


TECHNIQUES FOR SLOS AND ERROR BUDGETS AT SCALE

Fred Moyer
twitter/@phredmoyer
Conf42 Observability 2023



Have you used this in your career?

Hi, I'm Fred

SLOgician (like statistician)

Thinks about SLOs, SLIs, Error Budgets

Observability Hacker, Economist

TSDBs, Metrics/Logs/Traces, Histograms

Observability/Monitoring Engineer

20+ yrs C, Perl, Ruby, Go, Python, yadayada

Dad


Two kids, needs more sleep/coffee

Opinions in this talk are my own, not my employer's



@phredmoyer

How do you implement
SLOs for 1,000+
engineers?



Google Cloud
SRE implements DevOps


Presents:
SLIs, SLOs, SLAs, oh my!




Seth Vargo
@sethvargo



Liz Fong-Jones
@lizthegrey



Edited by Betsy Beyer,
Niall Richard Murphy, David K. Rensin,
Kent Kawahara & Stephen Thorne



Curated and edited by
David N. Blank-Edelman



Alex Hidalgo

CIRCONUS

Latency SLOs Done Right

SREcon19 Americas

@phredmoyer #SREcon



SLOconf
SERVICE LEVEL OBJECTIVE CONFERENCE

Cambrian
ExpSLOsion
2016-2022

SLI: Good vs bad requests

99th percentile home page latency over 5 minutes < 500ms

Home page request served in < 100ms

Home page request response code != 5xx

[Metric Identifier] [Operator] [Metric Value]

SLO: #good/#bad+time_range

99.95% of 99th percentile home page latency over 5 minutes < 500ms over the trailing month

99.95% of home page request response code != 5xx over last 7 days

99.95% of home page requests served in < 100ms over last 24 hours

[Success Objective] [SLI] [Period]

@phredmoyer

EB: 1-SLO, $1-0.9995 = 0.05\%$

Allow 0.05% failure of 99th percentile home page latency over 5 minutes < 500ms over trailing month

Allow 0.05% failure of home page request response code != 5xx over last 7 days

Allow 0.05% failure of home page requests served in < 100ms over last 24 hours

[Error Budget] [SLI] [Period]

@phredmoyer

Keys to SLO / Error Budget Democratization

Real world examples that are easy to reference

Formulas that can be parsed by humans and code

Be explicit; small details make big differences

Latency AND Availability

SLI

Home page request response code **!= 5xx**

or

Home page request served in **< 100ms**

99.95% of ((home page request response code **!= 5xx**) or (home page requests served in **< 100ms**)) **over last 7 days**


SLO

EB

Allow 0.05% failure of ((home page request response code **!= 5xx**) or (home page requests served in **< 100ms**)) **over last 7 days**

Measuring Availability is Easy


Measuring Latency is Not Easy



Quantifying Latency at Scale


<https://medium.com/p/9176cede3fe4>

A Common Mistake

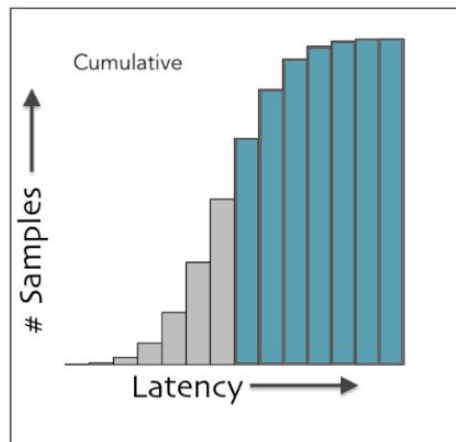


“Dr. Histogram - how I learned to stop worrying and love Latency Bands”


Latency = 125ms



Log Linear
Openhistogram.io
gt_100__lt_200



Cumulative, *Prometheus*
le_10000, ..., le_1000, le_200



Inverse Cumulative
Zendesk
gt_10, gt_50, gt_100
(but !gt_200)

SLO: #good/#bad+time_range

99.95% of home page requests served in < 100ms
over last 24 hours

Using histogram data for SLI

$(\text{sum}(\#reqs < 100ms) / \text{sum}(\#reqs)) * 100 \geq$


99.95%

[Success Objective] [SLI] [Period]


@phredmoyer

Use raw histograms, avoid sketches & approximations

Uniform Distribution / Accuracy Comparison




API Latencies / Accuracy Comparison



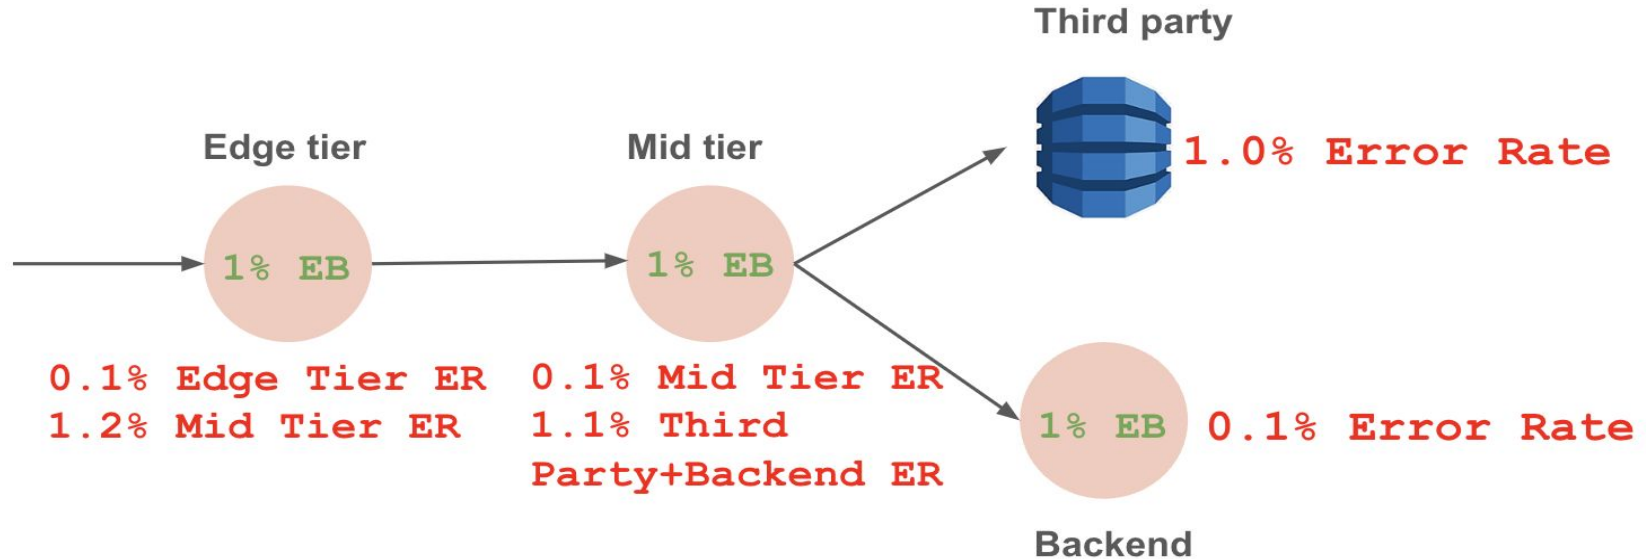
Circllhist - <https://arxiv.org/abs/2001.06561>

Decomposing Histogram Modes



“Percentiles don’t work...” - <https://adrianco.medium.com/>

Multi Service SL0s / Error Budgets



Thank you!

Questions?

[linkedin.com/redhotpenguin](https://www.linkedin.com/redhotpenguin)

[twitter/@phredmoyer](https://twitter.com/phredmoyer)