/Photos/ZH/MW-FD366_na_wor_20170109173158_ZH.jpg?uuid=6ef018c2-d6bb-11e6-acdb-001cc448aede

Mejorar lo que ya tenemos

https://cachelot.io/assets/img/blog/fine-tuning-memcached.jpg

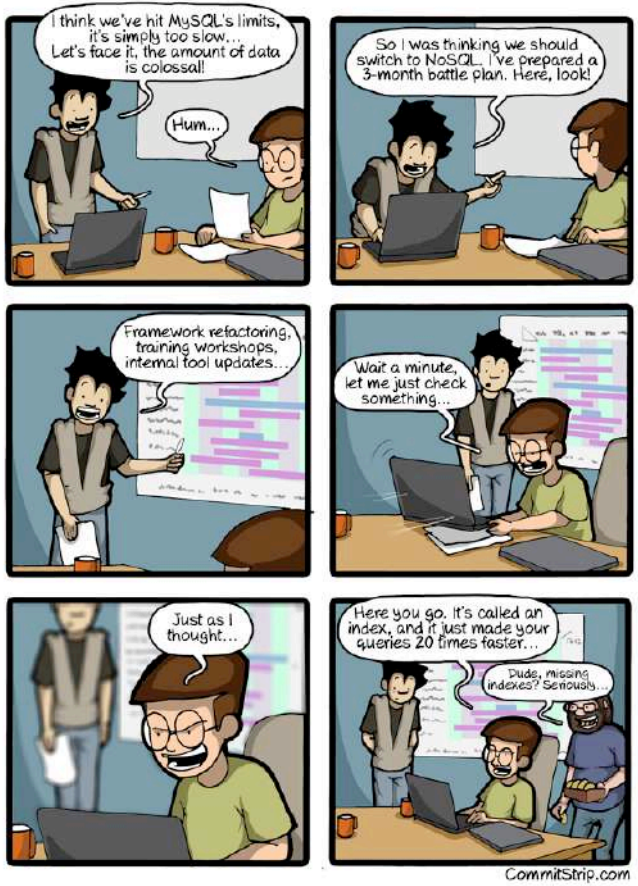

http://www.commitstrip.com/en/2014/06/03/the-problem-is-not-the-tool-itself/

Mejorar el hardware

https://cachelot.io/assets/img/blog/fine-tuning-memcached.jpg

Cambiar la arquitectura

https://publications.computer.org/micro/wp-content/uploads/sites/32/2016/11/architectural-plans.jpg

Hacer software es complejo

2008

DevOps

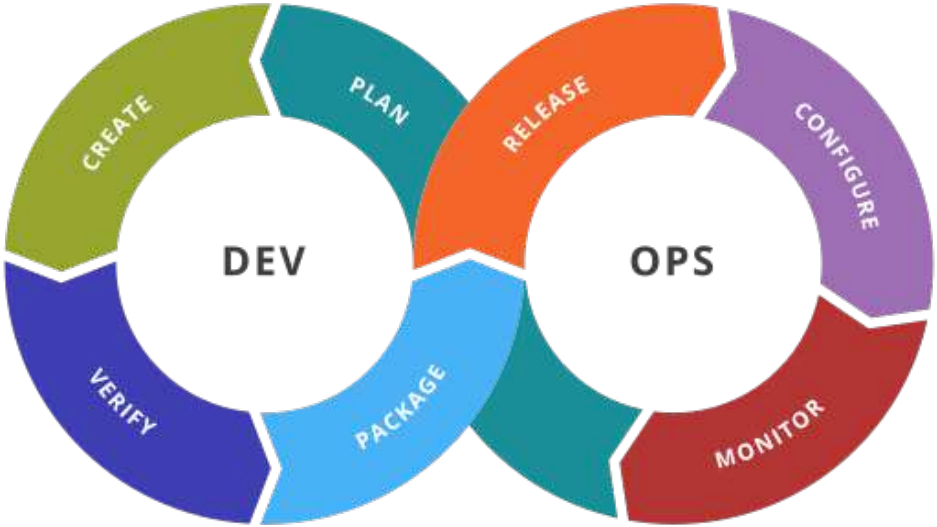

“DevOps is a set of practices intended to reduce the time between committing a change to a system and the change being placed into normal production, while ensuring high quality.”

https://en.wikipedia.org/wiki/DevOps#/media/File:Devops-toolchain.svg

Pets and cattle

Mascotas

Panela

Mascotas

Pascal

Panela

Mascotas

Ganado

Infrastructure as Code


puppet
labs


CHEF


SALTSTACK


ANSIBLE

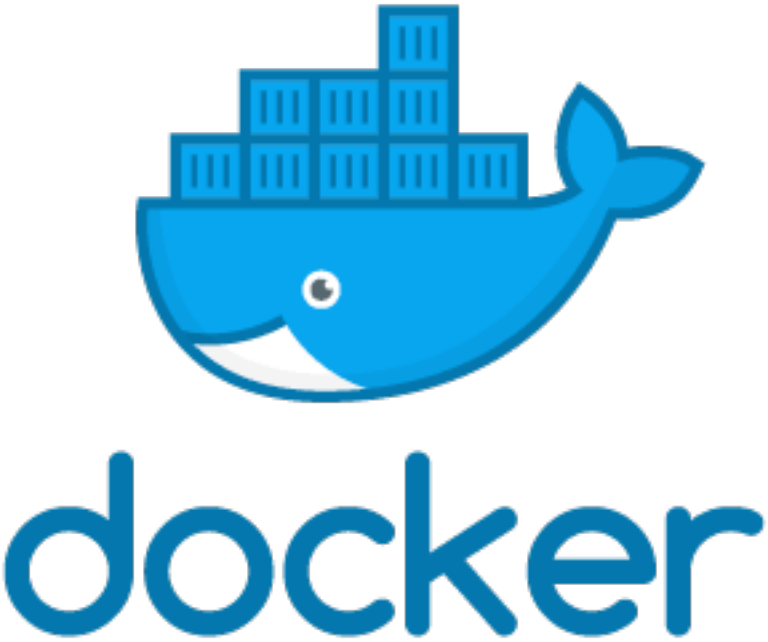
docker

The old way: Applications on host

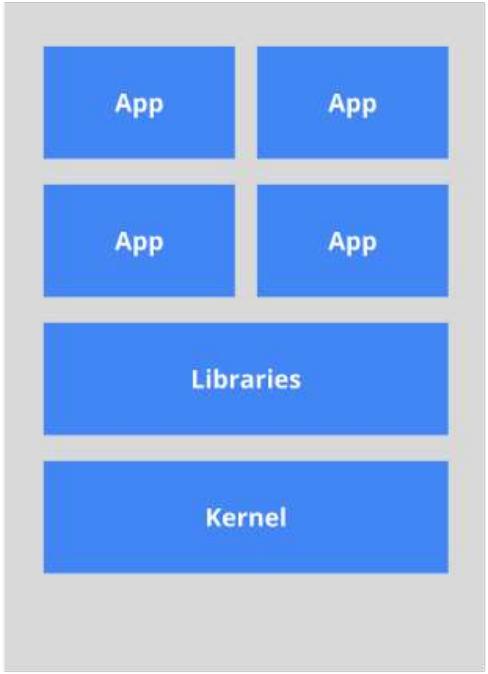

Heavyweight, non-portable
Relies on OS package manager

The new way: Deploy containers

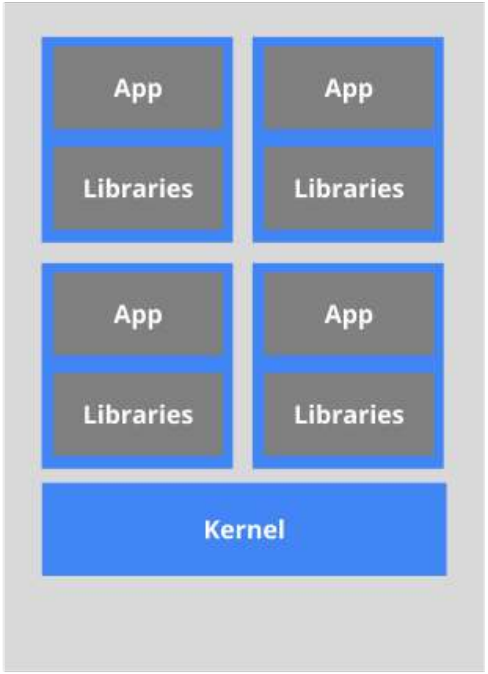

Small and fast, portable
Uses OS-level virtualization

https://d33wubrfki0l68.cloudfront.net/e7b766e0175f30ae37f7e0e349b87cfe2034a1ae/3e391/images/docs/why_containers.svg

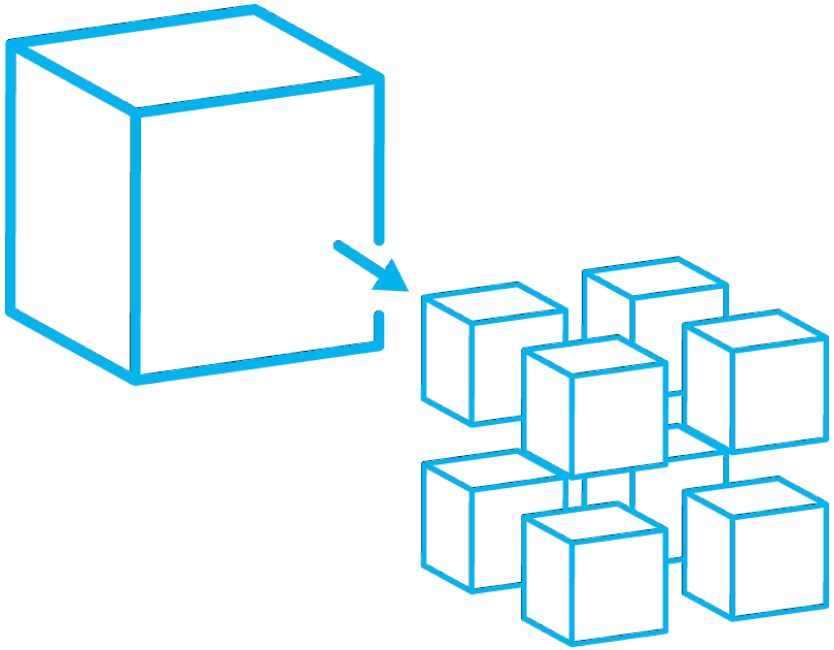

A monolithic application puts all its functionality into a single process...

... and scales by replicating the monolith on multiple servers

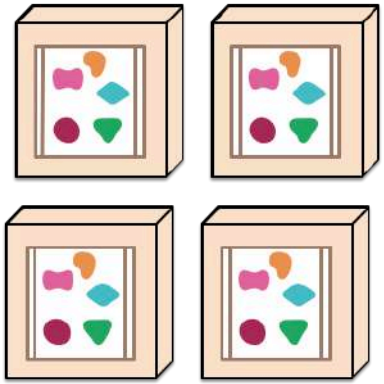

A microservices architecture puts each element of functionality into a separate service...

... and scales by distributing these services across servers, replicating as needed.

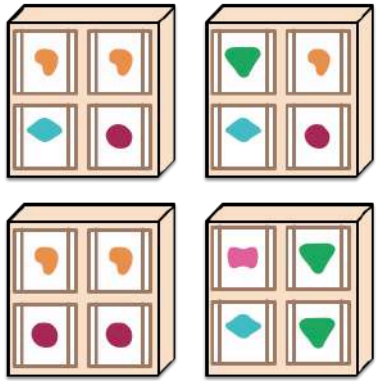

https://martinfowler.com/articles/microservices/images/sketch.png

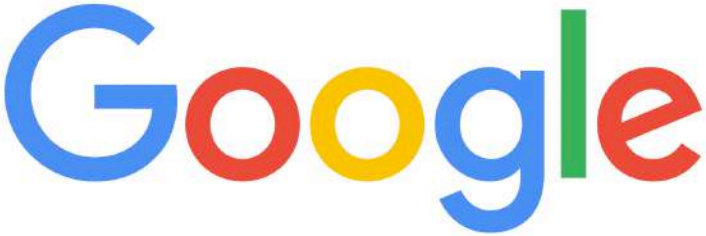
Google

Large-scale cluster management at Google with Borg

Abhishek Verma[†] Luis Pedrosa[‡] Madhukar Korupolu
David Oppenheimer Eric Tune John Wilkes

Google Inc.

Abstract

Google's Borg system is a cluster manager that runs hundreds of thousands of jobs, from many thousands of different applications, across a number of clusters each with up to tens of thousands of machines.

It achieves high utilization by combining admission control, efficient task-packing, over-commitment, and machine sharing with process-level performance isolation. It supports high-availability applications with runtime features that minimize fault-recovery time, and scheduling policies that reduce the probability of correlated failures. Borg simplifies life for its users by offering a declarative job specification language, name service integration, real-time job monitoring, and tools to analyze and simulate system behavior.

We present a summary of the Borg system architecture and features, important design decisions, a quantitative analysis of some of its policy decisions, and a qualitative examination of lessons learned from a decade of operational experience with it.

1. Introduction

The cluster management system we internally call Borg admits, schedules, starts, restarts, and monitors the full range of applications that Google runs. This paper explains how.

Borg provides three main benefits: it (1) hides the details

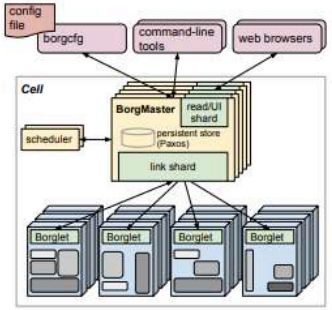

Figure 1: The high-level architecture of Borg. *Only a tiny fraction of the thousands of worker nodes are shown.*

cluding with a set of qualitative observations we have made from operating Borg in production for more than a decade.

2. The user perspective

Borg's users are Google developers and system administrators (site reliability engineers or SREs) that run Google's applications and services. Users submit their work to Borg

2014

kubernetes

Manejo automático de recursos

http://www.sanyglobal.com/skin/default_en/images/gangji/img-1.jpg

Crecimiento horizontal

http://metrolimousines.com/wp-content/uploads/Wineries-Vineyard-Long-Island-NY.jpg

Despliegues y retrocesos automatizados

https://static.timesofisrael.com/www/uploads/2017/05/red-carpet-e1495124255117.jpg

Orquestación de almacenamiento

https://www.google.com.mx/url?sa=i&rct=j&q=&esrc=s&source=images&cd=&cad=rja&uact=8&ved=2ahUKEwj0wYvE-IXdAhULLKwKHUI-

Auto curación

https://www.news965.com/news/local/new-jersey-nurse-caught-reusing-flu-shot-syringes/Bqu9jj9CRu2E0UgFeVjP7O/#eb656511-

Auto descubrimiento y balanceamiento de carga

https://www.bestbarcelonaapartments.com/wp-content/uploads/2017/08/que-ver-con-nin-os-en-barcelona_ljq5s5.jpg

Manejo de secretos y configuraciones

https://www.technobuffalo.com/wp-content/uploads/2014/04/top-secret-2-470x310@2x.jpeg

Demo

¡Gracias!

¿Tienes preguntas?

@FerPeralesM

hola@ferperales.net

ferperales.net