



kubernetes

Productivity and Efficiency for dev. teams



Who Am I?

Spas Atanasov

Lead DevOps Engineer



 /spasatanasov

!

DISCLAIMER

“

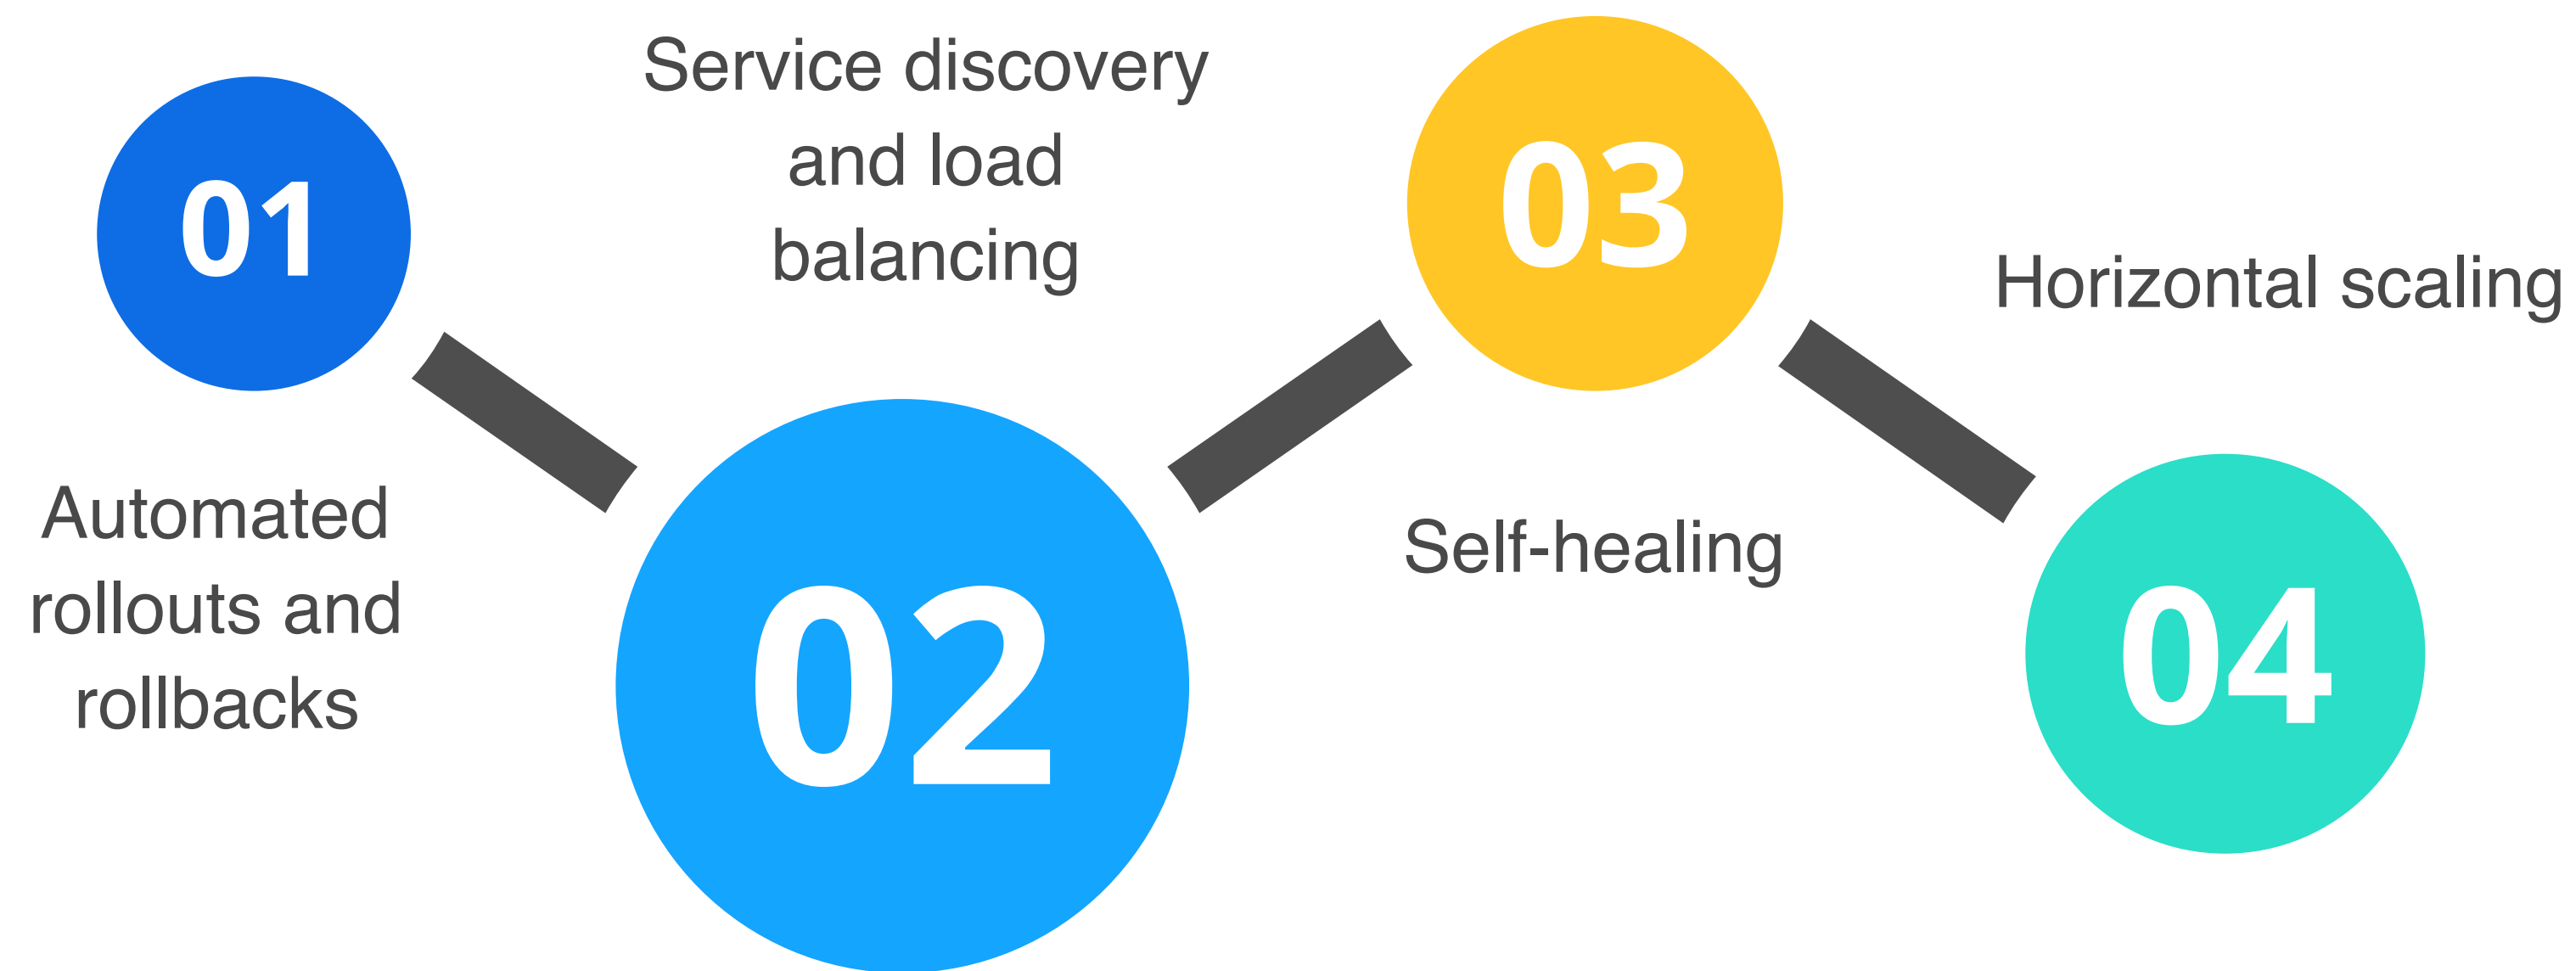
“A person who
never made a
mistake never
tried anything
new”




WHAT IS KUBERNETES?

- ✓ OpenSource Platform (by Google in 2014)
- ✓ Container manager and orchestrator
- ✓ Portable and Extendable
- ✓ Multi-cloud and on-premise

TOP FEATURES






Trend?



kubernetes

Google Trends / Worldwide / 5 years

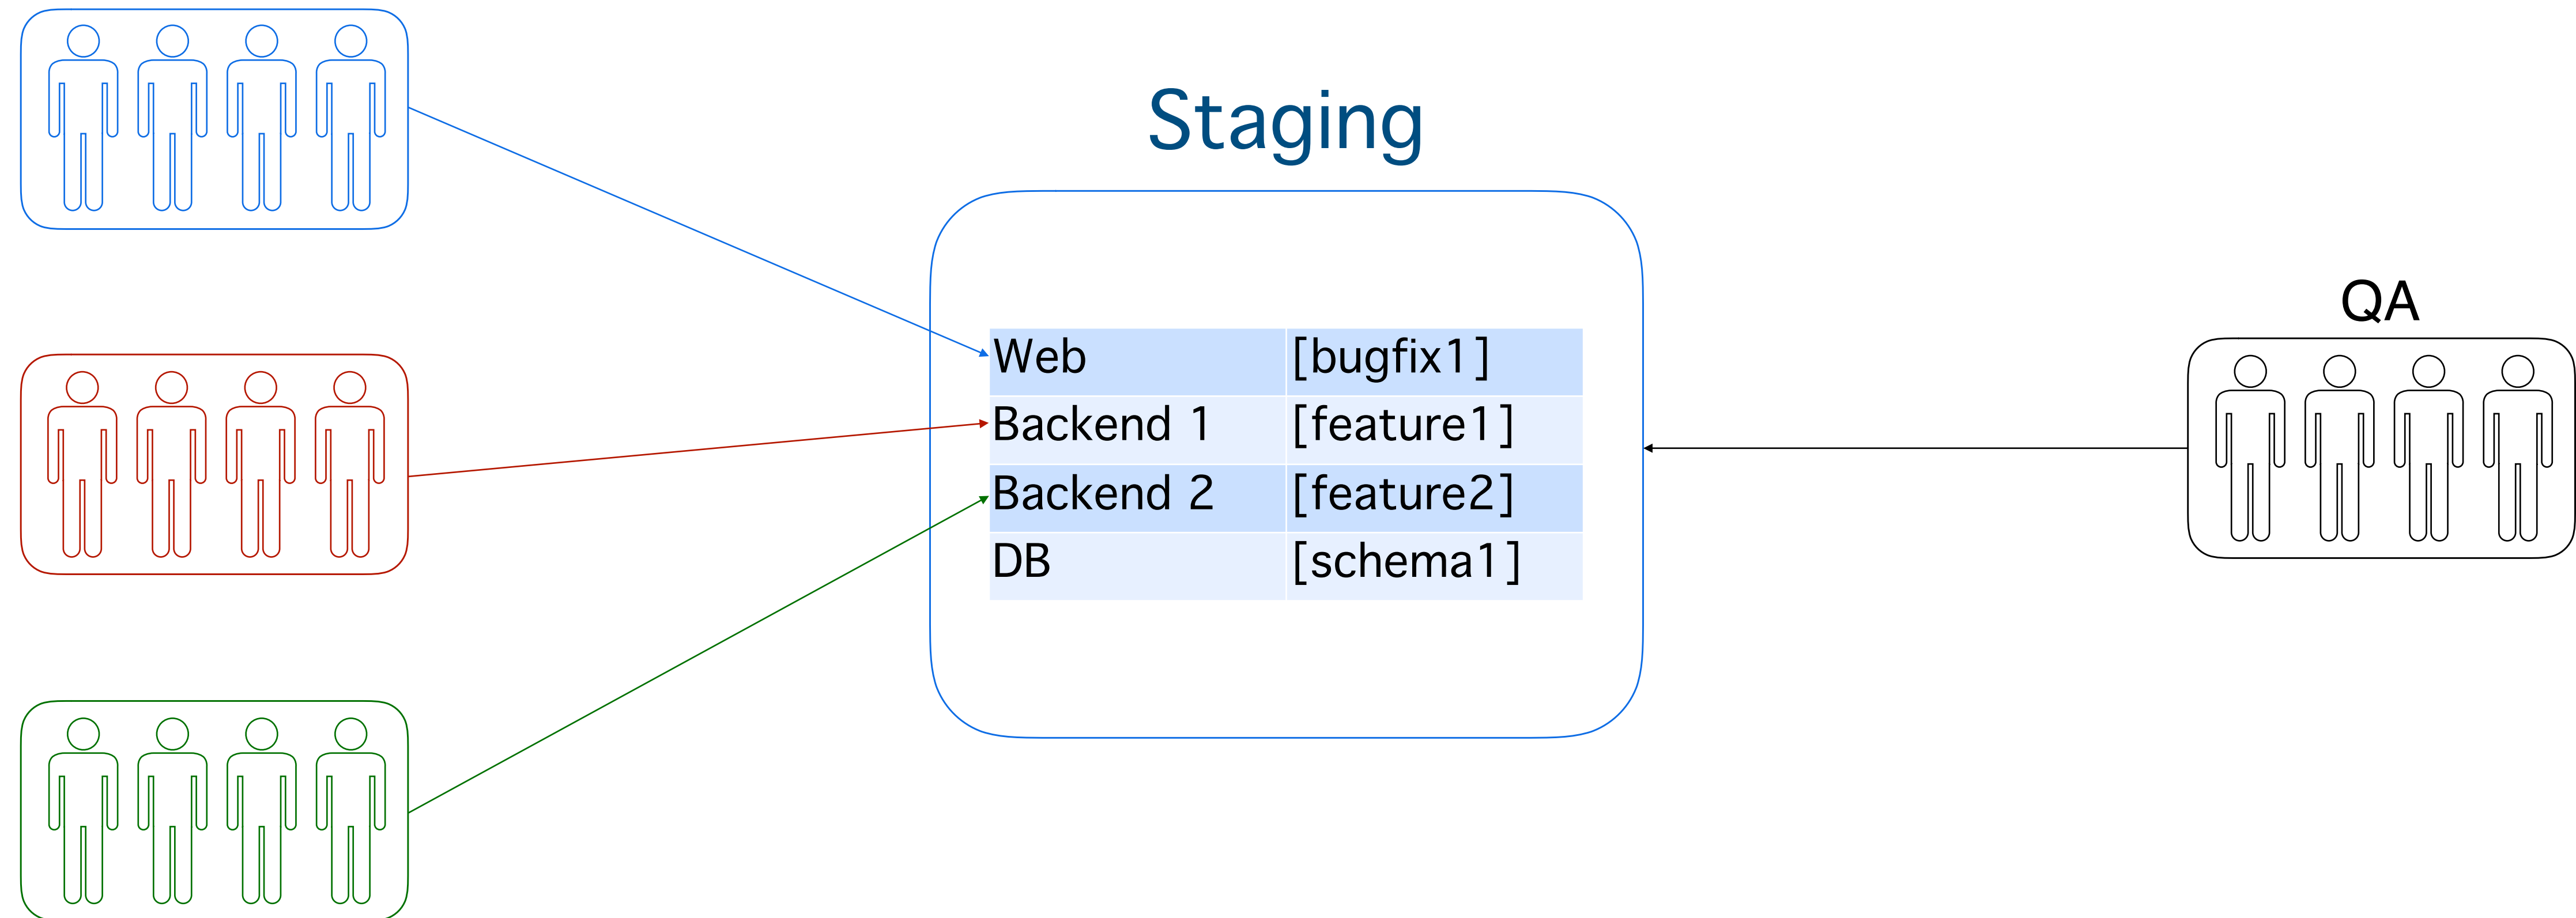


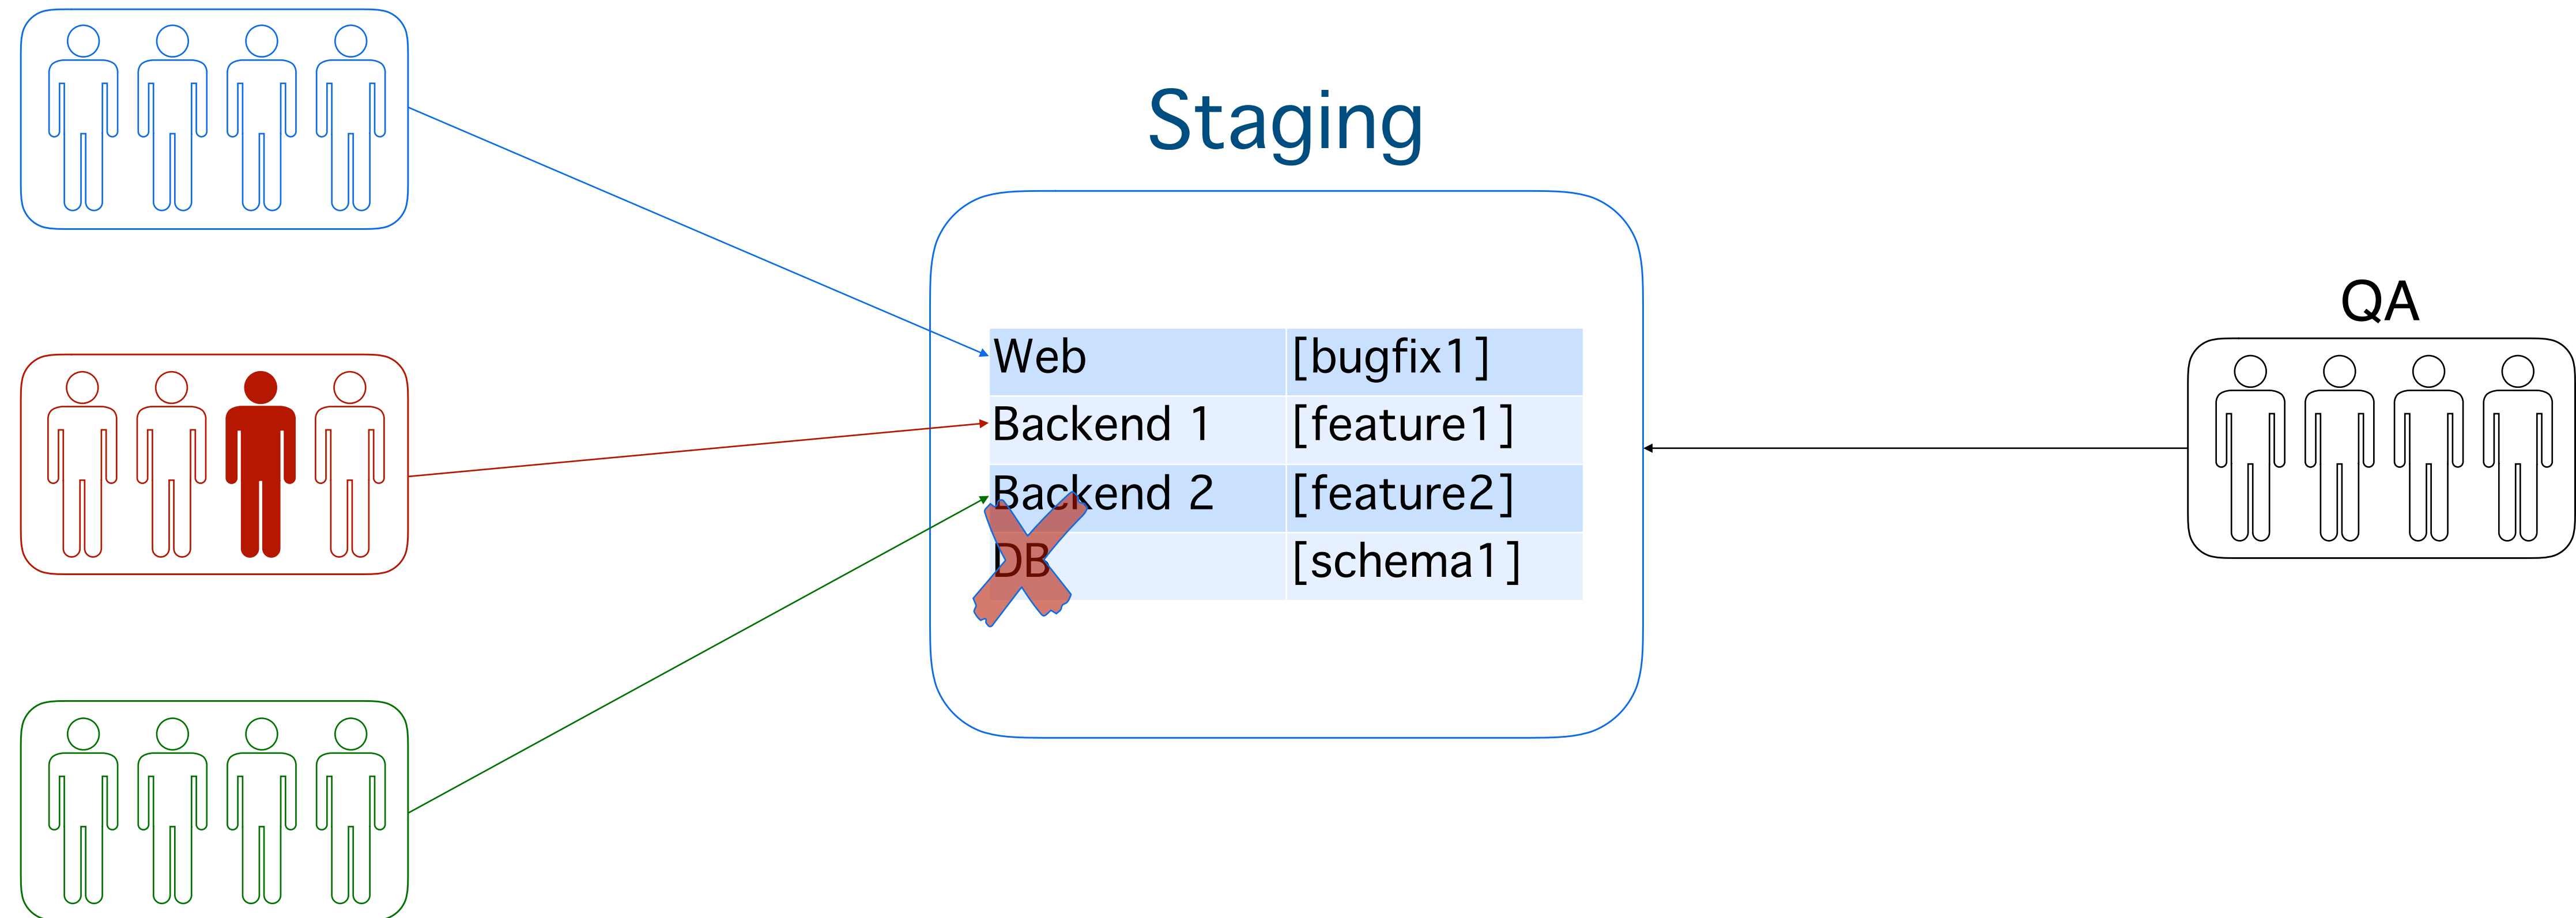
Staging

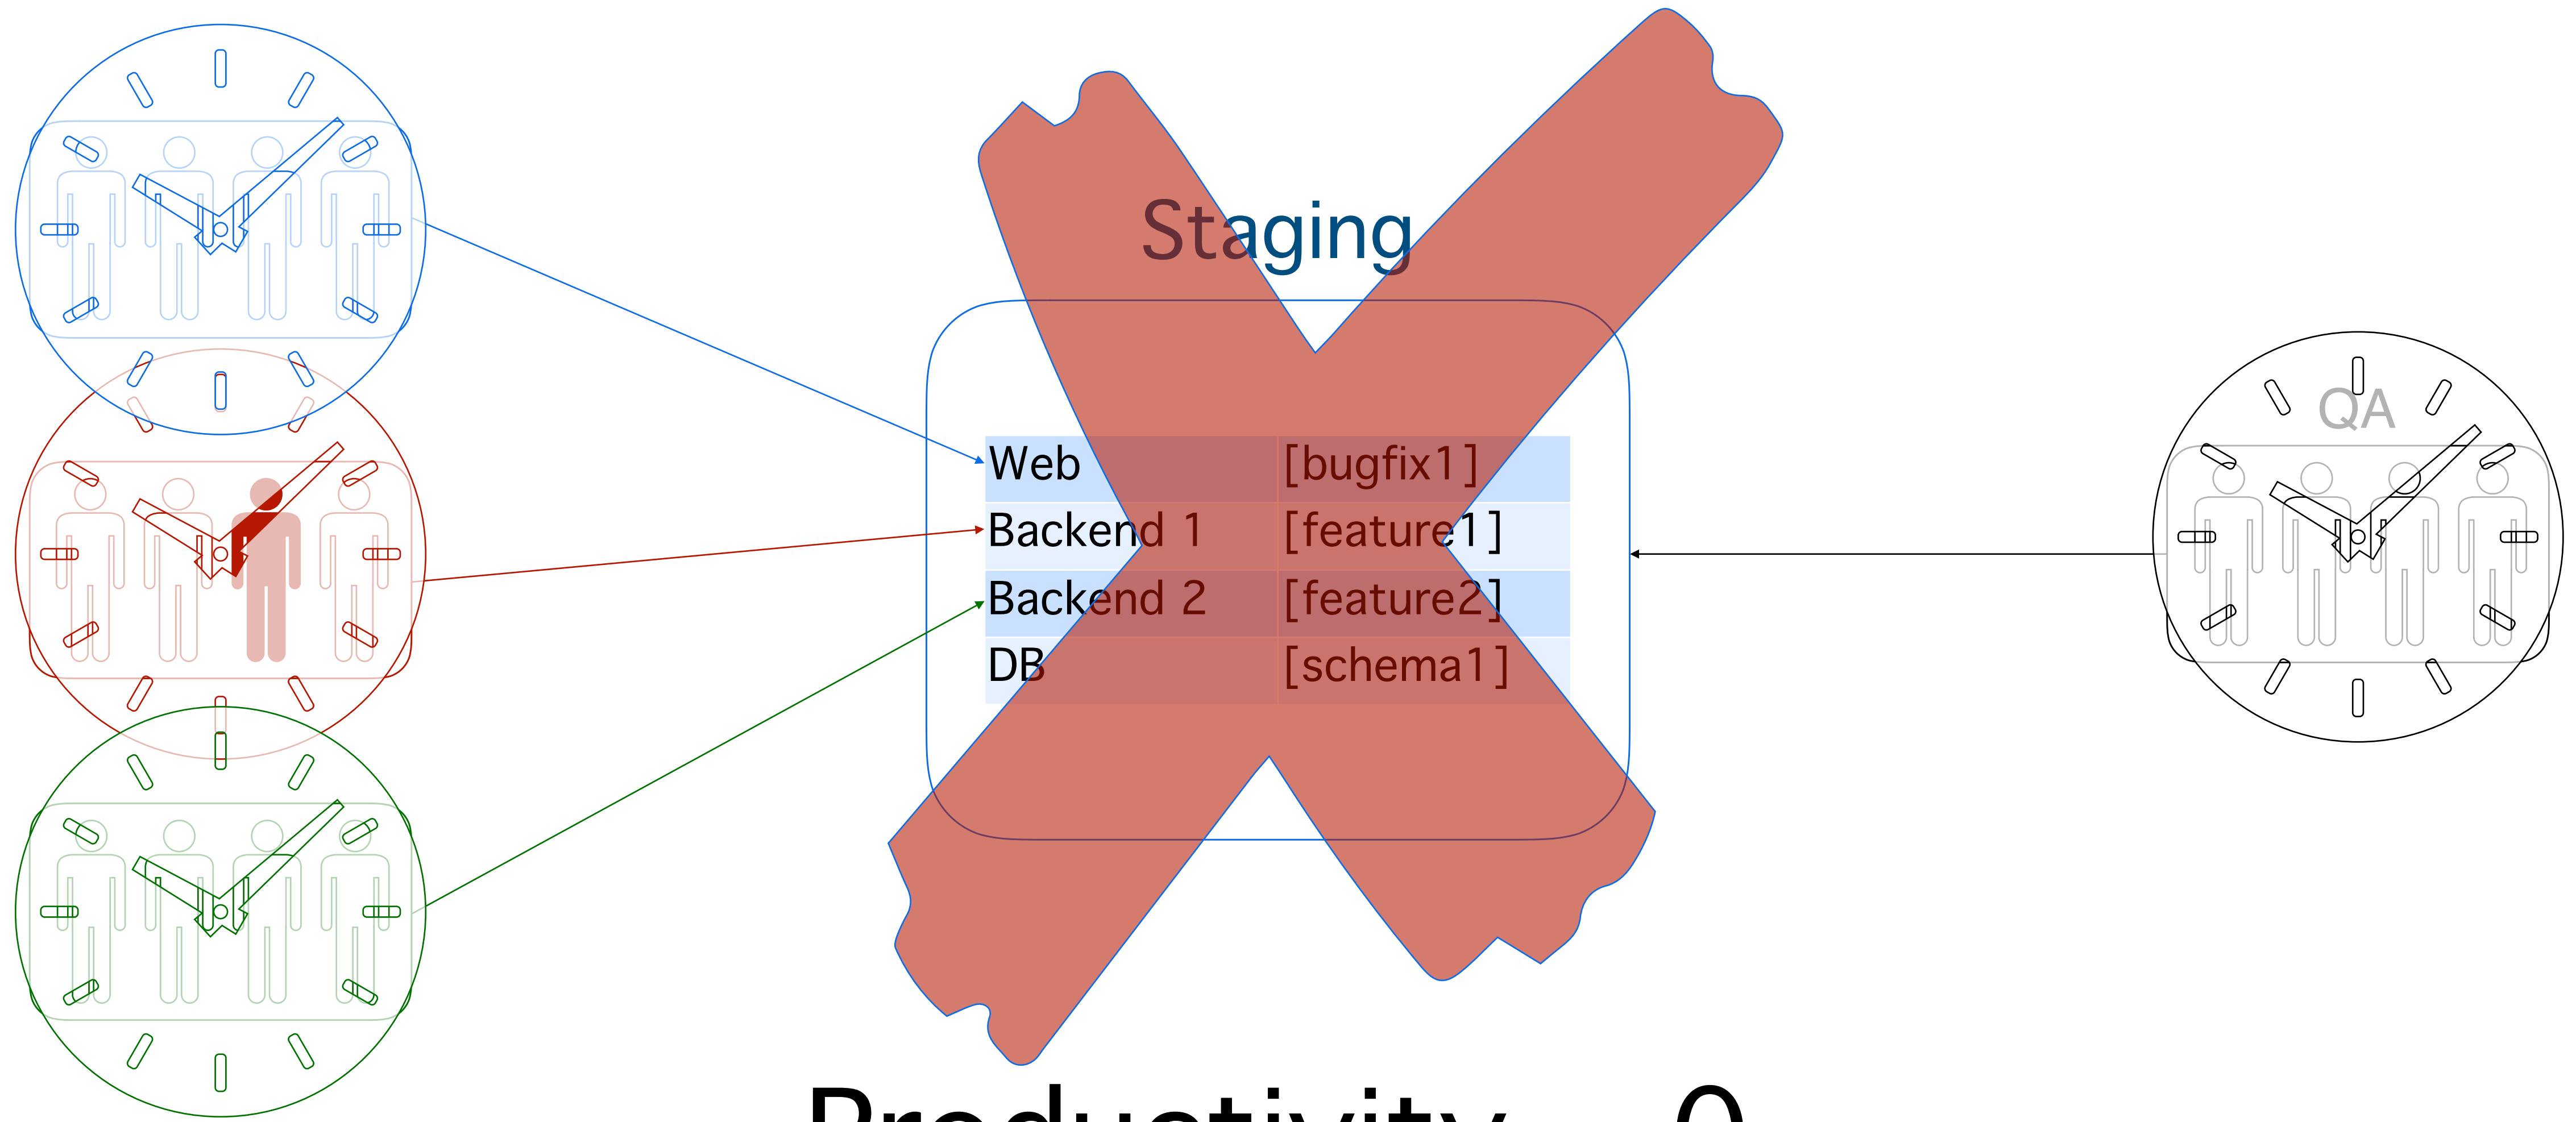
Web	[bugfix1]
Backend 1	[feature1]
Backend 2	[feature2]
DB	[schema1]

Production

Web	[master]
Backend 1	[master]
Backend 2	[master]
DB	[master]







Productivity = 0



Per Person Isolated Environment

Kubernetes namespaces are intended for use in environments with many users spread across multiple teams, or projects. For clusters with a few to tens of users, you should not need to create or think about namespaces at all. Start using namespaces when you need the features they provide.

Namespaces are a way to divide cluster resources between multiple users.

Kubernetes Cluster - Dev

Developer-1

Web	[master]
Backend 1	[feature1]
Backend 2	[feature2]
DB	[schema1]

Developer-2

Web	[feature1]
Backend 1	[master]
Backend 2	[staging]
DB	[schema2]

Developer-3

Web	[bugfix1]
Backend 1	[feature2]
Backend 2	[staging]
DB	[master]



**DON'T MAKE THEM
WAIT !**

DEPLOY TIME

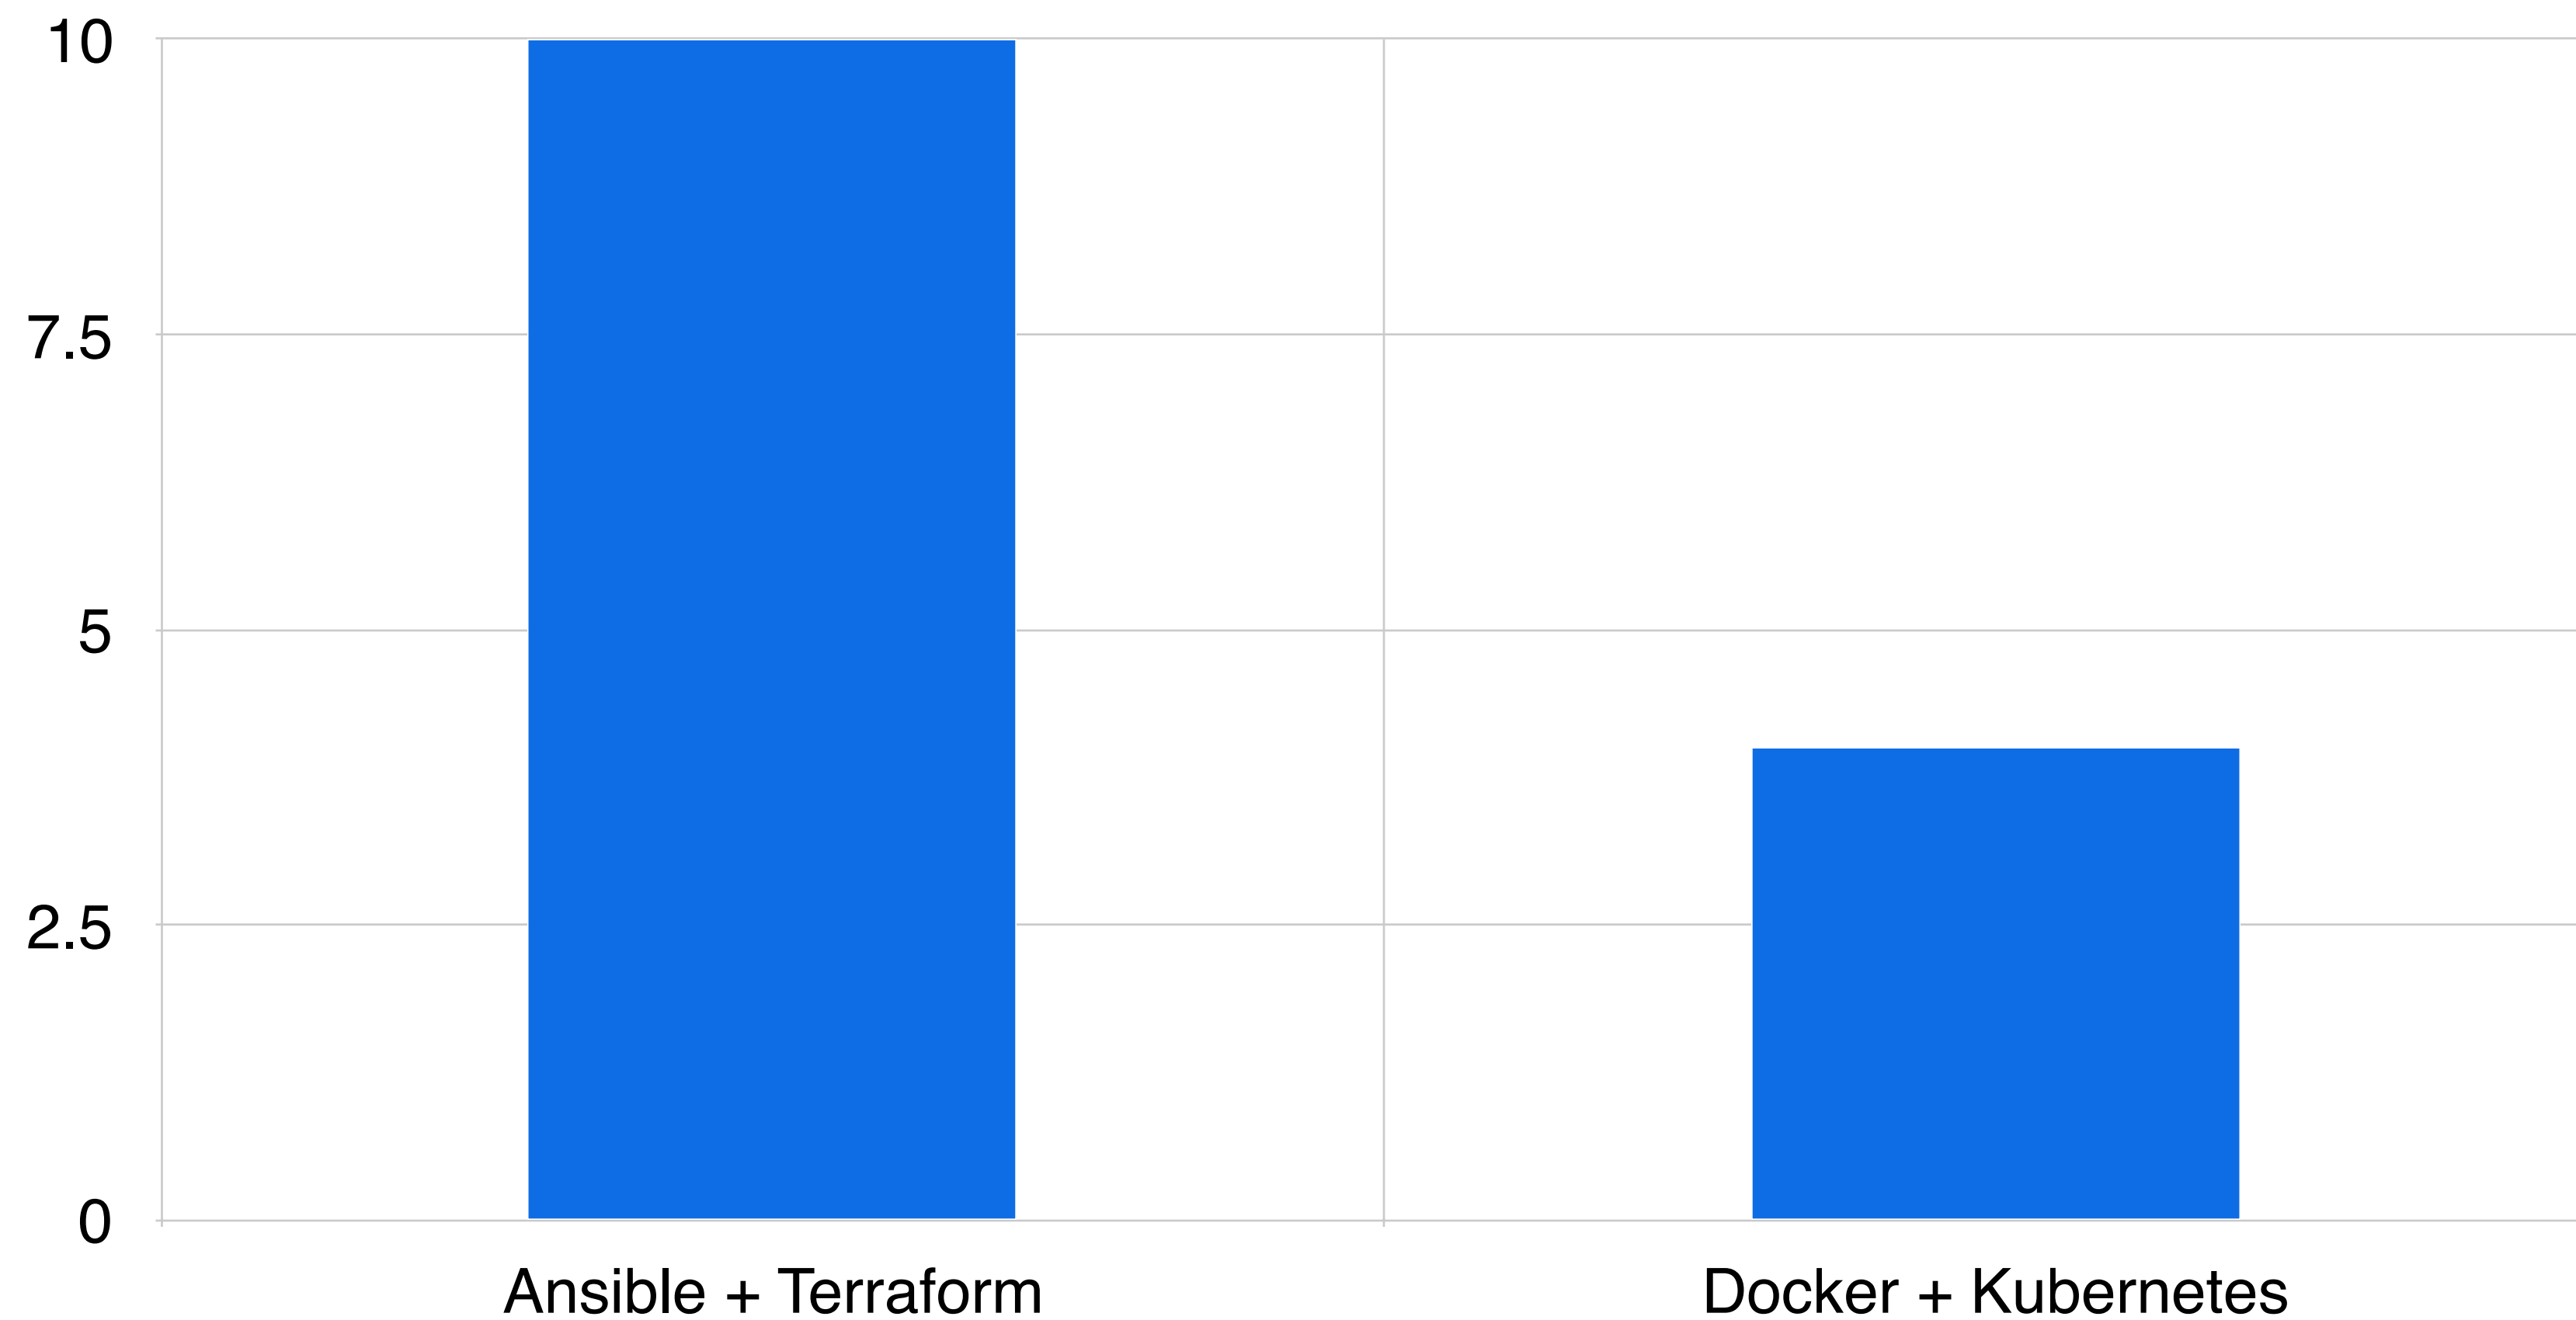
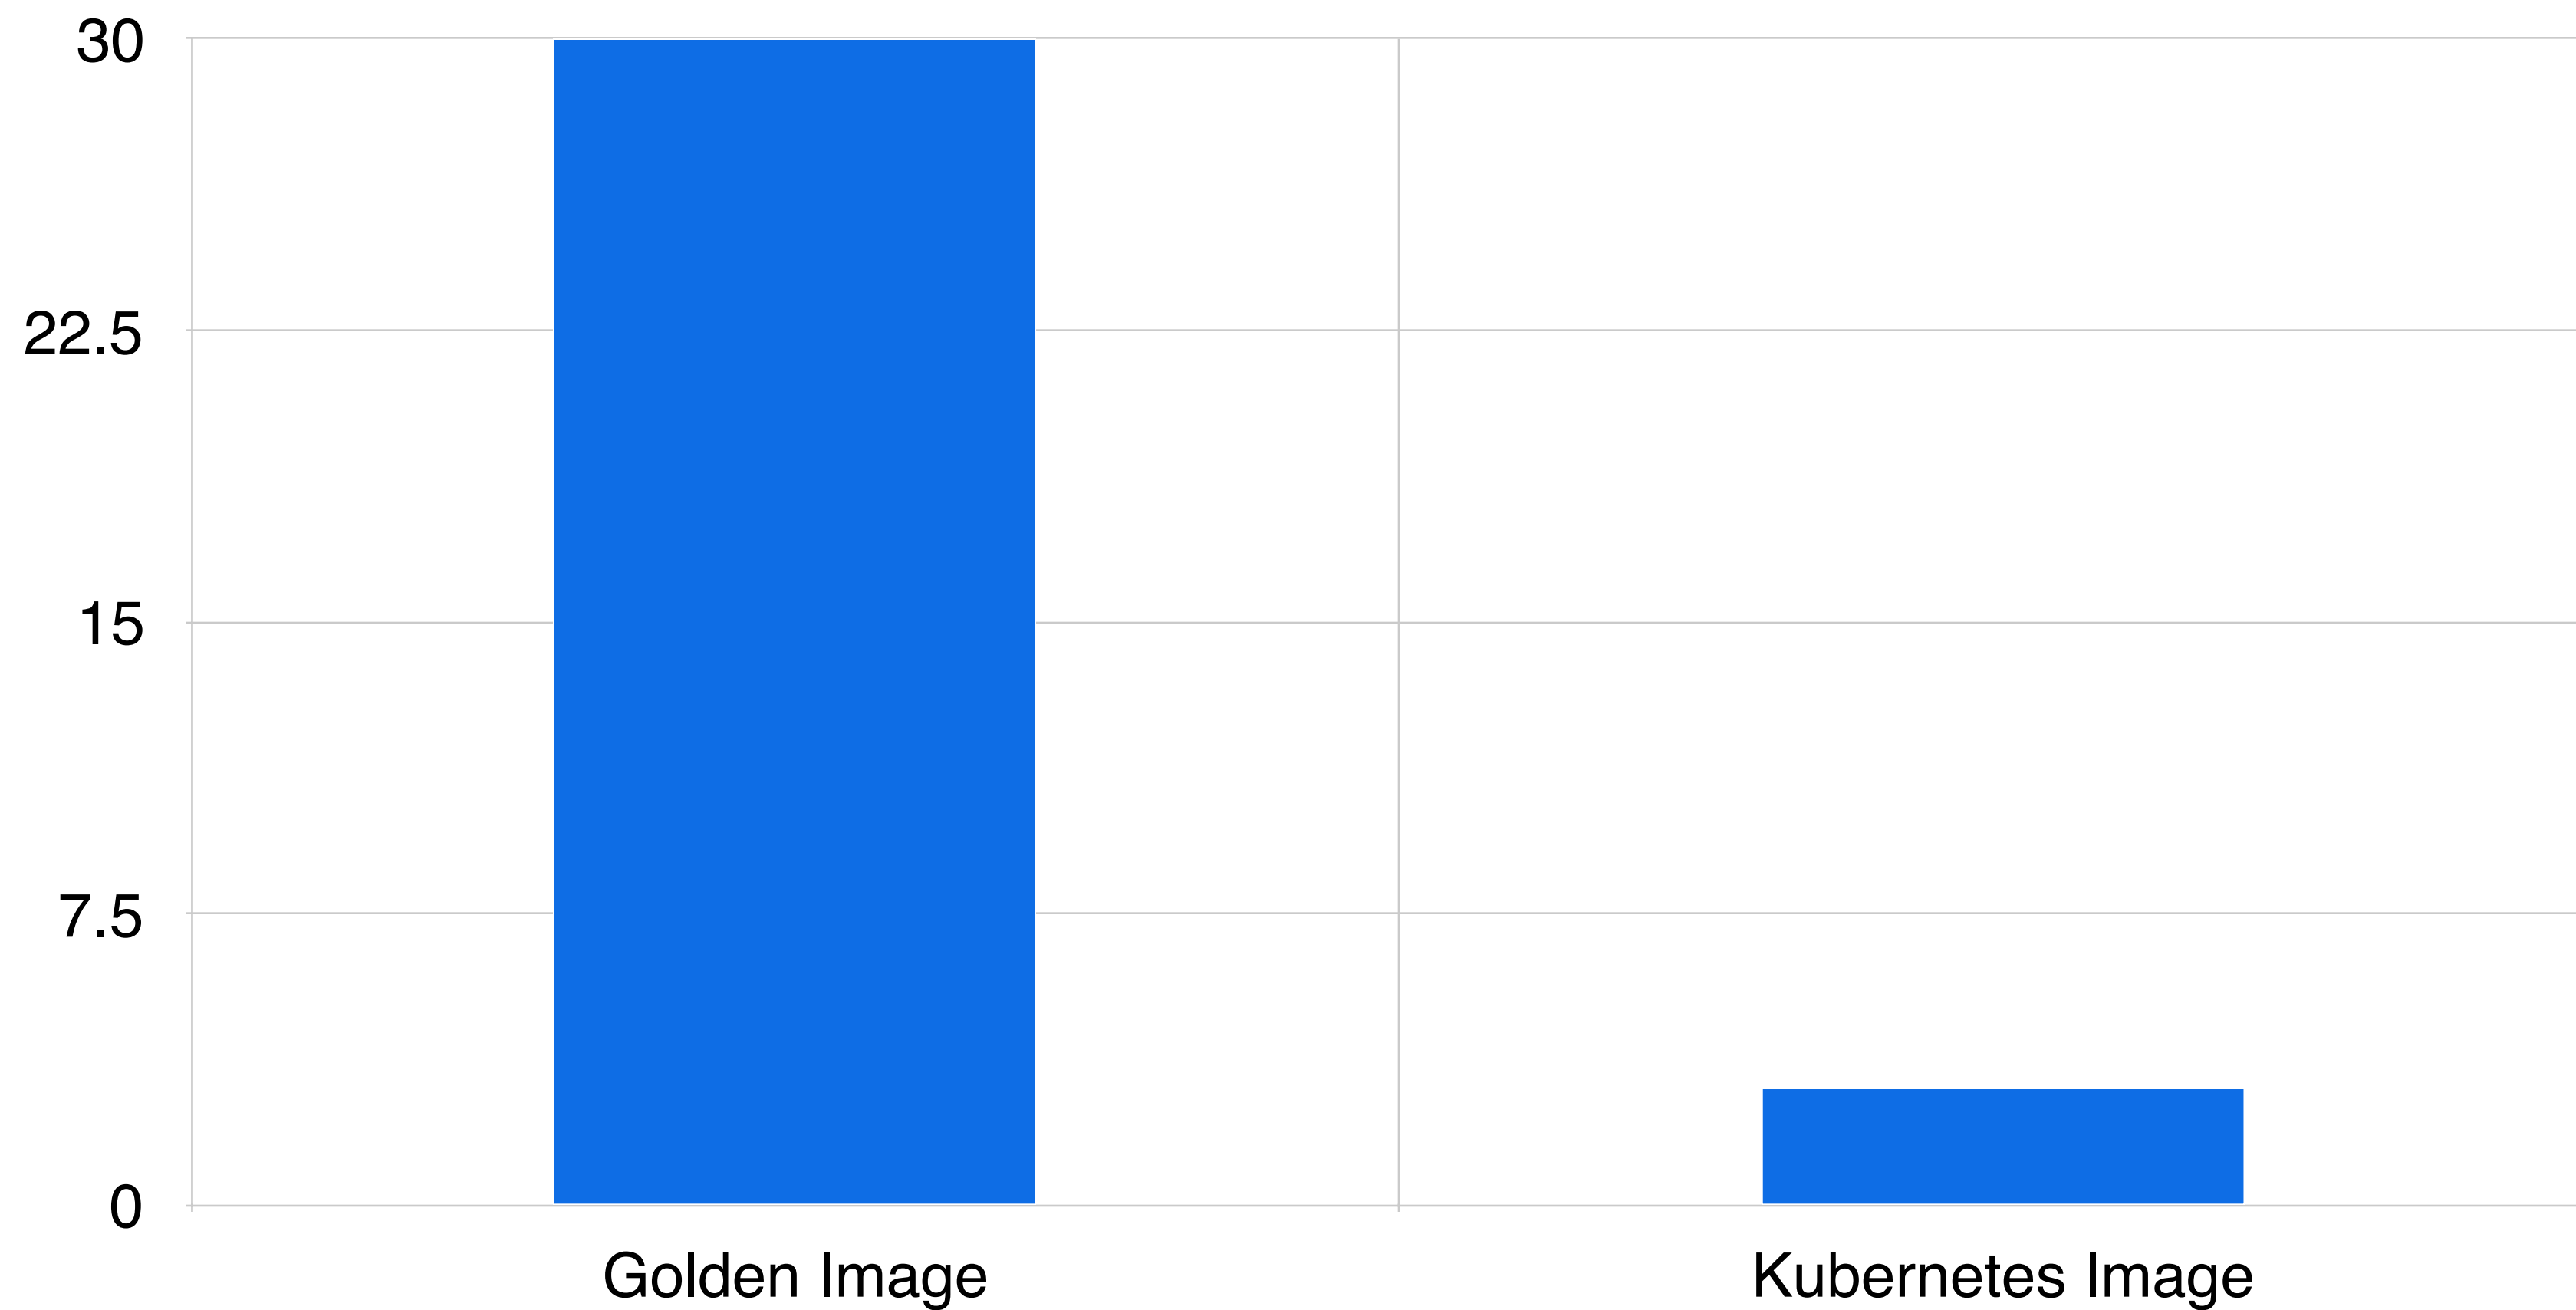



IMAGE BUILD TIME



FEATURES AND COSTS



KUBE-FLOW

The Machine Learning Toolkit for Kubernetes



Notebooks

A JupyterHub to create and manage interactive Jupyter notebooks. Project Jupyter is a non-profit, open-source project to support interactive data science and scientific computing across all programming languages.

TensorFlow model training

A TensorFlow Training Controller that can be configured to use either CPUs or GPUs and be dynamically adjusted to the size of a cluster with a single setting. We also provide a TensorFlow job operator.



Model serving

A TensorFlow Serving container to export trained TensorFlow models to Kubernetes. We also integrate with Seldon Core, an open source platform for deploying machine learning models on Kubernetes, and NVIDIA TensorRT Inference Server for maximized GPU utilization when deploying ML/DL models at scale.



Thanks!

Any questions?

 /spasatanasov

Meetup / Kubernetes & Cloud Native Computing Bulgaria