A line chart with a dark background and a light grid. It features three data series: a blue line at the top, an orange line in the middle, and a green line at the bottom. All three lines show a similar trend, with a significant peak around the center of the chart, marked by a vertical dashed line. The blue line peaks highest, followed by the orange line, and then the green line. The lines are connected by small dots at each data point.

Devstats

A Full-Stack, All-the-Things, Clear Example End-to-End Application on Kubernetes

KubeCon EU 2019
#devstats

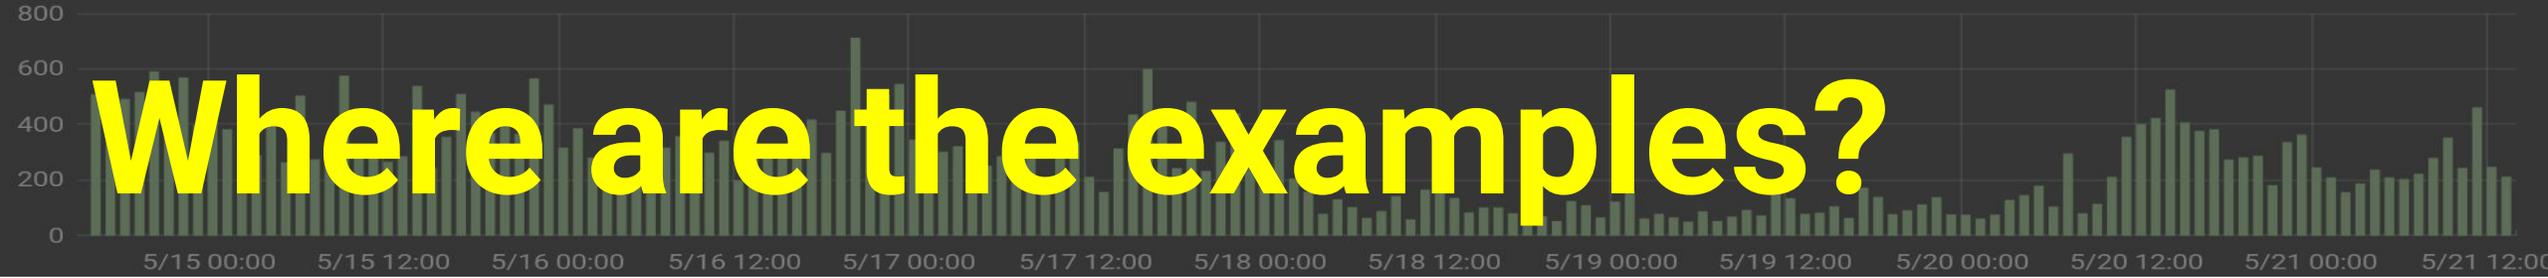
Łukasz Gryglicki, CNCF
Josh Berkus, Red Hat

Who?



Łukasz Gryglicki
Developer for the CNCF
Primary author of
DevStats

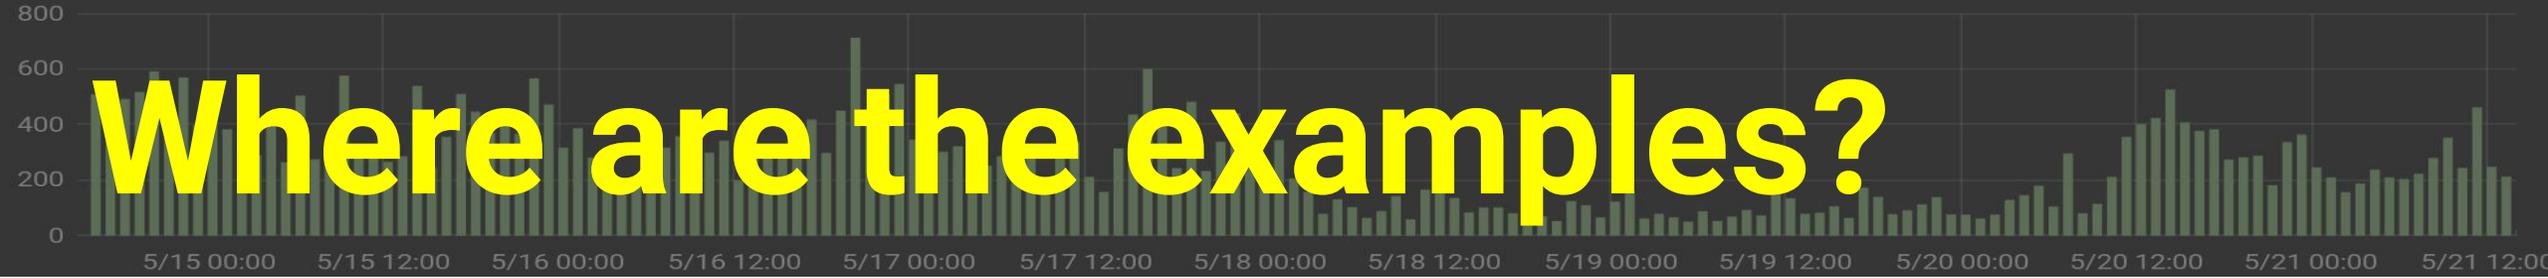
Josh Berkus
Community Manager for Red Hat
Contributor to SIG-ContribEx and
Devstats
@fuzzychef



Where are the examples?

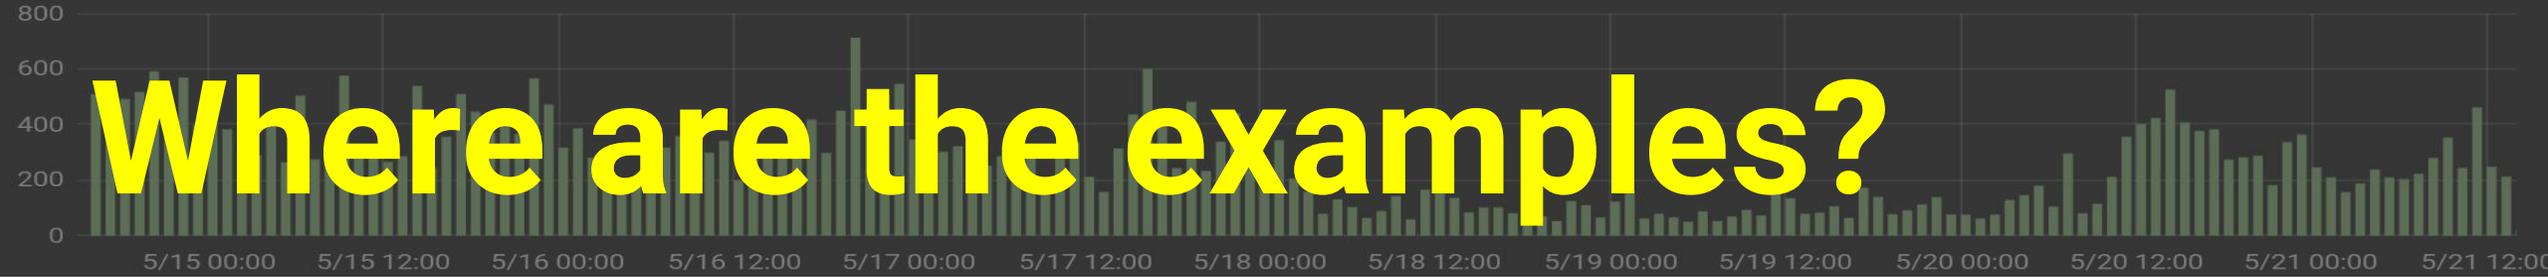
So, we decided to build a Kubernetes app using current practices and tools ...

- *Deployment*
- *Configuration*
- *CI/CD*
- *Packaging (Helm)*
- UI
- Database
- Routing
- Jobs



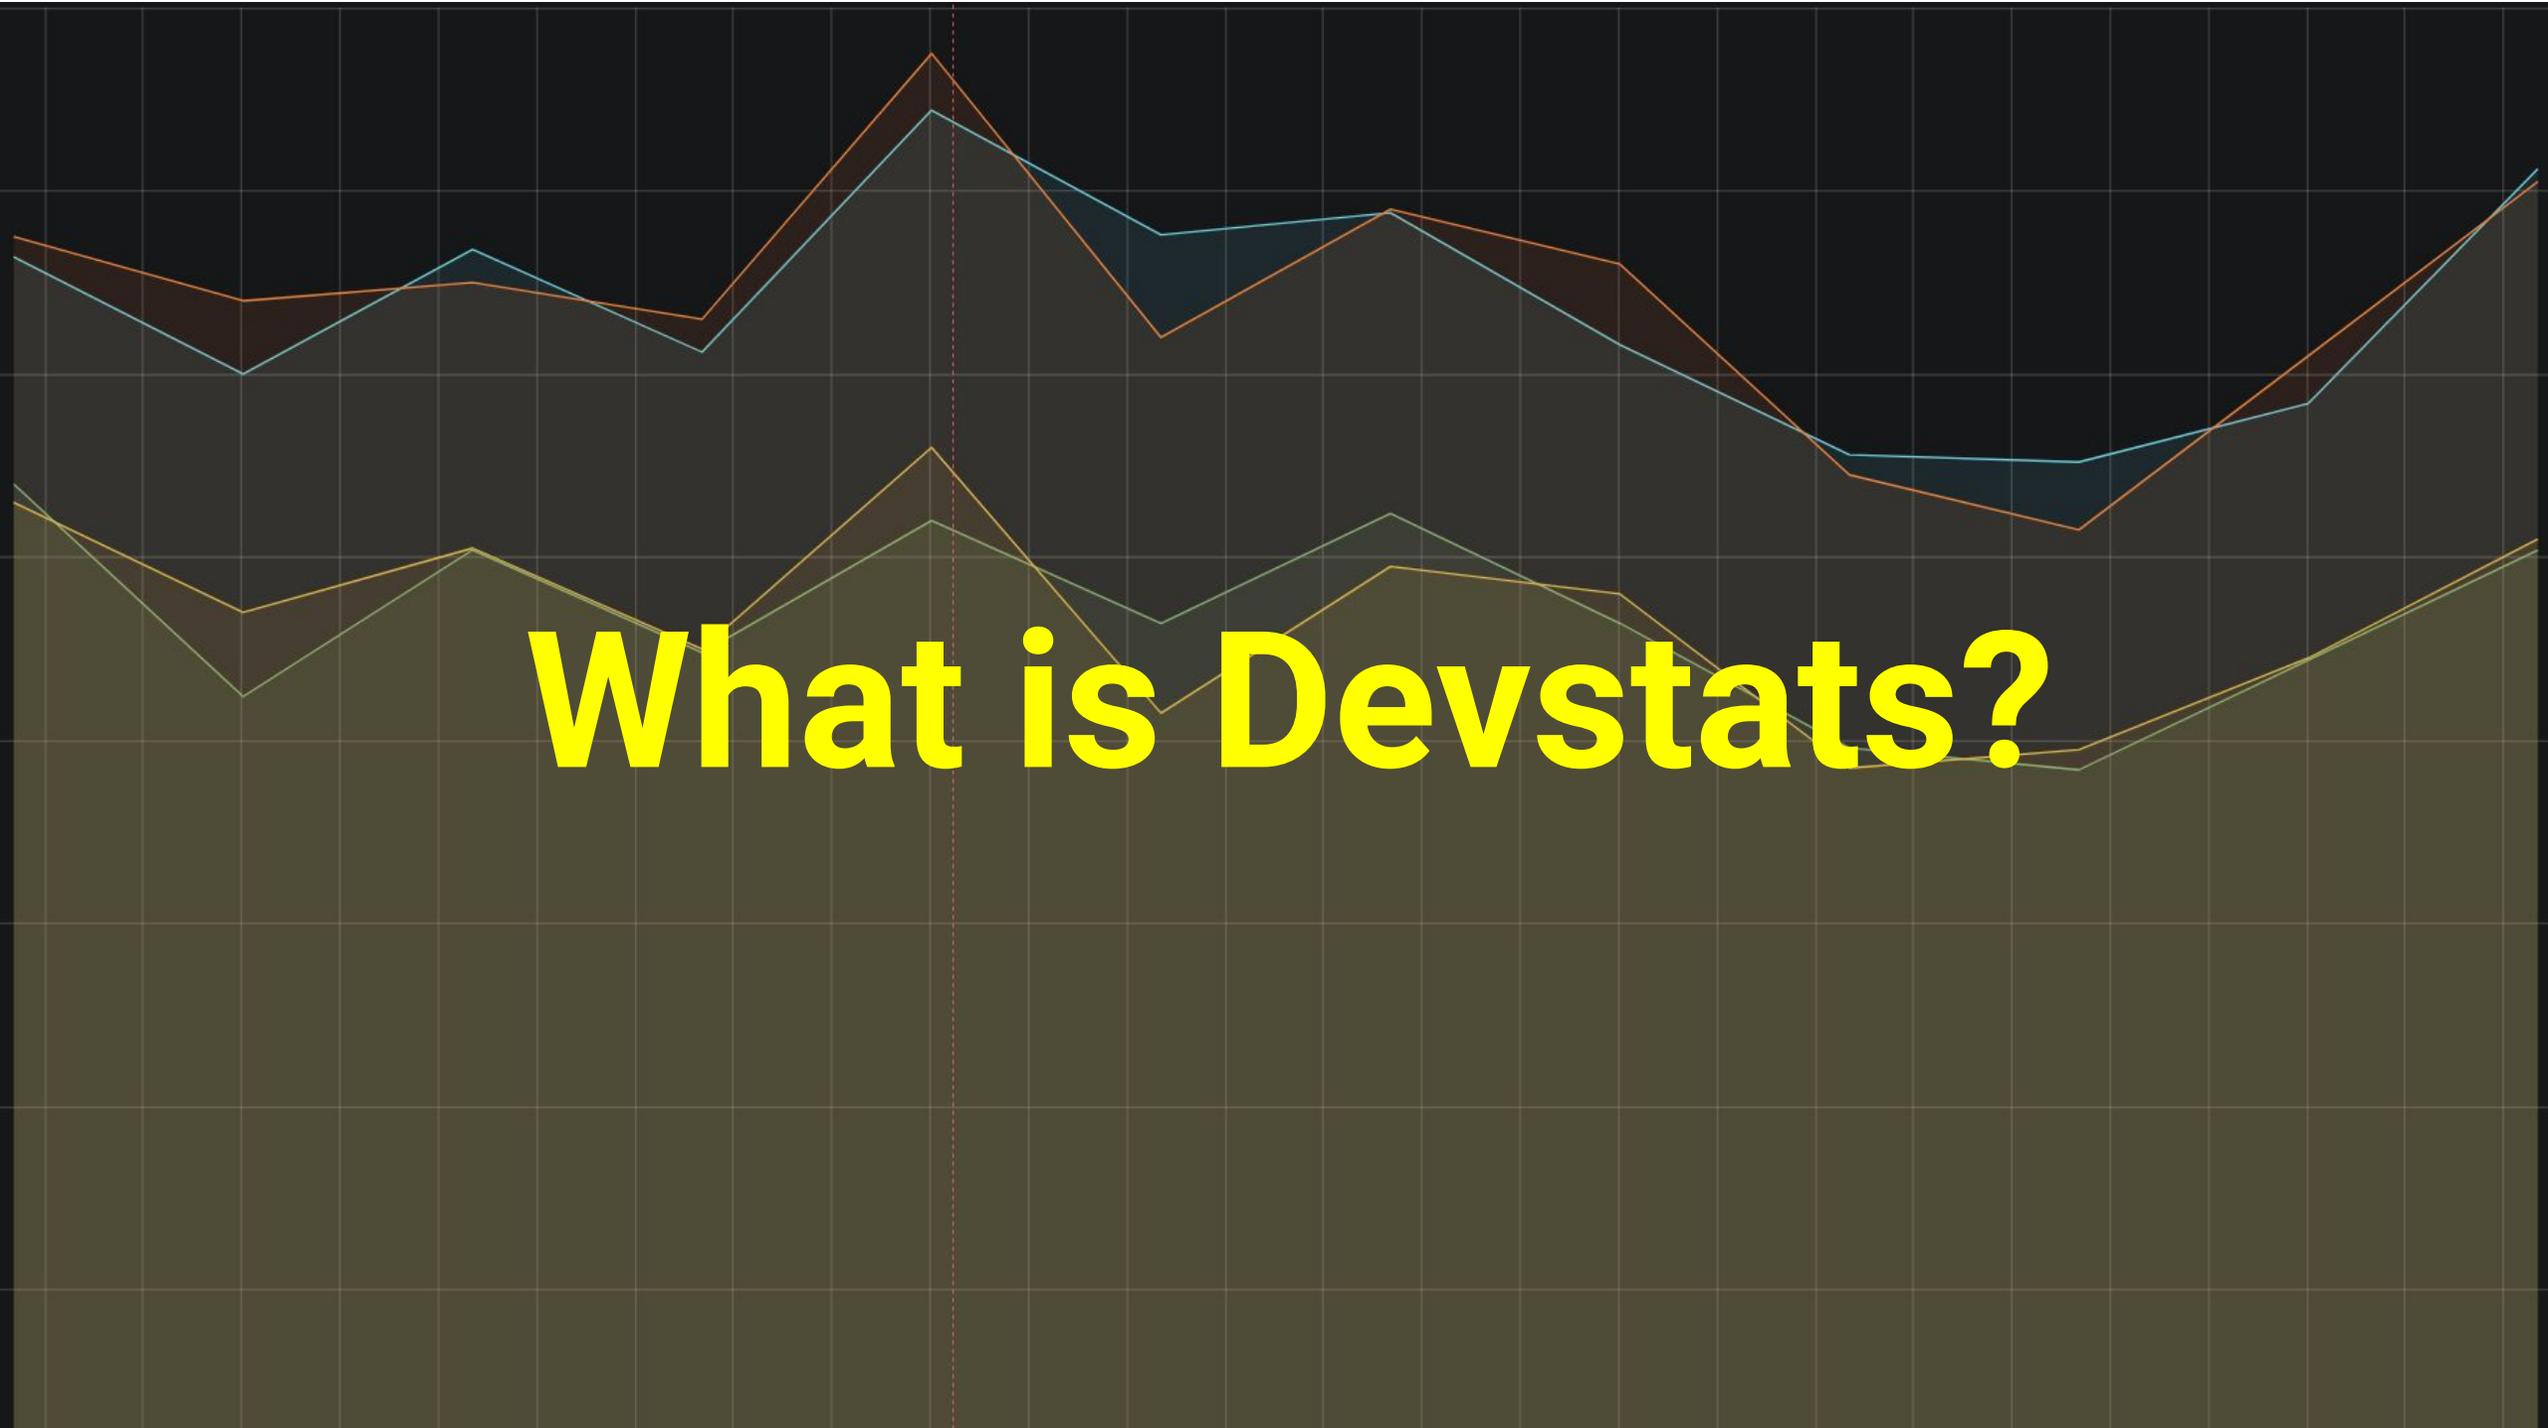
- Guestbook

- A couple replicaset
- A secret
- A couple services

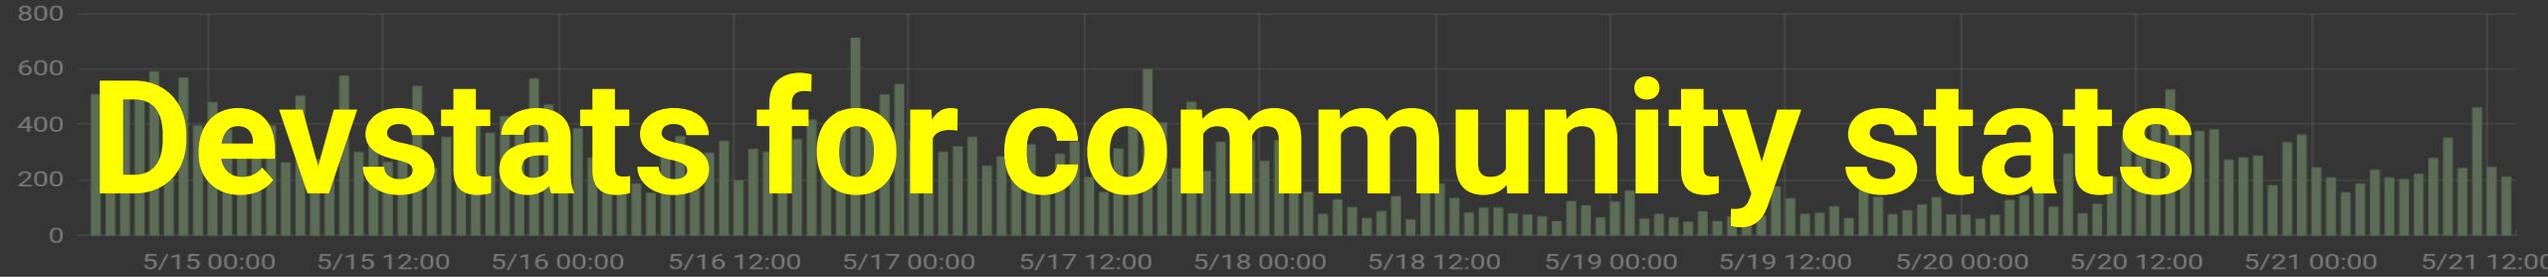


- Guestbook

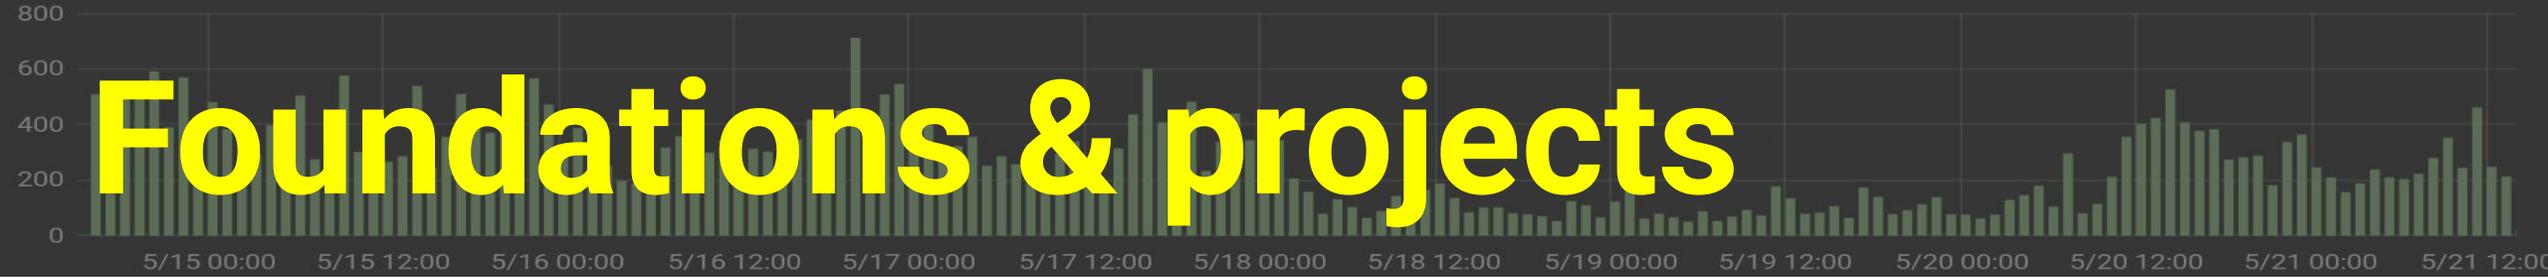
- A couple replicaset
- A secret
- A couple services
- no Ingress
- non-persistent database
- no Helm/Operator
- no Jobs
- no CI/CD
- not current for 1.13
- no SSL

A line chart with a dark background and a light gray grid. The chart features four data series: a blue line, an orange line, a yellow line, and a green line. The blue and orange lines are positioned higher on the y-axis, while the yellow and green lines are lower. All lines show a similar trend, with a notable peak around the center of the chart. A vertical dashed orange line is drawn at the position of the peak. The text 'What is Devstats?' is overlaid in the center in a large, bold, yellow font.

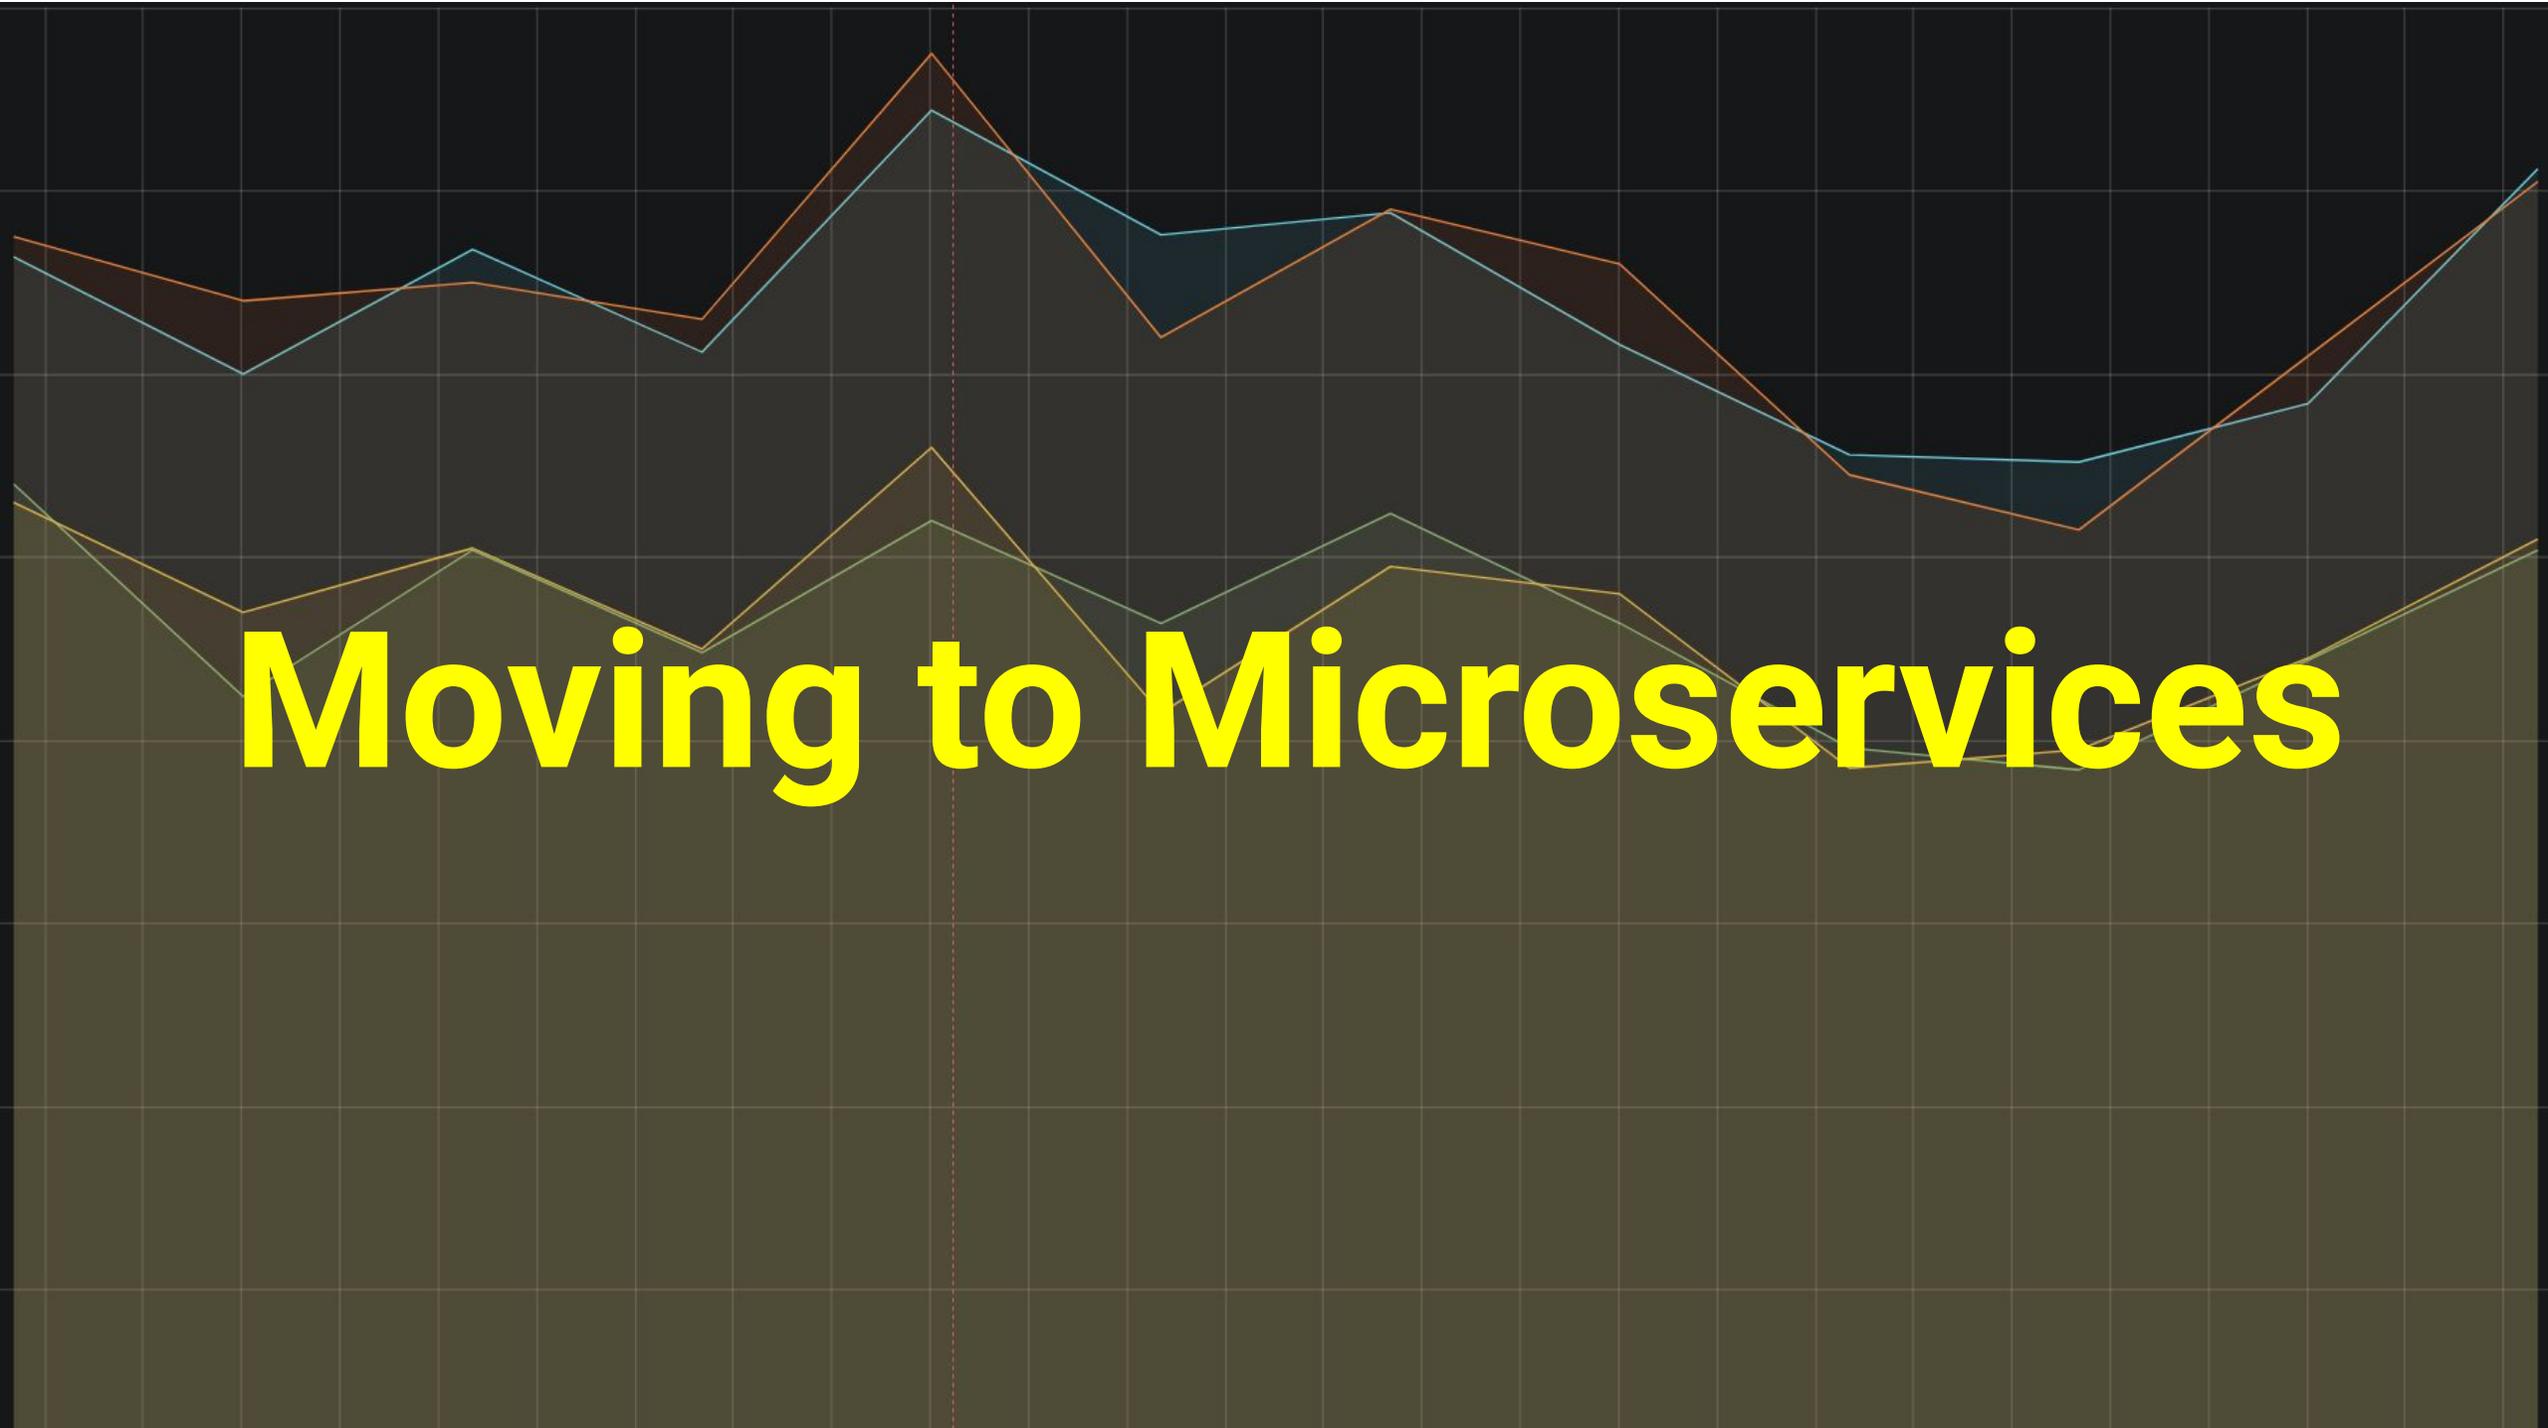
What is Devstats?



1. A service that takes data from Git and Github and turns it into graphs reporting community activity.
2. A CNCF/LF project & service for all their supported projects and foundations.
3. And Now: an example Kubernetes application.



- Each Foundation has its own Instance of Devstats
- Foundations support multiple Projects that share a Devstats Instance;
- Each Project has one or more Github Orgs, Repositories that are included in its stats.



The image features a line chart with a dark background and a light gray grid. There are four data series represented by lines of different colors: orange, light blue, yellow, and light green. The orange and light blue series are positioned higher on the chart, while the yellow and light green series are lower. A vertical dashed orange line is drawn at approximately the 35% mark on the x-axis. The text 'Moving to Microservices' is overlaid in the center in a large, bold, yellow font.

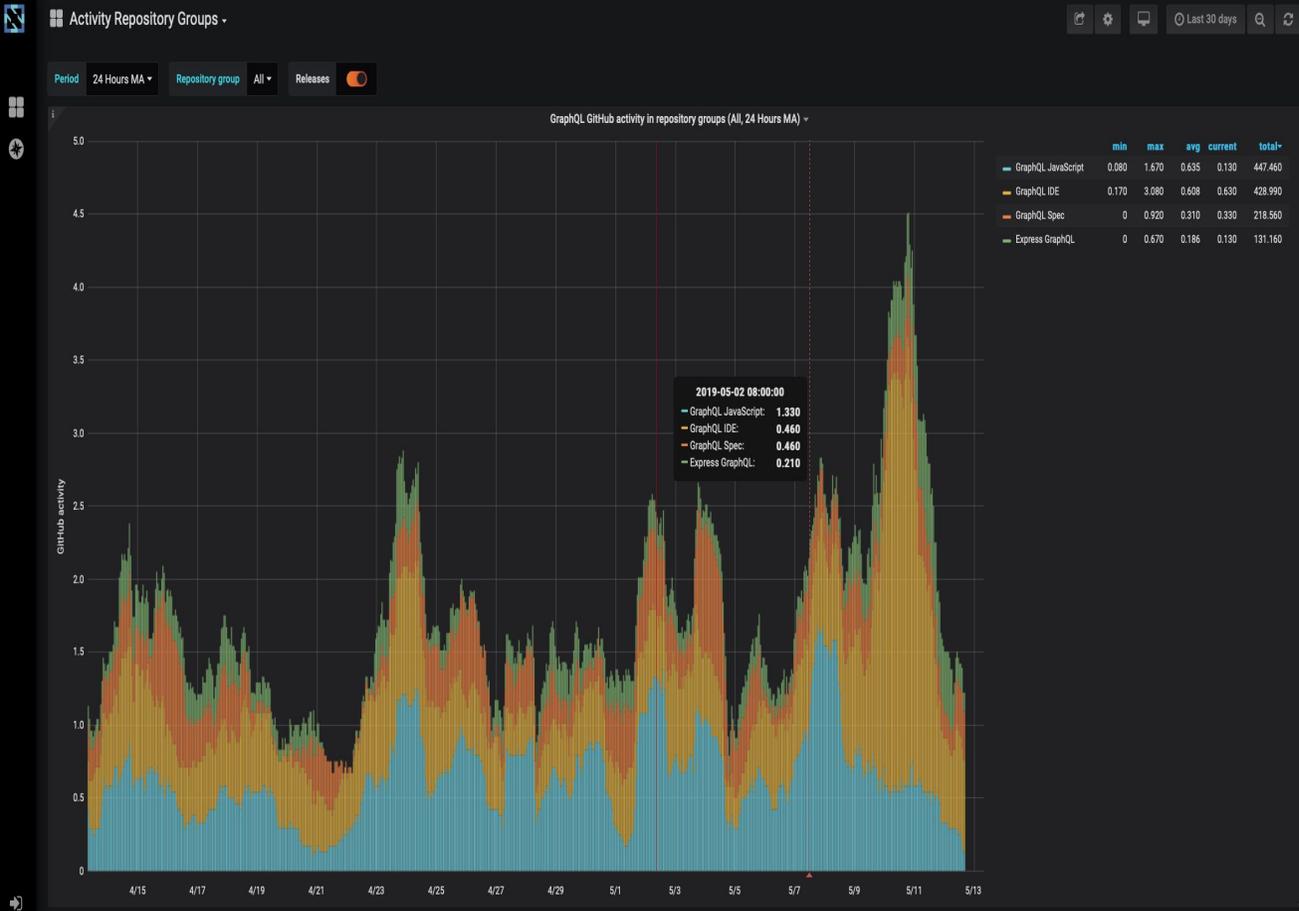
Moving to Microservices



- Two 48-core Packet.net servers running Ubuntu
- 1 Postgres instance with many databases
- Apache Proxy + LetsEncrypt
- 1 Grafana per project, routed by hostname
- Local disk for git/DB storage
- Hourly sync cron job



Hey, we're looking to install an instance of Devstats for our projects. Can you refer me to someone who can help me install it?



Projects Health Table

DevStats Projects

Metric	Projects			
	GraphQL JavaScript	GraphQL IDE	Express GraphQL	GraphQL Spec
Commits: Days since last commit	3 days	0 days	1 days	173 days
Commits: Last commit date	05/08/2019 05:51:00 pm	05/11/2019 05:31:11 pm	05/11/2019 05:45:00 am	11/20/2018 02:41:47 pm
Commits: Number of commits in the last 12 months	559	58	20	59
Commits: Number of commits in the last 3 months	131	33	6	0
Commits: Number of commits in the last 3 months (previous 3 months)	176	4	0	3
Commits: Number of commits in the last 3 months vs. previous 3 months	↓	↑	↑	↓
Commits: Number of commits in the last 6 months	307	37	6	3
Committers: Number of committers in the last 12 months	23	19	5	4
Committers: Number of committers in the last 3 months	5	12	4	0
Committers: Number of committers in the last 3 months (previous 3 months)	11	2	0	2
Committers: Number of committers in the last 3 months vs. previous 3 months	↓	↑	↑	↓
Committers: Number of committers in the last 6 months	13	13	4	2
Companies: Percent of all commits authors from top committing company (last 3 months)	0.00% Apple	-	-	-
Companies: Percent of all commits authors from top committing company (last year)	0.18% Apple	-	-	-
Companies: Percent of all commits from top committing company (last year)	-	-	-	-
Companies: Percent of all commits from top committing company (previous 3 months)	-	-	-	-
Companies: Percent of all commits pushers from top committing company (last 3 months)	-	-	-	-
Companies: Percent of all commits pushers from top committing company (last year)	-	-	-	-
Companies: Percent of known commits authors from top committing company (last 3 months)	-	-	-	-
Companies: Percent of known commits authors from top committing company (last year)	100.00% Apple	-	-	-



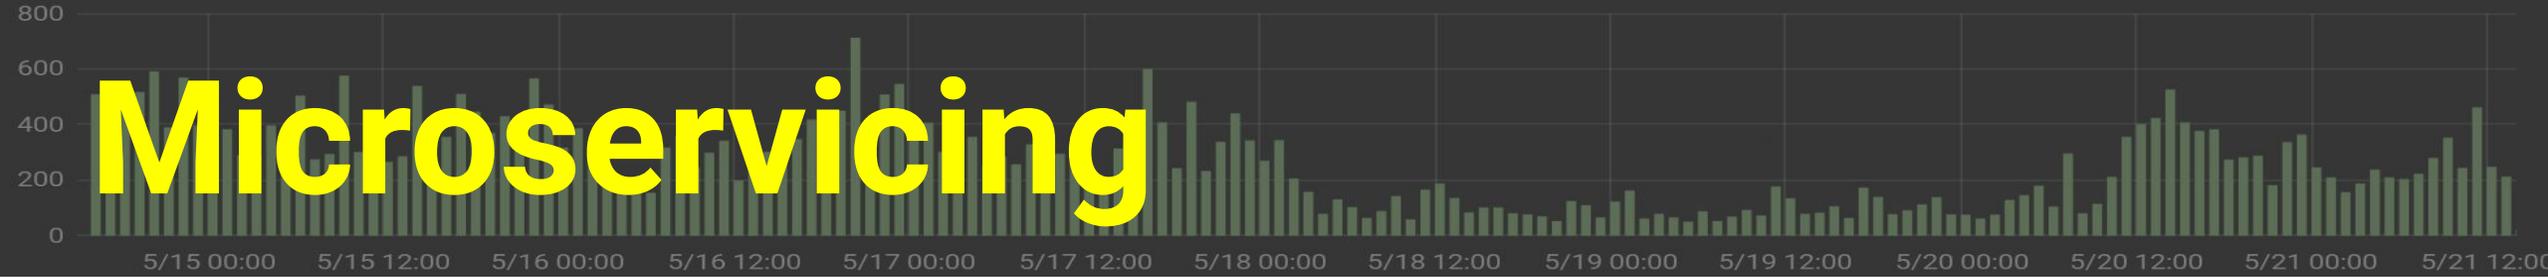
New projects/foundations

etcd, Buildpacks, Falco,
Dragonfly, Virtual
Kubelet, KubeEdge,
Brigade, CRI-O, Network
Service Mesh,
OpenEBS

OpenTelemetry,
Continuous Delivery
Foundation, GraphQL
Foundation, Kubeflow,
Istio



EYODF



Microservicing

- Switch from local storage to EBS volumes
- Containerize devstats image, grafana image, hourly sync job, tests
- Replace local Postgres DB with Patroni cluster
- Separate RO and RW DB connections
- Create Helm templates for everything (including tests)



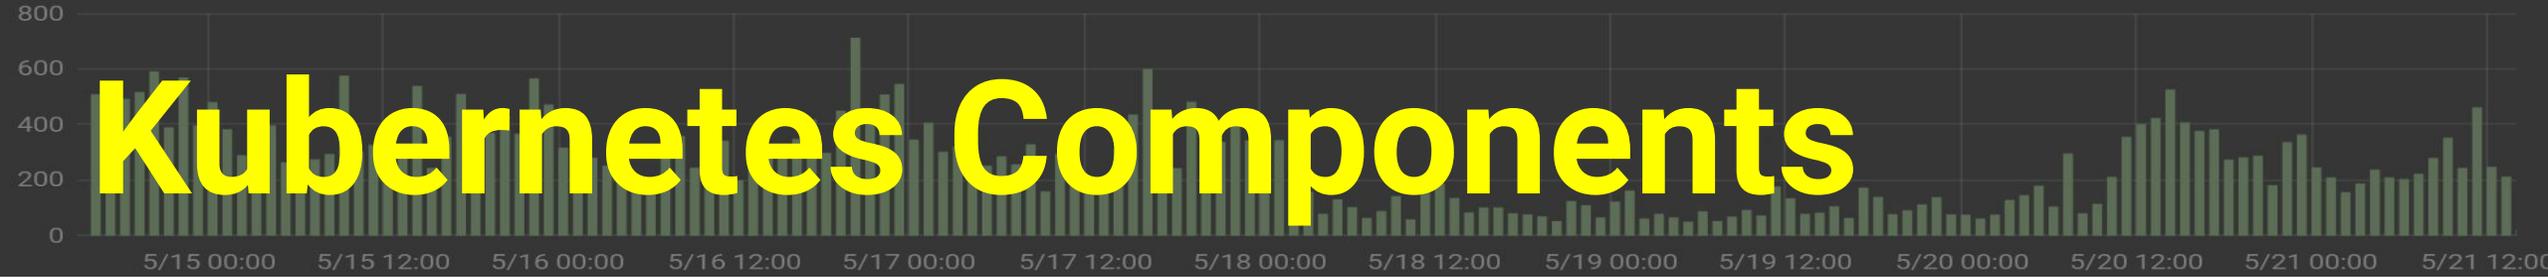
Microservicing Follies

- Look out for Docker shared memory limits
- Permissions on Persistent Volumes
- More nodes > larger instances
- Learning curve for EKS/IAM



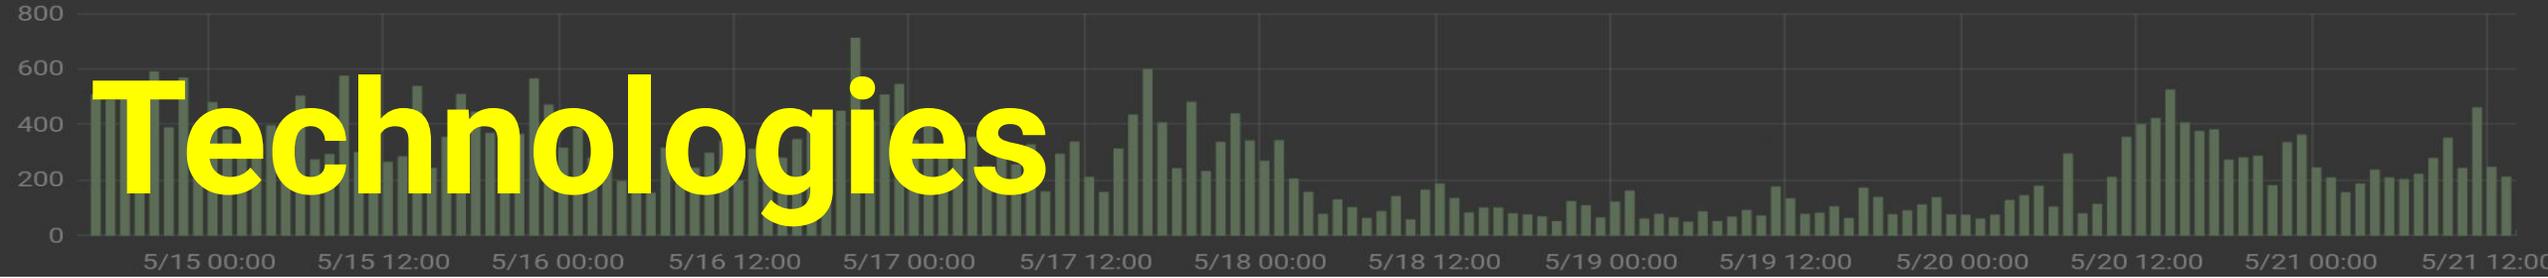
The image features a line chart with a dark background and a light gray grid. There are four data series represented by lines of different colors: orange, cyan, yellow, and green. The orange and cyan lines are positioned higher on the chart, while the yellow and green lines are lower. A vertical dashed orange line is drawn at approximately the 35% mark on the x-axis. The word 'Architecture' is written in a large, bold, yellow font across the middle of the chart, overlapping the lines.

Architecture

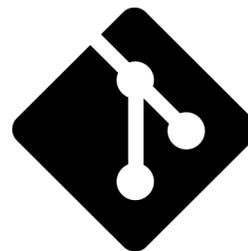
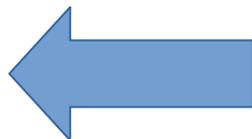
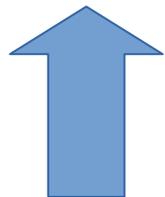
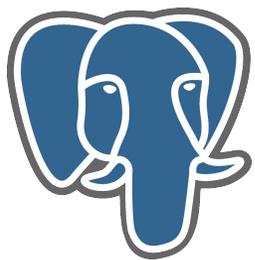
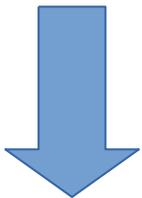


- Ingress
- Services
- StatefulSet
- ReplicaSets
- Volumes
- CronJobs

- Secrets
- Certificates
- Helm
- Tests
- RBAC



- Ingress-ELB, Ingress Nginx
- AWS Cloud Provider
- Grafana
- Patroni Postgres HA
- Travis CI
- GitHub
- Docker Hub
- AWS
- LetsEncrypt CertManager



cncf / [devstats-docker-images](#) Unwatch 4 ★ Star 2 Fork 3

[Code](#) [Issues 0](#) [Pull requests 0](#) [Projects 0](#) [Wiki](#) [Insights](#) [Settings](#)

Branch: [master](#) [devstats-docker-images / README.md](#) Find file Copy path

 [lukaszgrylicki](#) Update README.md ad9c650 just now

[1 contributor](#)

31 lines (16 sloc) | 1.05 KB Raw Blame History   

devstats-docker-images

DevStats docker images: minimal (hourly cron job sync), full (provisioning/bootstrapping), Grafana (UI endpoint), Patroni (database), tests.

Create and test images

Create and remove docker images:

- To create DevStats docker container images and publish them, use: `DOCKER_USER=... ./images/build_images.sh`.
- To drop local DevStats docker container images use: `DOCKER_USER=... ./images/remove_images.sh`.

Containers Repo

lukaszgryglicki Update docs 3

67ae0d4 3 days ago

1 contributor

205 lines (140 sloc) | 14.9 KB

Raw

Blame

History



devstats-helm-example

DevStats Deployment on Kubernetes using Helm. This is an example deployment of few CNCF projects.

Helm chart in `devstats-helm-example`.

Charts Repo

Branch: master devstats / README.md

Find file Copy path

lukaszgryglicki Update README.md 7c3ceab just now
2 contributors

135 lines (74 sloc) 6.82 KB

Raw Blame History

build passing ci best practices passing

GitHub archives and git Grafana visualization dashboards

Authors: Łukasz Gryglicki lgryglicki@cncf.io, Justyna Gryglicka lgryglicka@cncf.io.

This is a toolset to visualize GitHub [archives](#) using Grafana dashboards.

GHA2DB stands for GitHub Archives to DashBoards.

More information about Kubernetes dashboards [here](#).

Kubernetes and Helm

Please see [example Helm chart](#) for an example Helm deployment.

Please see [LF Helm chart](#) for the LF Helm deployment (it is a data deployment, has no Grafana and uses Elasticsearch in addition to Postgres to store data).

Please see [GraphQL Helm chart](#) for GraphQL foundation DevStats deployment.

Please see [bare metal example](#) to see example bare metal deployment.

Presentations

- Presentations are available [here](#).

Branch: master devstatscode / README.md

Find file Copy path

lukaszgryglicki Updated docs [wip] 1797f0b a minute ago
1 contributor

19 lines (12 sloc) 1.43 KB

Raw Blame History

build passing ci best practices passing

DevStats code

This is a code repository for DevStats used to display [CNCF projects dashboards](#), [CDF projects dashboards](#), [GraphQL projects dashboards](#) and [example Kubernetes/helm deployment](#).

Authors: Łukasz Gryglicki lgryglicki@cncf.io, Justyna Gryglicka lgryglicka@cncf.io, Josh Berkus jberkus@redhat.com.

Building and installing

- Follow [this guide](#) to see how to deploy on Kubernetes using Helm.
- Follow [this guide](#) to see GraphQL deployment using Kubernetes & Helm.
- Follow [this guide](#) for installing on bare metal.
- Follow [this guide](#) to deploy your own project on bare metal (this example deploys Homebrew statistics).
- Fetch dependency libraries.
- `make` then `make test` finally `make install`.

DevStats Repo

Helm charts

A bit complicated:

- variables to omit components (such as DB)
- loop for deploying variable numbers of projects

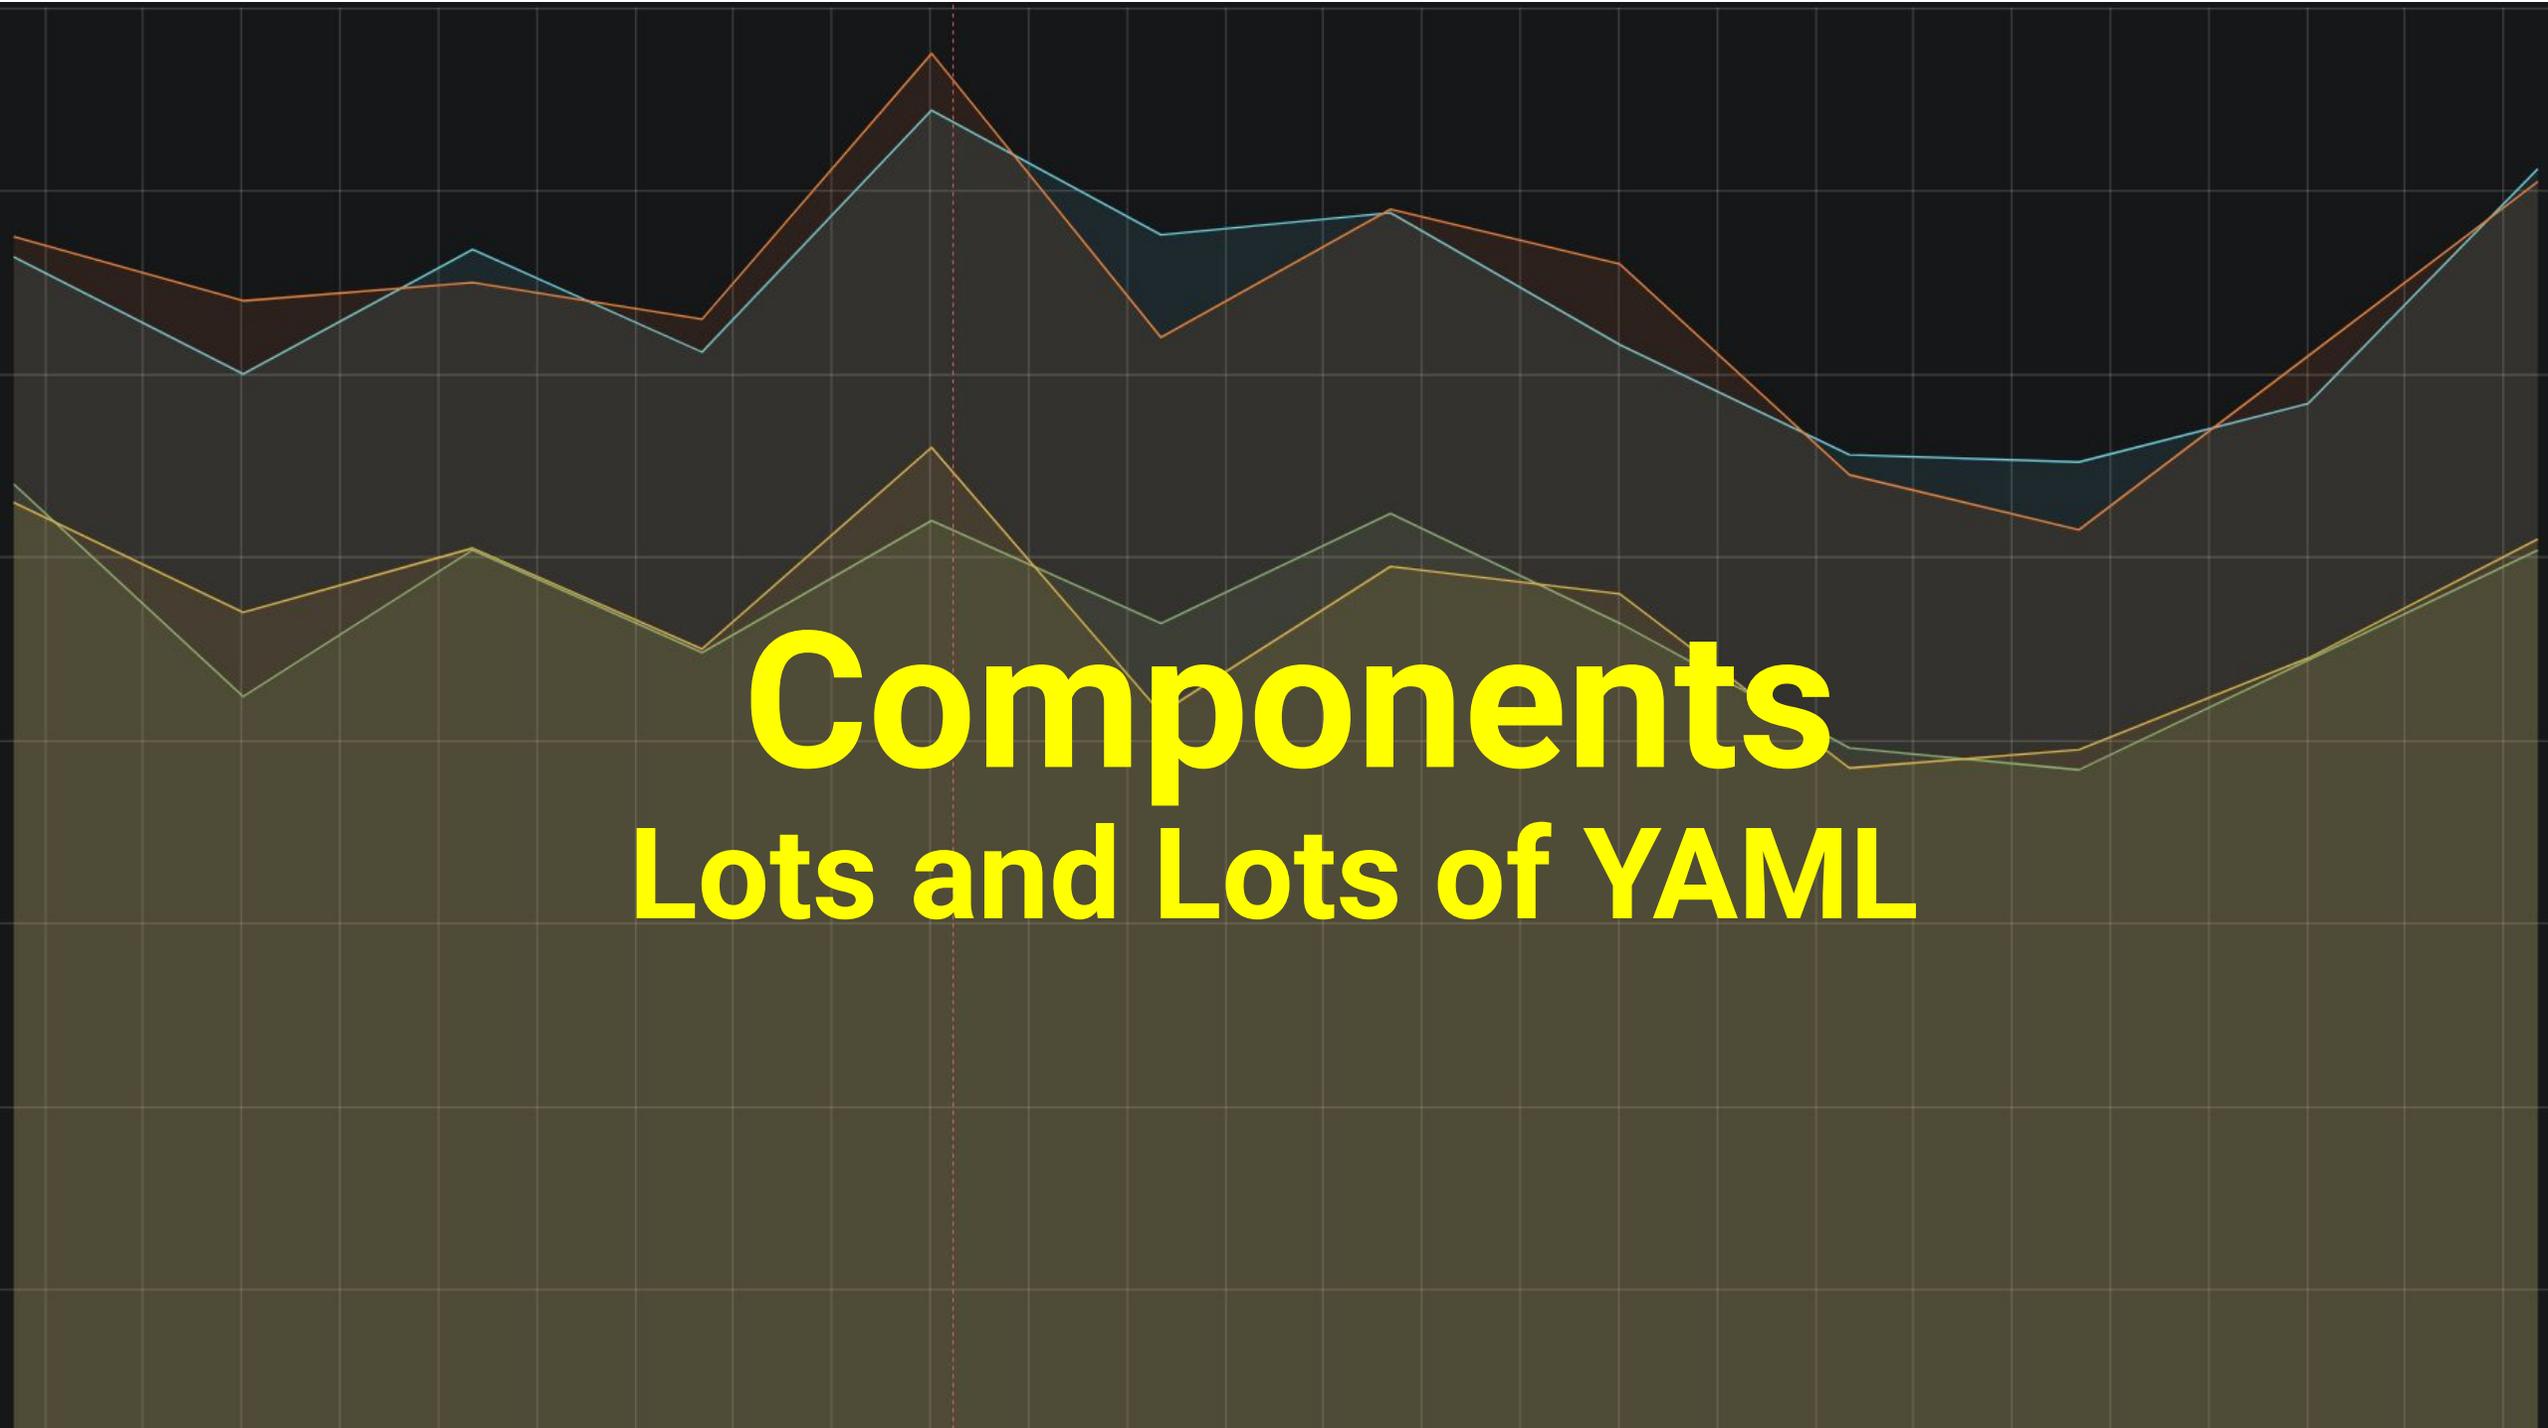
Maybe should use an operator



Deployment platform

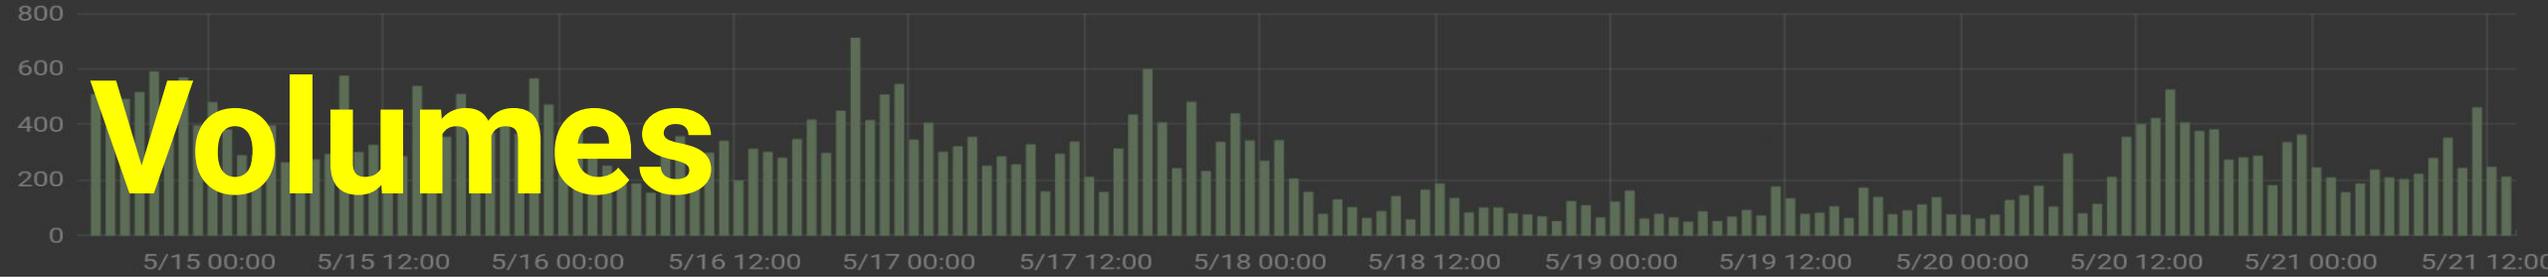
- Kubernetes 1.12
- Amazon EKS
 - Fully portable though
- Advantages:
 - Experience with AWS
 - Easy handling of storage, Ingress
- Move to bare metal?





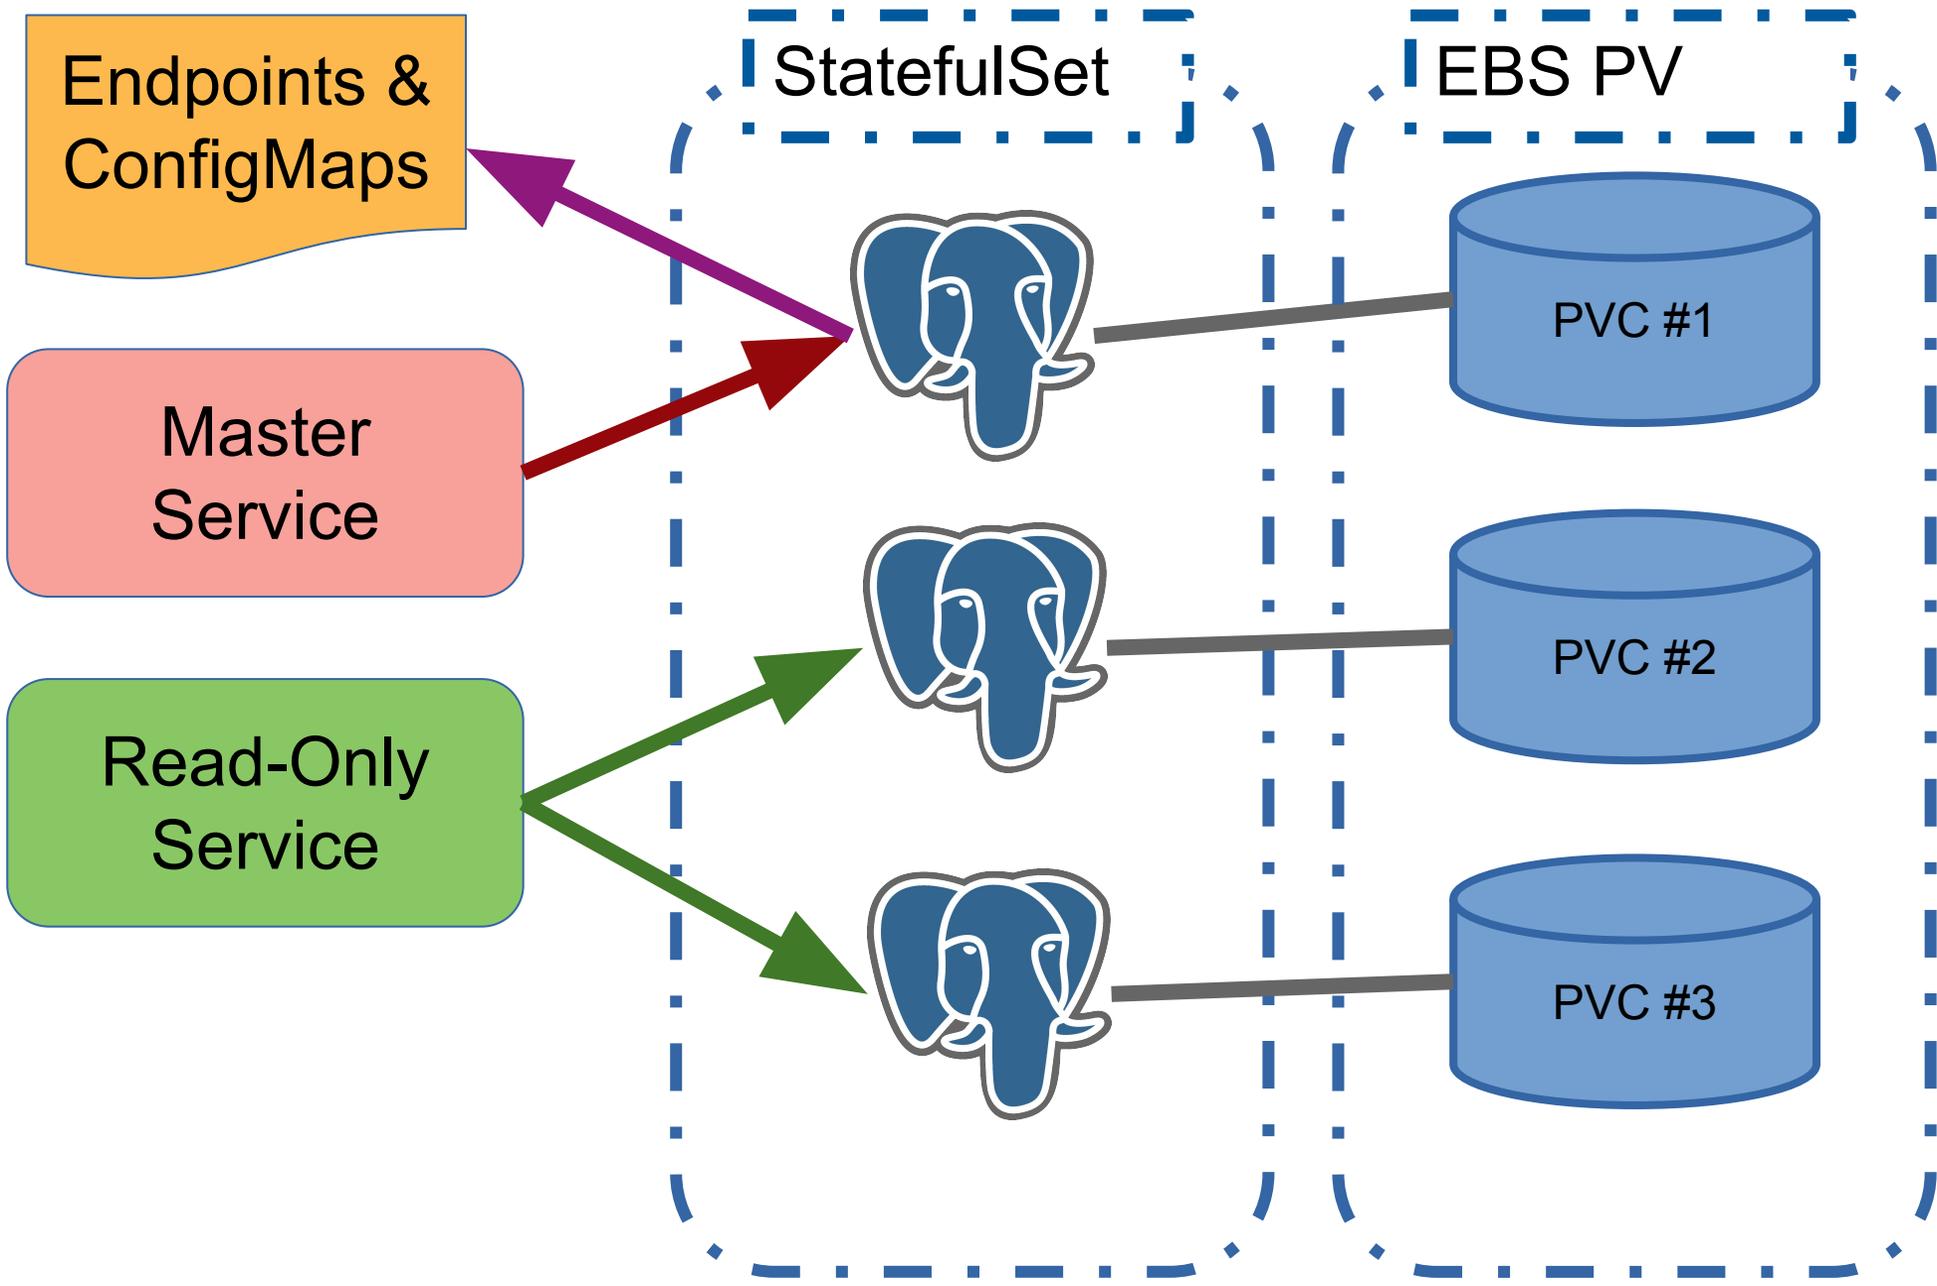
Components

Lots and Lots of YAML



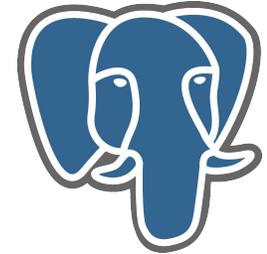
3 Volumes

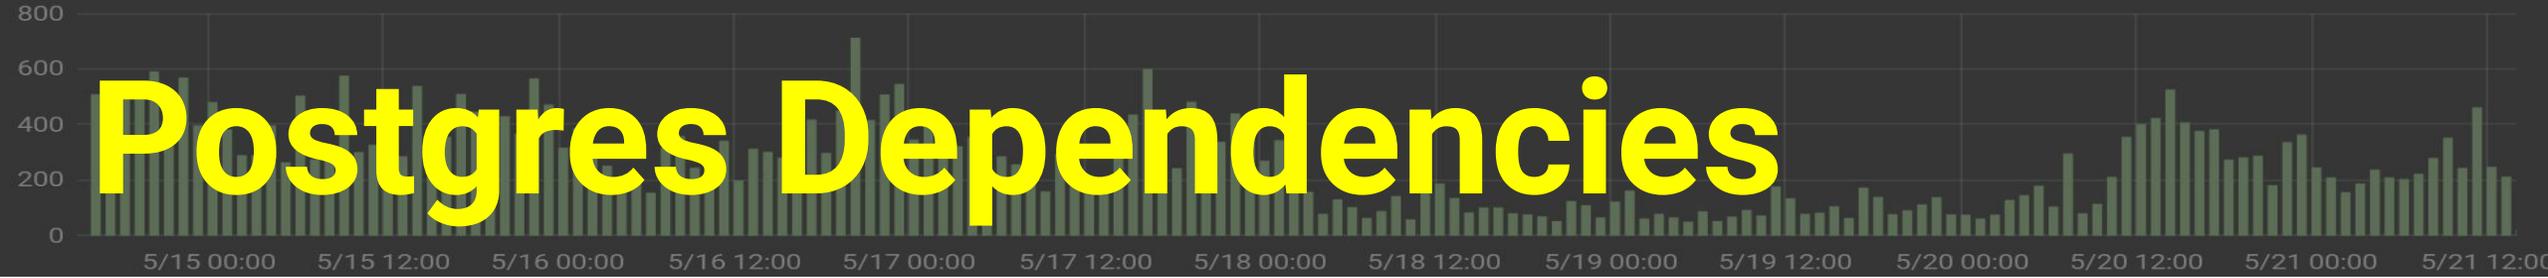
- 1 for provisioning
 - temporary
 - for **each** project git clones storage
- 1 for hourly sync, temporary
- 1 for Database nodes, permanent





- StatefulSet
- Uses Patroni for automated HA
 - (current example includes Anti-Affinity to make sure each patroni pod runs on a different node)
- 1 database instance shared by all Projects
- Uses PVT





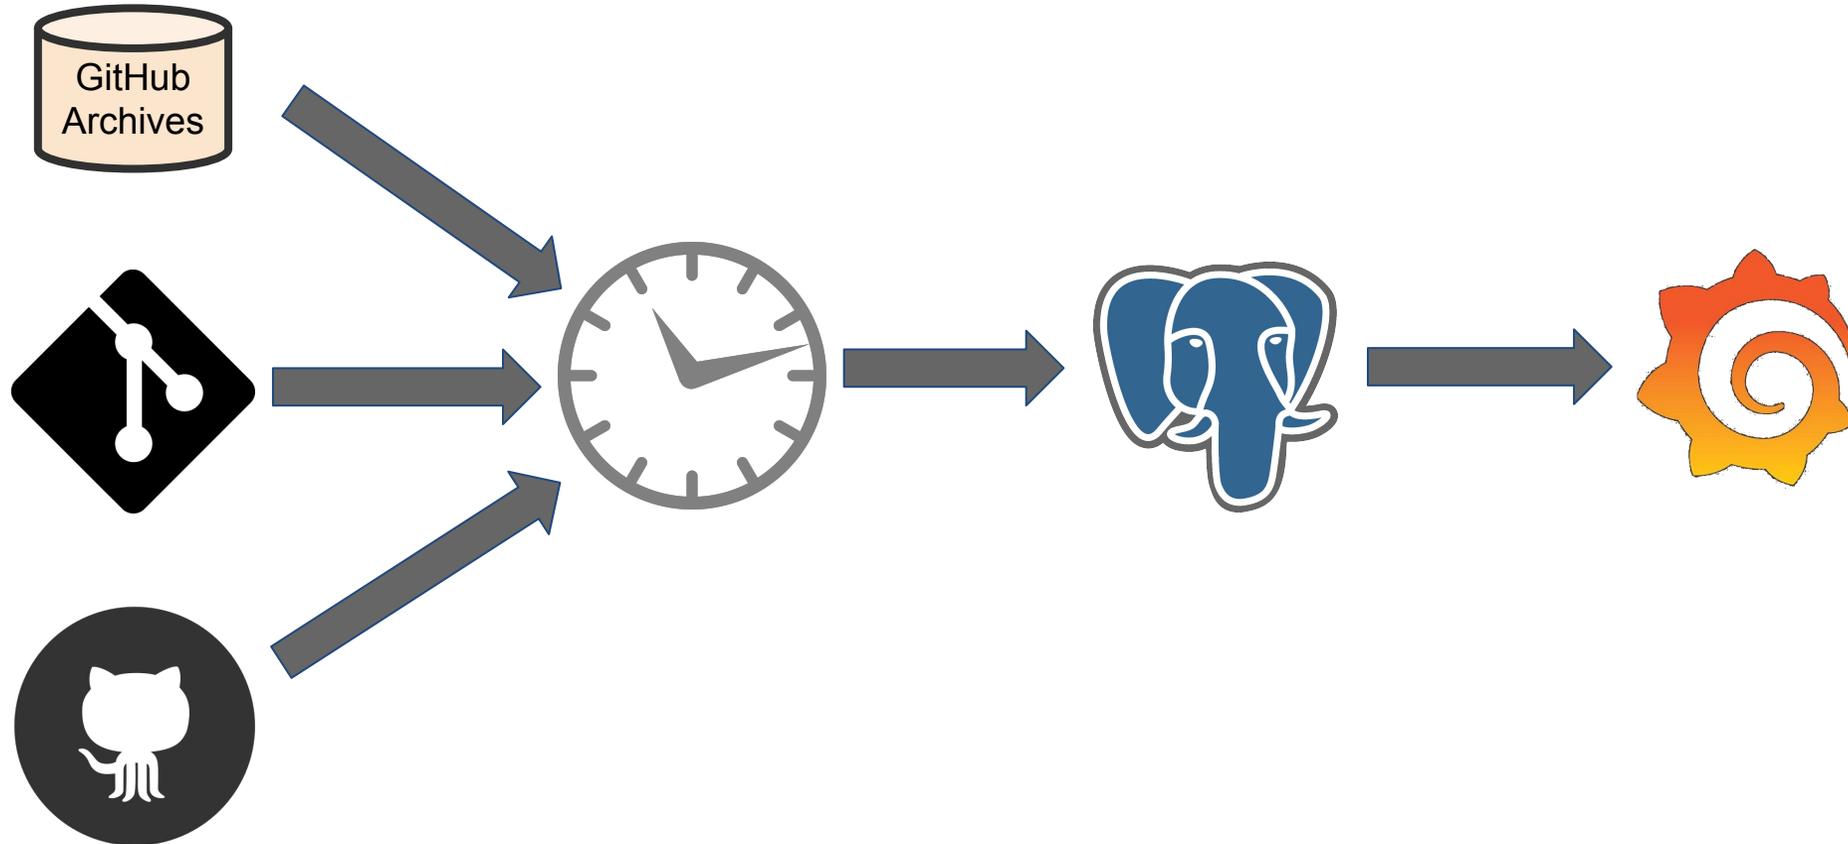
- RBAC

- Patroni needs special roles to work HA magic

- Services

- 1 Master Service for Read-Write
- 1 General Service for Read-Only load balancing

Where the data comes from





- One job per Project
- Pulls data from GitHub/Archives, API and from git
- Pushes data into Postgres
- Runs timeseries-building go programs



- One per hosted project
 - Use loop in Helm template to create sets
- Graph designs from Git
 - add a graph by generating new Grafana image (manually)



Secrets: Postgres & Grafana

- Secrets are loaded by Helm from files
 - (not contained in chart)
- Multiple DB connections



- Nginx Ingress provider
- LetsEncrypt service for certs
 - cert-manager CRD
- Generates subdomain of “devstats-demo.net” per Project

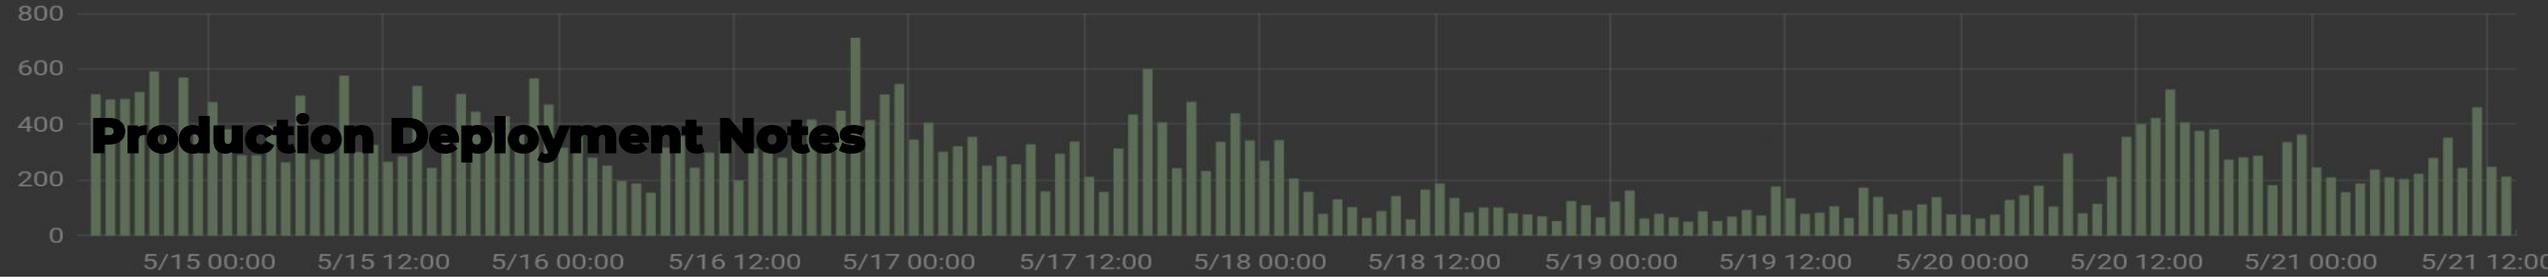


- Builds Code from Devstats Repos
- Builds Images from Images Repo
- Manual triggering required for deployment

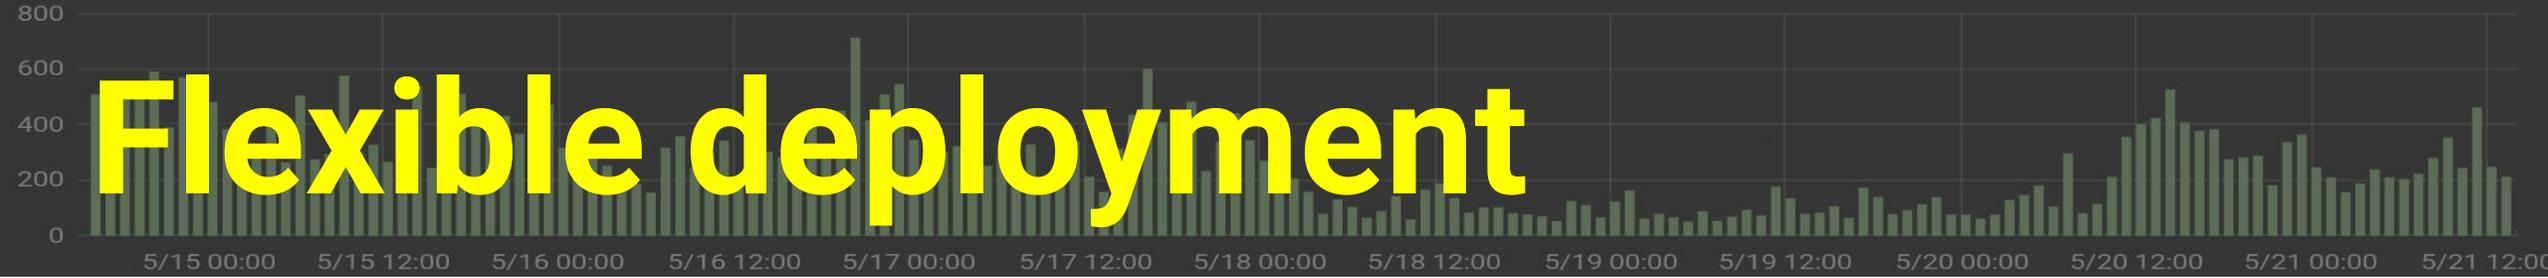
A line chart with a dark background and a light gray grid. It features four data series: a blue line, an orange line, a yellow line, and a green line. All series show a similar trend with a significant peak at a specific point marked by a vertical dashed orange line. The blue and orange lines are the highest, while the yellow and green lines are the lowest. The text 'Running Devstats' and 'Production Setup Steps' is overlaid in the center in a bright yellow font.

Running Devstats

Production Setup Steps



- register a domain
- use ingress
 - with a cloud provider
- re-generate certificates
 - should attach to domain

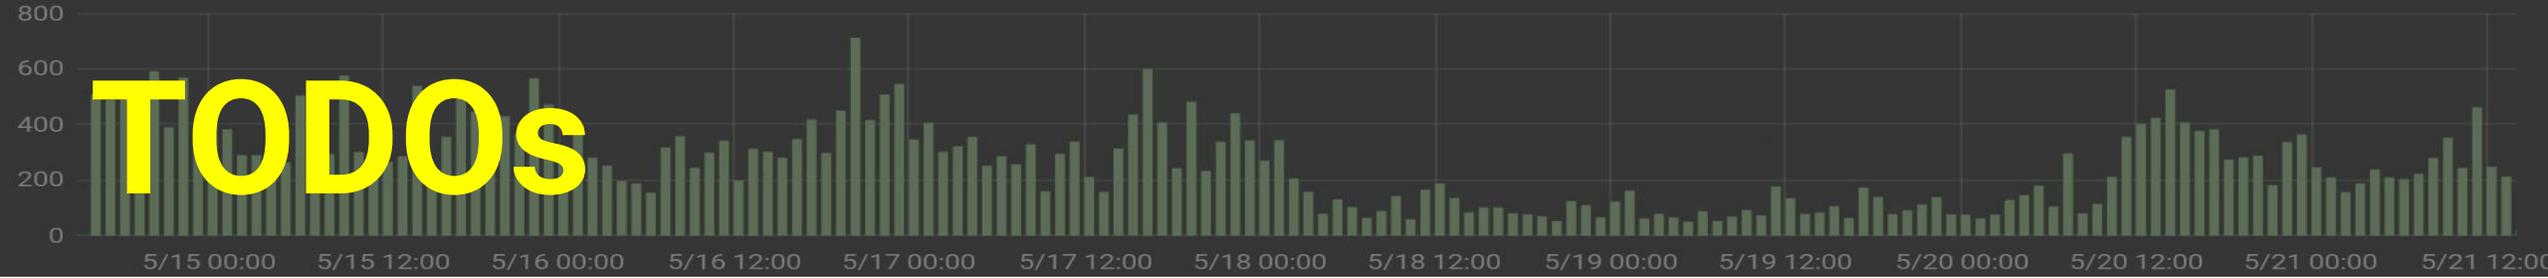


- Swap out Patroni for hosted Postgres
- Don't use Grafana (just have data API)
- ElasticSearch output

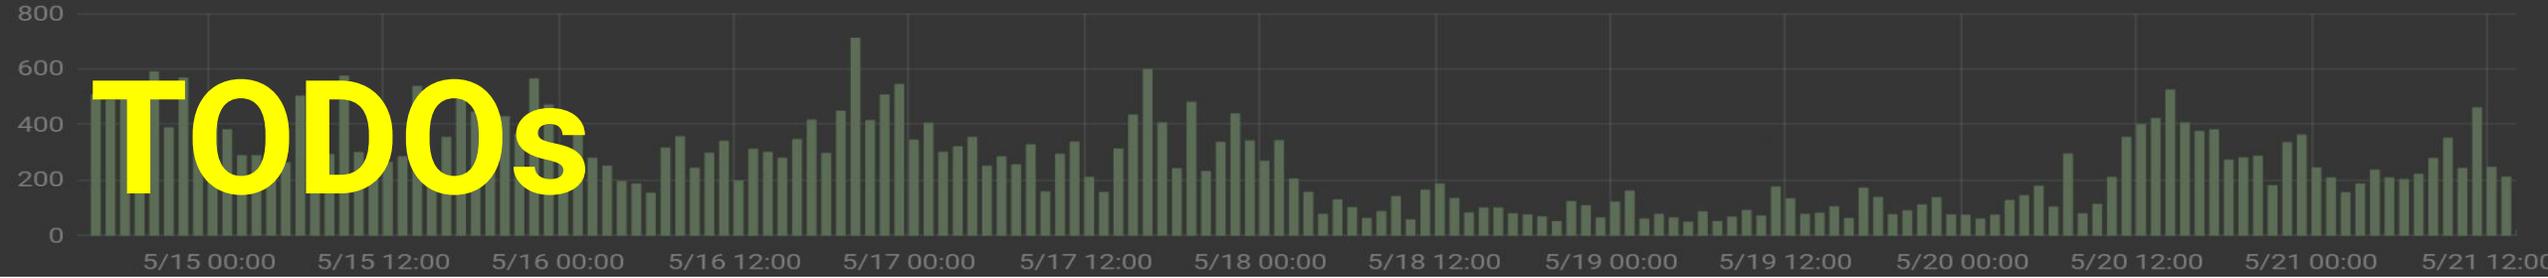


Future Plans

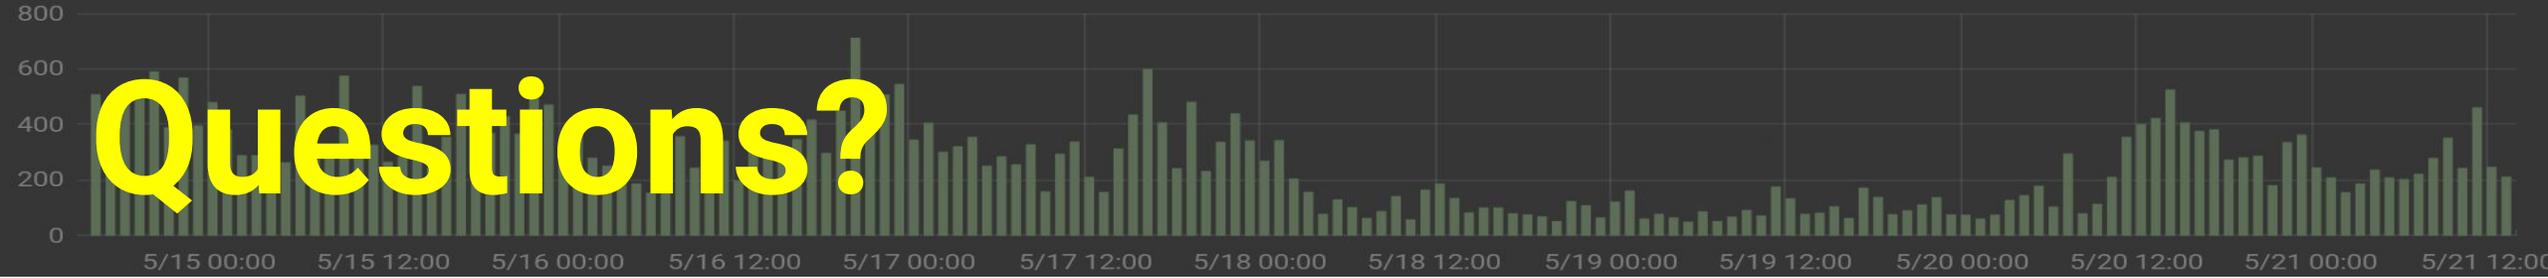
(stuff you can contribute to)



- Documentation!
- Register the Helm chart
- Add monitoring
- Create an operator
- CI/CD for Grafana
- More sophisticated secrets management



Get more users and contributors!



Questions?

- DevStats Repos:

- `gh/cncf/devstats`
- `gh/cncf/devstatscode`
- `gh/cncf/devstats-helm-example`
- `gh/cncf/devstats-docker-images`
- `gh/cncf/devstats-kubernetes-dashboard`

- `#devstats` on `k8s.slack.com`

`#devstats`

- Lukasz:

- lgryglicki@cncf.io
- lukaszgryglicki@o2.pl
- `@lukaszgryglicki`

- Josh:

- `jberkus@redhat.com`
- `@fuzzychef`