CERN openlab

Service Level Objectives : What's working and what's not ?

CERN Openlab Lightning Talks 2020

Felipe Bastos

Supervisor: Nikolay Tsvetkov

24/09/2020

Introduction

A reliable computer system must :

- Be available when needed
- Deliver results as quickly as possible
- Provide correct information

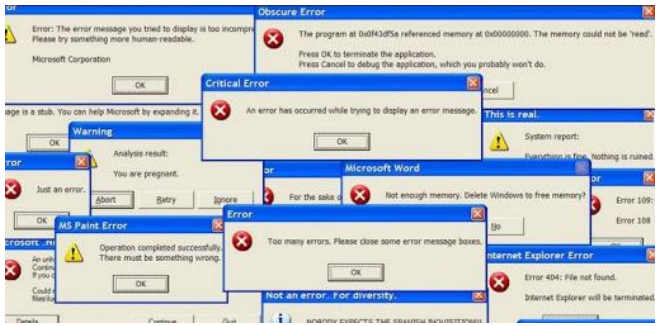

Site Reliability Engineering

Service Level Indicator – SLI

- Is the system reliability ?
- How long does it take to process ?
- Is the data being well processed?
- Are the results correct ?

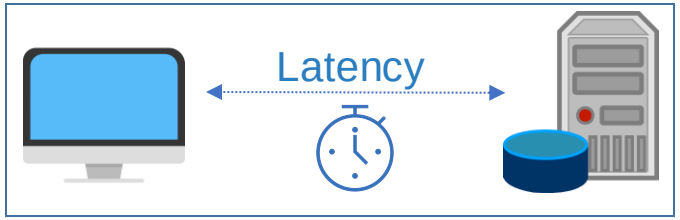

$$\text{SLI} = \frac{\text{successful events}}{\text{valid events}} * 100$$

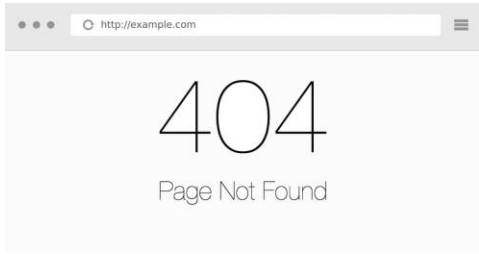

SLO

ERROR BUDGET

Site Reliability Engineering

SLI

Service Level Objective – SLO

- Defines target level for the SLIs.
- Is the system available enough?
- Is it fast enough?

e.g. lantency < 100ms for 95% responses

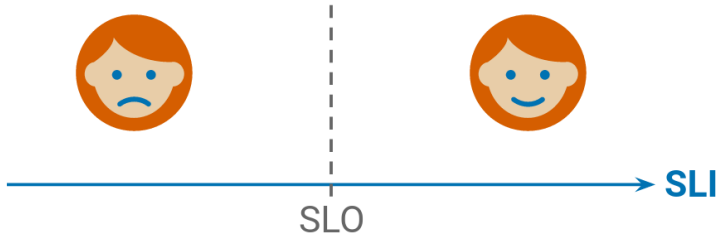

Error Budget

- How much error is acceptable in a certain period?
- Error Budget = 100% - SLO

e.g. The HTTP success rate can fail for 3 minutes per day

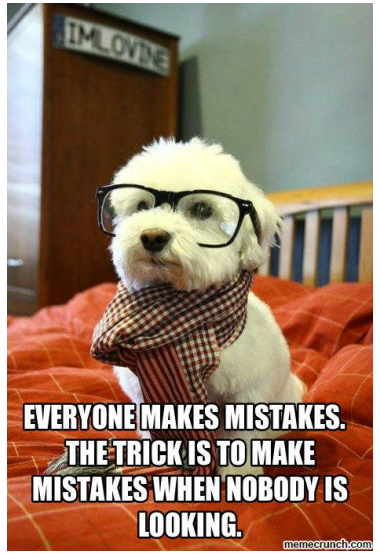

CERN's Monitoring SLIs

Data Access

- How long does it take to access those services?
- How often they fail?

Data Pipeline

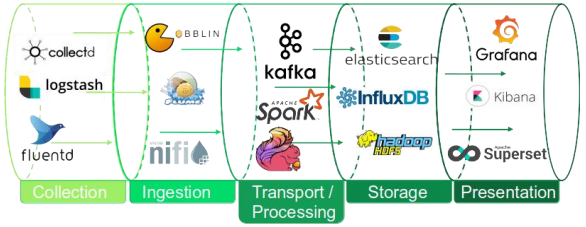

- Are incoming data being well stored?
- How fresh is that data?
- How long does it take to ingest the information?

Implementation

Data Access

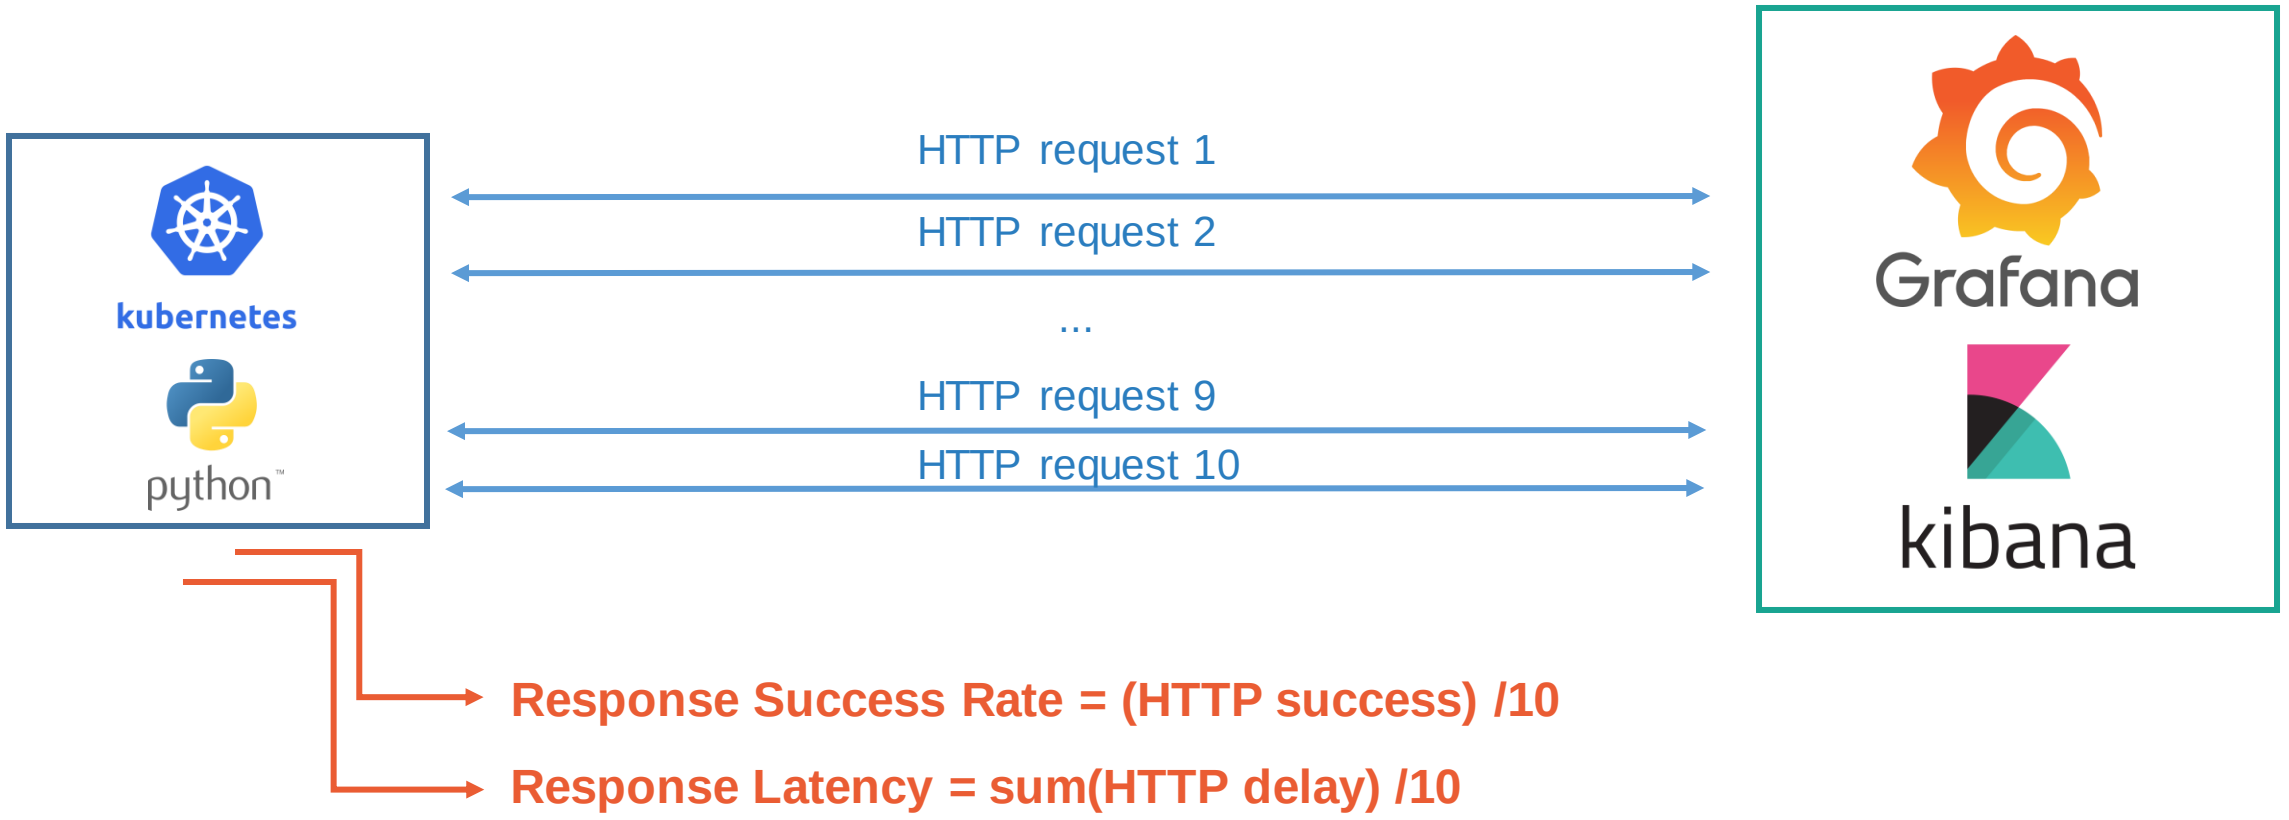

Implementation

Completness and Freshness

Send random documents every minute

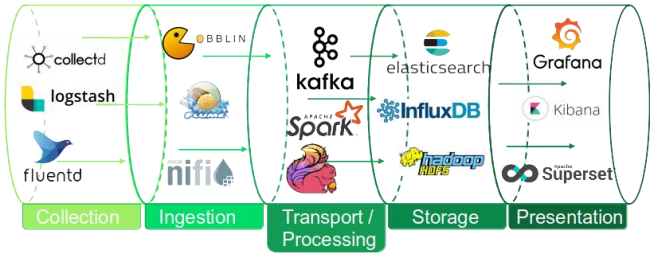

Count the number of stored documents

Check most recent document timestamp

Completness = (number of sent documents) – (number of stored documents)

Freshness = (actual time) – (last document timestamp)

Implementation

InfluxDB Integration

Continuous queries runs automatically and periodically to produce the resulted data

- Convert raw metrics into percentage for the SLIs and the Error Budget
- Resample every minute to keep the dashboard updated

$$\textbf{SLI} = \frac{(250 - \text{number of stored docs})}{250} * 100$$

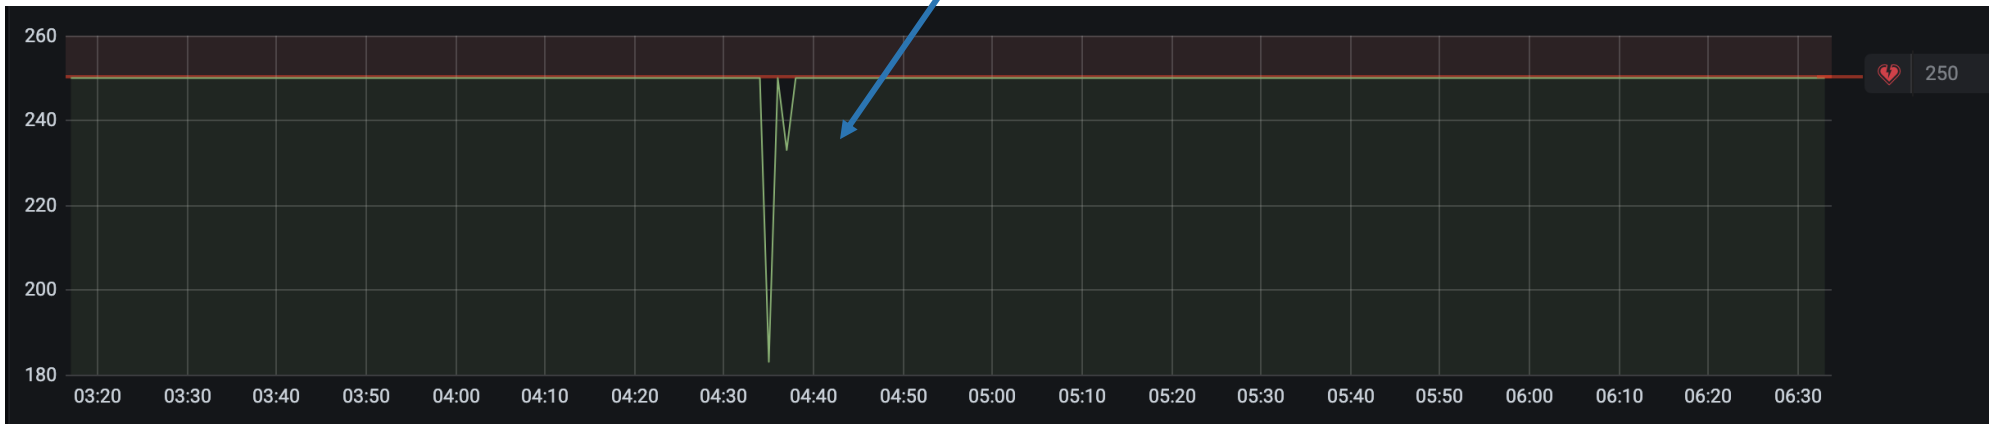

** we expect having 250 documents stored per minute*

Results

Dashboard & Alerts

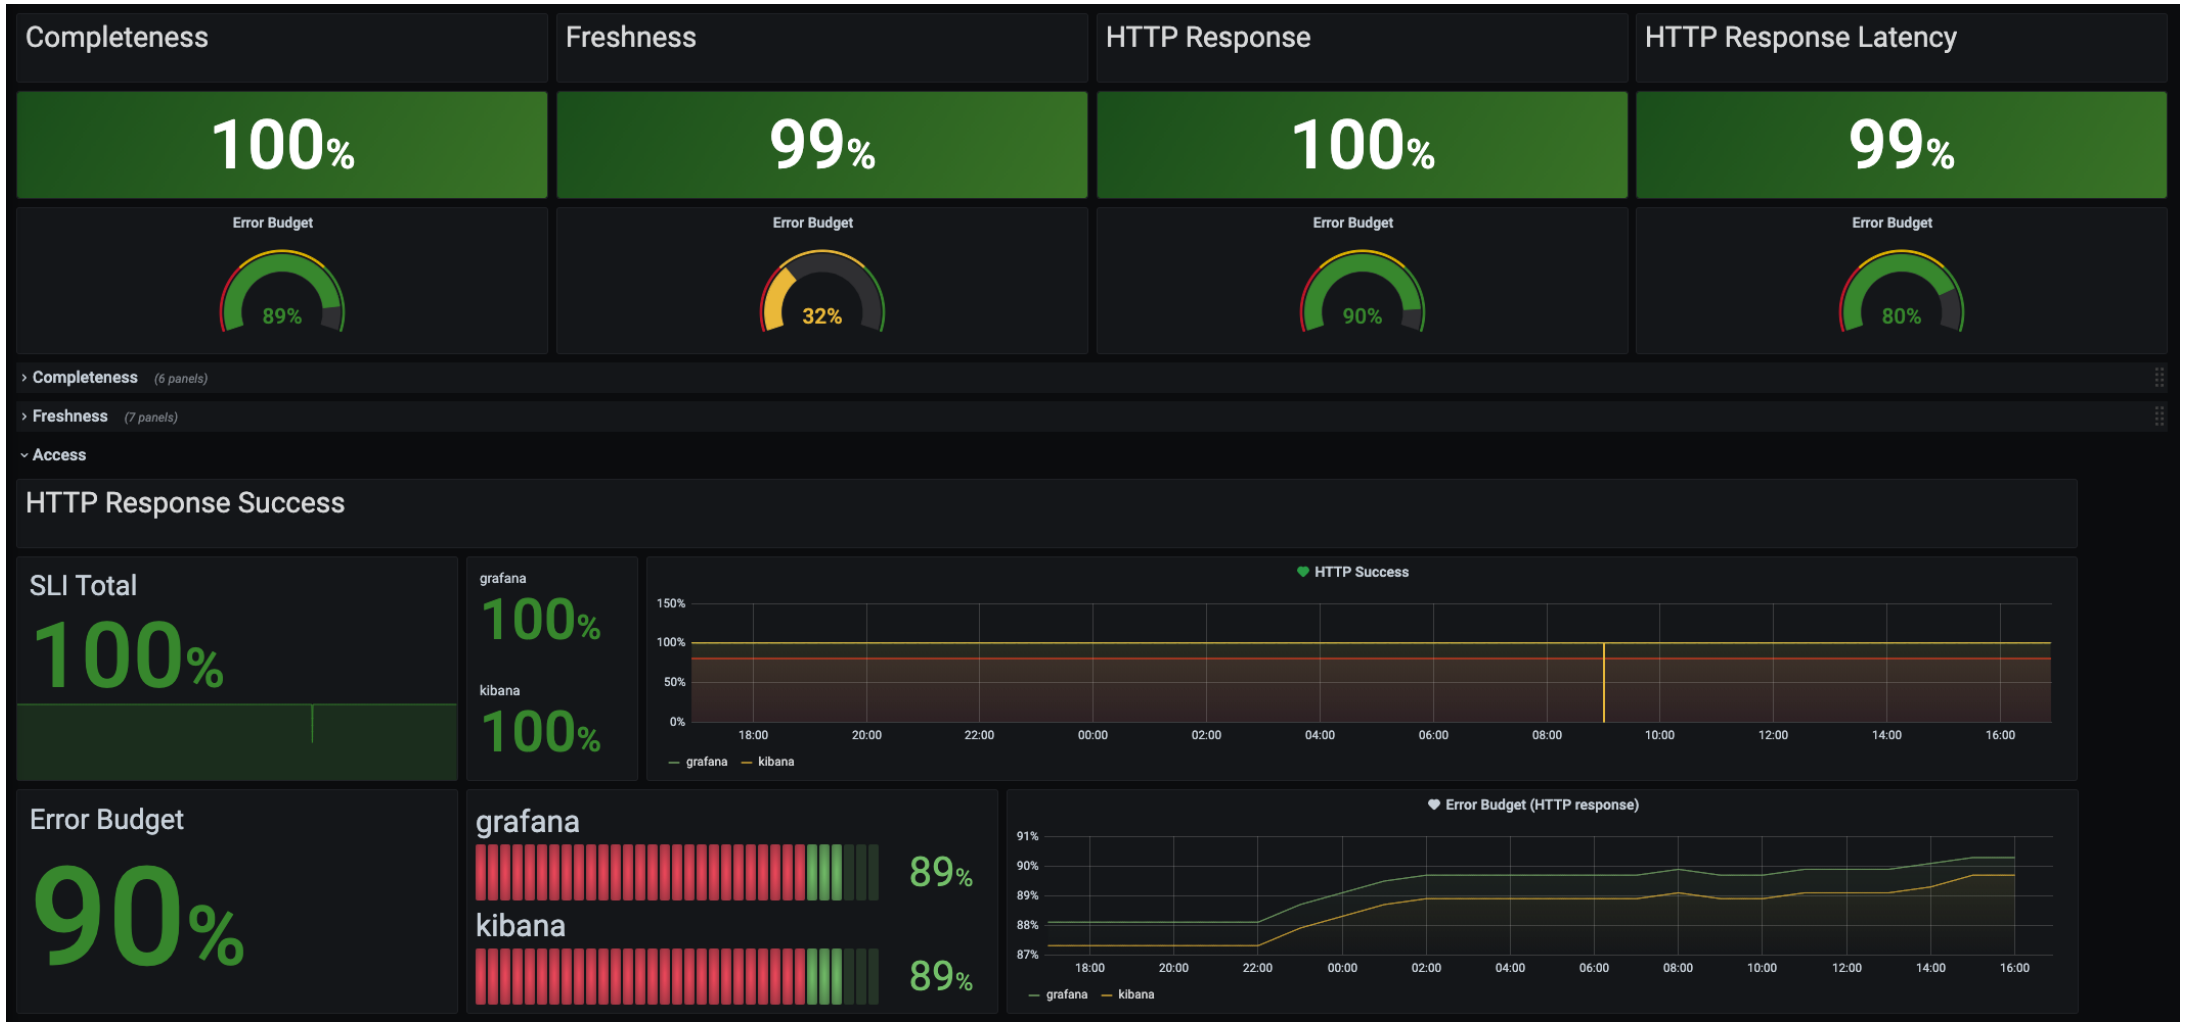

Conclusion

SLI/SLO monitoring is key for providing reliable services :

- Better user experience
- Reduce maintenance costs
- Help identifying misfunctioning

Thanks :)


felipe.bastos@cern.ch


felipe.bastos@ymail.com


linkedin.com/in/felipebastosb/


felipebastosb.github.io

References

- Beyer, Betsy, et al. *Site Reliability Workbook: Practical Ways to Implement SRE*. O'Reilly Media, Incorporated, 2018.
- Google. “The Art of SLOs - .” Google, 2020.
- Holloran, Dan. “The Comprehensive Site Reliability Engineering (SRE).” *VictorOps*, 20 Mar. 2020, victorops.com/blog/the-comprehensive-site-reliability-engineering-sre-pdf.
- Mehta, Sushant. “10 Tips To A Faster PC.” *Drivers.com*, 23 June 2015, www.drivers.com/update/pc-fix-tips/10-tips-to-a-faster-pc/.
- Rahman, Aber. “Getting Started with Site Reliability Engineering (SRE) .” Toronto Agile Conference , Oct. 2018.
- Thorne, Stephen. “Getting Started with SRE.” *It Revolution*, 29 Sept. 2018, itrevolution.com/getting-started-with-sre-stephen-thorne-google/.
- Tsvetkov, Nikolay. “Building Scalable & Reliablemonitoring System.” CERN, 2020.