



SREcon20 Americas

Latency and Availability Error Budgets Done Right at Scale

Fred Moyer

12/8/2020

Hi, I'm Fred

SLOgician (like statistician) @Zendesk

Thinks about SLOs, SLIs, Error Budgets

Observability Hacker

TSDBs, Metrics/Logs/Traces, Histograms

Software Engineer (SRE)

15+ yrs C, Perl, Ruby, Go, Python, blabla

Dad

Two kids, needs more sleep/coffee



AGENDA



Error Budget Refresher

Usage and Prior Art

SLIs, SLOs, Error Budgets

Formulas and Implementation Details

Multi-Service Error Budgets




Scaling across Service Dependencies

Error Budgets

prioritize reliability

work vs feature work

Doing the research



Great comprehensive overview of SLOs / Error Budgets

Implementation details an exercise for the reader

Doing the research



Presents:

SLIs, SLOs, SLAs, oh my!



Seth Vargo
@sethvargo



Liz Fong-Jones
@lizthegrey

CIRCUNUS

Latency SLOs Done Right

SREcon19 Americas

@phredmoyer #SREcon

usenix

Requires some knowledge of SLIs, SLOs, EBs Solid implementation details

AGENDA

Error Budget Refresher

Usage and Prior Art



SLIs, SLOs, Error Budgets

Formulas and Implementation Details

Multi-Service Error Budgets

Scaling across Service Dependencies

SLIs

Delineates 'Good' vs 'Bad' Requests

EXAMPLE SLIS

95th percentile home page latency over 5
minutes < 500ms

Home page request response code != 5xx

Home page request served in < 100ms

EXAMPLE SLIS

Metric Identifier

95th percentile home page latency over 5
minutes < 500ms

Home page request response code != 5xx

Home page request served in < 100ms

[Metric Identifier] [Operator] [Metric Value]

EXAMPLE SLIS

Operator

95th percentile home page latency over 5 minutes < 500ms

Home page request response code != 5xx

Home page request served in < 100ms

[Metric Identifier] [Operator] [Metric Value]

EXAMPLE SLIS

Metric Value

95th percentile home page latency over 5 minutes < 500ms

Home page request response code != 5xx

Home page request served in < 100ms

[Metric Identifier] [Operator] [Metric Value]

EXAMPLE SLIS

95th percentile home page latency over 5
minutes < 500ms

Home page request response code != 5xx

Home page request served in < 100ms

[Metric Identifier] [Operator] [Metric Value]

SLOs

Binding target for SLIs

$$\text{SLO} = \frac{\text{\#goodreqs}}{\text{\#totalreqs}} + \text{Time range}$$

EXAMPLE SLOS

99% of 95th percentile home page latency
over 5 minutes < 500ms over the trailing
month

99% of home page request response code
!= 5xx over last 7 days

95% of home page requests served in <
100ms over last 24 hours

EXAMPLE SLOS

Success Objective

99% of 95th percentile home page latency
over 5 minutes < 500ms over the trailing
month

99% of home page request response code
!= 5xx over last 7 days

95% of home page requests served in <
100ms over last 24 hours

[Success Objective] [SLI] [Period]

EXAMPLE SLOS

SLI

99% of **95th percentile home page latency**
over 5 minutes < 500ms over the trailing
month

99% of **home page request response code**
!= 5xx over last 7 days

95% of **home page requests served in <**
100ms over last 24 hours

[Success Objective] [SLI] [Period]

EXAMPLE SLOS

Period

99% of 95th percentile home page latency
over 5 minutes < 500ms **over the trailing
month**

99% of home page request response code
!= 5xx **over last 7 days**

95% of home page requests served in <
100ms **over last 24 hours**

[Success Objective] [SLI] [Period]

EXAMPLE SLOS

99% of 95th percentile home page latency
over 5 minutes < 500ms over the trailing
month

99% of home page request response code
!= 5xx over last 7 days

95% of home page requests served in <
100ms over last 24 hours

[Success Objective] [SLI] [Period]

Nobody's Perfect

Error Budget = 1-SLO

Success Objective == 99%
Error Budget = $1 - 0.99$ == 1%

EXAMPLE EBS

Allow 1% failure of 95th percentile home page latency over 5 minutes < 500ms over the trailing month

Allow 1% failure of home page request response code != 5xx over last 7 days

Allow 5% failure of home page requests served in < 100ms over last 24 hours

EXAMPLE EBS

Error Budget

Allow 1% failure of 95th percentile home page latency over 5 minutes < 500ms over the trailing month

Allow 1% failure of home page request response code != 5xx over last 7 days

Allow 5% failure of home page requests served in < 100ms over last 24 hours

[Error Budget] [SLI] [Period]

EXAMPLE EBS

SLI

Allow 1% failure of **95th percentile home page latency over 5 minutes < 500ms** over the trailing month

Allow 1% failure of **home page request response code != 5xx** over last 7 days

Allow 5% failure of home **page requests served in < 100ms** over last 24 hours

[Error Budget] [SLI] [Period]

EXAMPLE EBS

Period

Allow 1% failure of 95th percentile home page latency over 5 minutes < 500ms **over the trailing month**

Allow 1% failure of home page request response code != 5xx **over last 7 days**

Allow 5% failure of home page requests served in < 100ms **over last 24 hours**

[Error Budget] [SLI] [Period]

EXAMPLE EBS

Allow 1% failure of 95th percentile home page latency over 5 minutes < 500ms over the trailing month

Allow 1% failure of home page request response code != 5xx over last 7 days

Allow 5% failure of home page requests served in < 100ms over last 24 hours

[Error Budget] [SLI] [Period]

Keys to 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

Latency AND Availability

SLI

Home page request response code \neq 5xx
or
Home page request served in $<$ 100ms

99% of ((home page request response code \neq 5xx) or (home page requests served in $<$ 100ms)) over last 7 days

SLO

Allow 1% failure of ((home page request response code \neq 5xx) or (home page requests served in $<$ 100ms)) over last 7 days

Error Budget

Latency AND Availability

1

Define the source

Monitor Based

Metric Based

Good events (numerator)

a Metric zendesk.classic.app.sli.request.A... from bin:gt_0 *

sum by (everything) as count

+

b Metric zendesk.classic.app.sli.request.A... from bin:gt_10000 *

sum by (everything) as count

+

c Metric zendesk.classic.app.sli.request.A...
from status_code_range:5xx * sum by (everything) as count

+

→ a - b - c

Add Query +

Error Budget (1%) usage for all Tick

25.502%

2

Set your targets

Target: 99.95 % Time Window: 7 Days Warning: 99.99 %

Target: 99.95 % Time Window: 30 Days Warning: 99.99 %

Target: 99.95 % Time Window: 90 Days Warning: 99.99 %



AGENDA

Error Budget Refresher

Usage and Prior Art

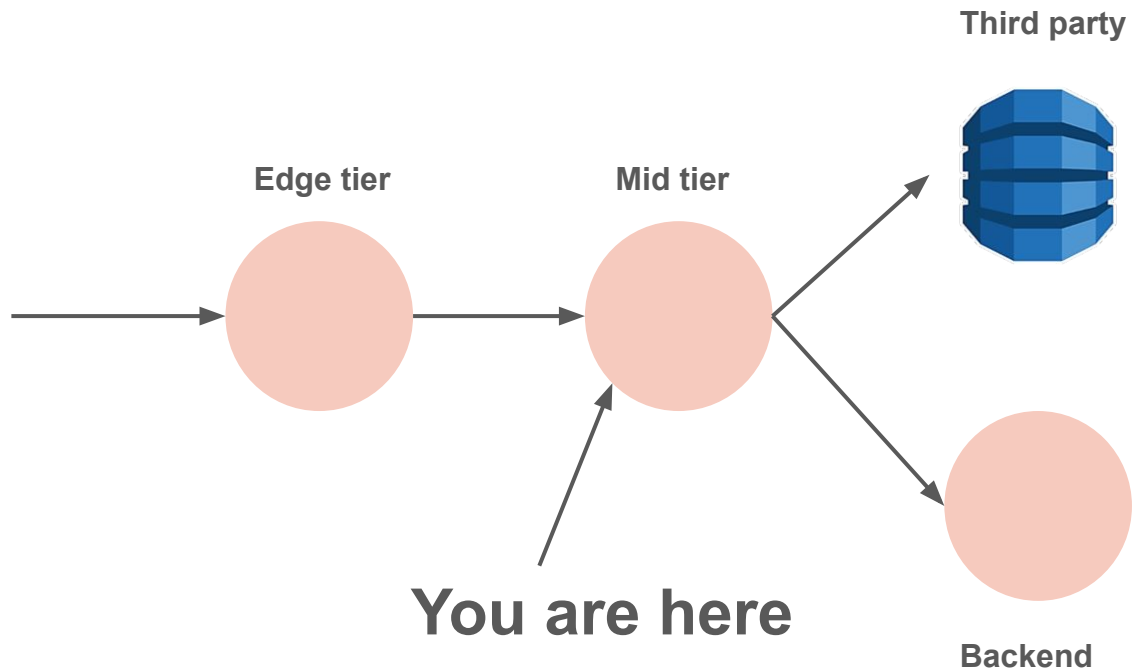
SLIs, SLOs, Error Budgets


Formulas and Implementation Details

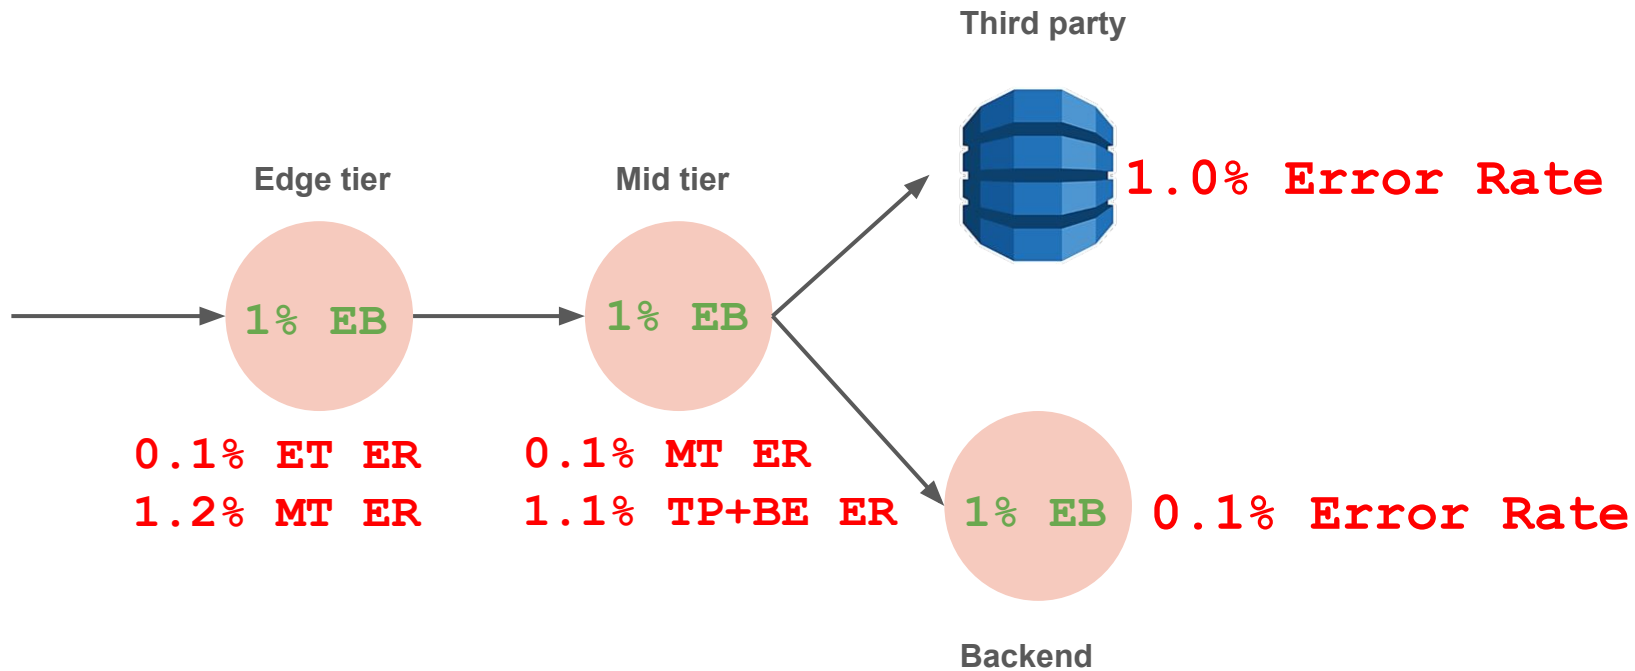


Multi-Service Error Budgets

Scaling across Service Dependencies









#srecon



@phredmoyer

Thank you

twitter/@phredmoyer
SREcon Slack - Fred Moyer

We're hiring!
zendesk.com/jobs

