nCloudsSM

How DevOps Teams Use SRE to Innovate Faster with Reliability

GOLD DATADOG PARTNER | MSP

www.nclouds.com

DATADOG

Datadog is the essential monitoring and security platform for cloud applications. We bring together end-to-end traces, metrics, and logs to make your applications, infrastructure, and third-party services entirely observable. These capabilities help businesses secure their systems, avoid downtime, and ensure customers are getting the best user experience.

nClouds℠

Deliver Innovation Faster

nClouds is an **AWS Premier Consulting Partner** and award-winning provider of AWS and DevOps consulting and implementation services. Our mission is to partner with our customers, as extensions of their teams, to build and manage modern infrastructure solutions that deliver innovation faster. We leap beyond the status quo.

THE CHANNEL CO. CRN MSP 500 2021

THE CHANNEL CO. CRN FAST GROWTH 150 2021

GOLD DATADOG PARTNER | MSP

AWS Well-Architected

aws partner network

Premier Consulting Partner

- Data & Analytics Competency
- DevOps Competency
- MSP Partner
- Migration Competency
- SaaS Competency

Trusted by **INNOVATIVE BRANDS**

nClouds℠

algo
SONY
Informatica
TUNE IN
RADIVISION POWERED BY ENTREPRENEURS
Invenger
SportsRecruits
THUUZ
AVAYA
VIOLET GREY THE INDUSTRY'S BEAUTY EDIT
SOPRIS HEALTH
SleepIQ
Algorand THE BORDERLESS ECONOMY
Welkin
GENABILITY
efables.com
ondot
metamoto
GuideSpark
FenixCommerce CONVERSION.DELIVERED.
K4 CONNECT
june
WEXUS
SPECIALIZED
elara
eZdia
COLLECTIVE LABORATORIES
Captiv8
Onriva
FREEBIRD RIDES
unitQ
flexReceipts GO BEYOND THE SALE
burningglass TECHNOLOGIES
Ople
CREDITSTACKS
AgileMD
BrainCheck
ASCO CANCERLINQ Learning Intelligence Network for Quality
Alation
HYLA mobile RECONSIDERED
Merkle Data
edcast
BKJ
walks
FOND
lumina networks
Revere
businesswire A BERKSHIRE HATHAWAY COMPANY
AdviNOW Medical AUTOMATING HEALTHCARE
REBAG
KINETIC DATA Your business. Your process.
ALLUME
CLOUDCAR
Vineti
6CONNEX
5AM
GEOTIX
nDimensional
LeaseLabs FULL STACK MARKETING
Yewno
wizehive
Civil Maps
POWER
immersv
onscale
TETRASCIENCE
prodea unleash the internet of services
GSN
komodohealth
ZMB Technologies
arlo
AutoGrid
MPS Monolithic Power Systems
DESIGN REACTOR
motion MATH
nextcapital
VELOCICAST
CyberCube
endpoint
raytrix
LendingHome
AUGMEDIX
TENMARKS an amazon company
Torchmark Corporation
magicpin
Notable Labs
OpenInvest
OPERATIONSRX Simple Solutions for a Complex World
White Ops
trooly
tapClicks
sensehealth
PADI
punchh
Recognize
REVINATE
ECO FACTOR
blueshift
SoHalo
FIDELIS
THE pill CLUB
Cirrent
earnup
Lambda
HighWire
withme
TapHeaven MOBILE UA: EASIER + FASTER + SMARTER
Superside
TEAMWORKS
Unifyed
ZABBLE Reduce Waste. Save Money. Protect Earth.

How DevOps Teams Use SRE to Innovate Faster with Reliability

PRESENTERS

Mark Solomon

VP, DevOps Practice Lead

nClouds℠

Waldo Grunenwald

Tech Evangelist

DATADOG

Gagan Goswami

DevOps Engineer

nClouds℠

How DevOps Teams Use SRE to Innovate Faster with Reliability

AGENDA

DETAILS ***(All times PT)***

- 10:00 - 10:20 am - **Setting Expectations of SRE** - Mark Solomon, nClouds
- 10:20 - 10:35 am - **Modern Observability** - Waldo Grunenwald, Datadog
- 10:35 - 10:50 am - **Putting SRE into Practice for DevOps Teams** - Gagan Goswami, nClouds
- 10:50 - 11:00 am - **Live Q&A**

Special OFFERS

nClouds

Free SRE Assessment

for all eligible attendees

Poll

Setting Expectations for SRE

Mark Solomon

VP, DevOps Practice Lead

nClouds

A Bad **Forecast**

> “Through 2023, **90%** of DevOps initiatives will fail …”

- George Spafford, Analyst & Author, 2018

A Bad **Forecast**

> “Through 2023, **90%** of DevOps initiatives will fail …
>
> … to ***fully meet expectations*** due to the limitations of leadership approaches, not technical reasons.”

- George Spafford, Analyst & Author, 2018

SRE Expectations Abound

Everything will get easier

The Team Size will Double

We'll only have a very small team

Revenue Will Increase By 20%

Costs Will Be Cut By 50%

We're Gonna Use Terraform

We're Gonna Use Error Budgets

Chaos Engineering, Here we come!

I'm gonna put SRE on my resume

Good riddance Toil!!

Everything Will Be Automated

Everyone Will Be Happy Forever

No More Outages

We're Gonna Use K8S!

SRE is Not a Template of Practices

If You're At Level 0,
Does SRE = Google?

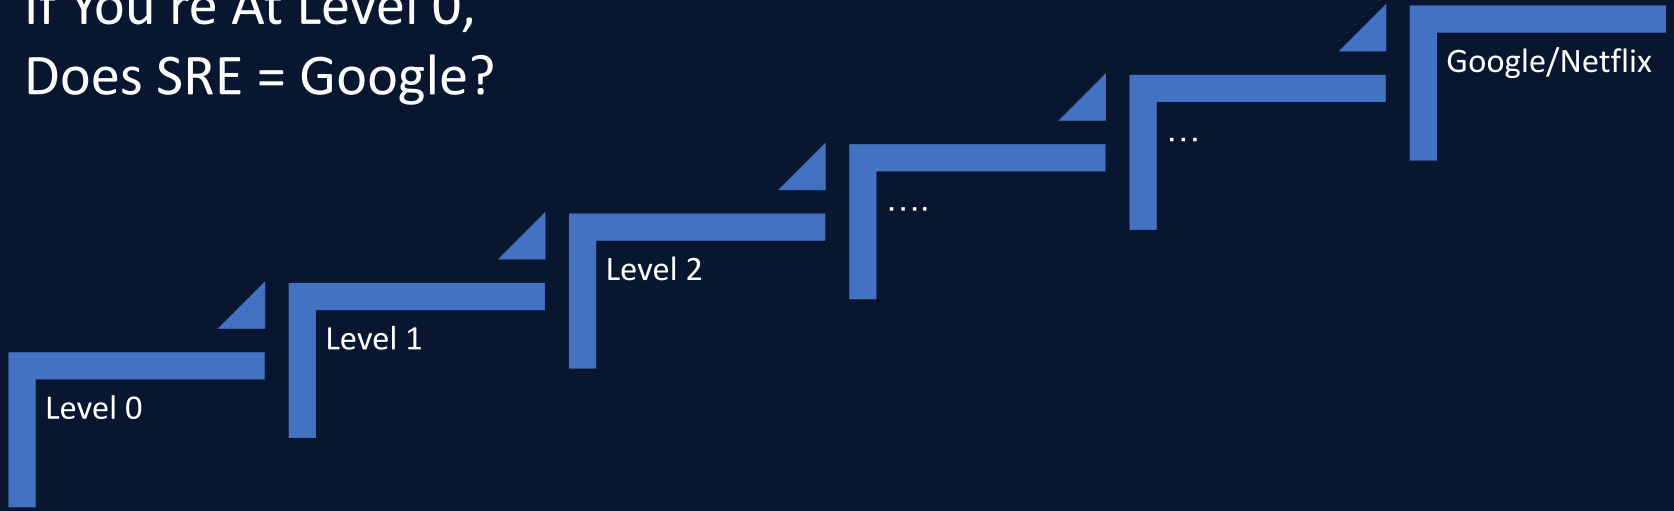

...It Is a Journey of Applying Principles

A Level 0 Story

The Good

Brilliant engineers

We knew how to fix things

The Bad

Revenue dropping, yet high operating costs

Team morale low/ attrition high

The Ugly

> 1 outage per day

> 500 update tickets in backlog

Level 1 – Getting the Right Expectations

nClouds℠

The team owned the system and its backlog.

The team must spend 50% time on innovations.

Users could recommend enhancements, and could expect a response.

The team must halt Feature changes if availability below 99.99% (new Service Level Objective).

The leadership must fund this new operating model.

5 Months Later ...

nClouds℠

- Critical automation technologies introduced.
- Critical updates complete for all subsystems.
- No outages for full month.
- Costs stable/revenue rebounding.
- 0% attrition, high morale.

SRE Principles & Expectations

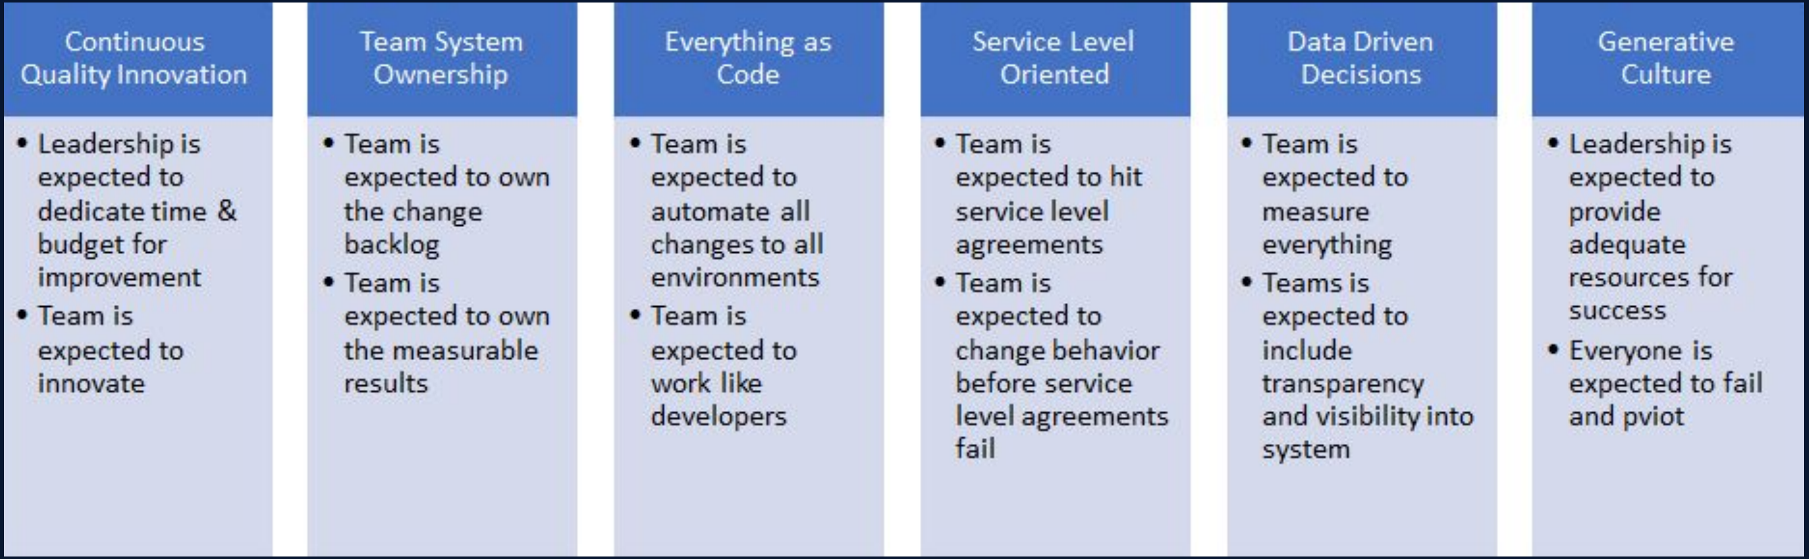

Continuous Quality Innovation	Team System Ownership	Everything as Code	Service Level Oriented	Data Driven Decisions	Generative Culture
• Leadership is expected to dedicate time & budget for improvement • Team is expected to innovate	• Team is expected to own the change backlog • Team is expected to own the measurable results	• Team is expected to automate all changes to all environments • Team is expected to work like developers	• Team is expected to hit service level agreements • Team is expected to change behavior before service level agreements fail	• Team is expected to measure everything • Teams is expected to include transparency and visibility into system	• Leadership is expected to provide adequate resources for success • Everyone is expected to fail and pviot

Poll

Modern Observability

Waldo Grunewald

Tech Evangelist

DATADOG

DATADOG

Building a SLO Culture

“S. L. [A, E, I. O, U, and definitely Y].”

DATADOG

What we’ll cover today

- Why
- What you need
- The vocabulary
- Problems you might encounter

Building a SLO Culture

Why Bother?

DATADOG

“How to balance new feature development with availability?”

What do we need?

What do we need?

Visibility

What do we need?

Transparency

What about all of those letters?

DATADOG

Service Level [A, E, I, O, U, and definitely Y]

Service Level [A, E, I, O, U, and definitely Y]

A **S**ervice **L**evel **A**greement is a contract that defines the results (and consequences) of meeting (or missing) one or more promises.

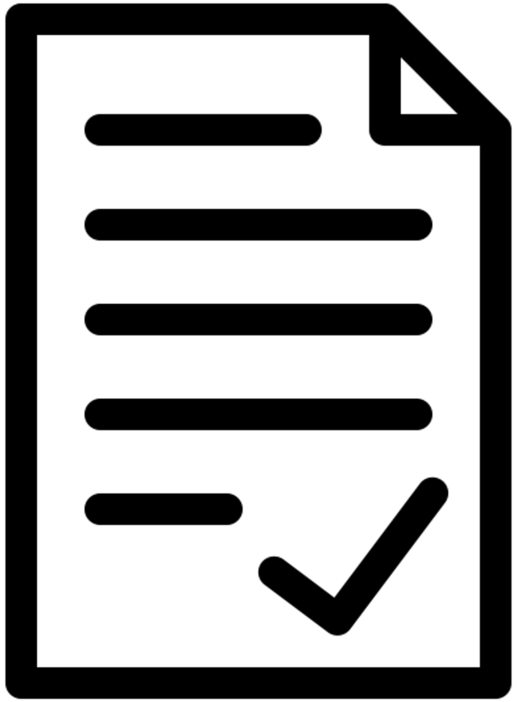

SLA

Image credit: Checked document by Rockicon from the Noun Project

Service Level [A, E, **I**, O, U, and definitely Y**]**

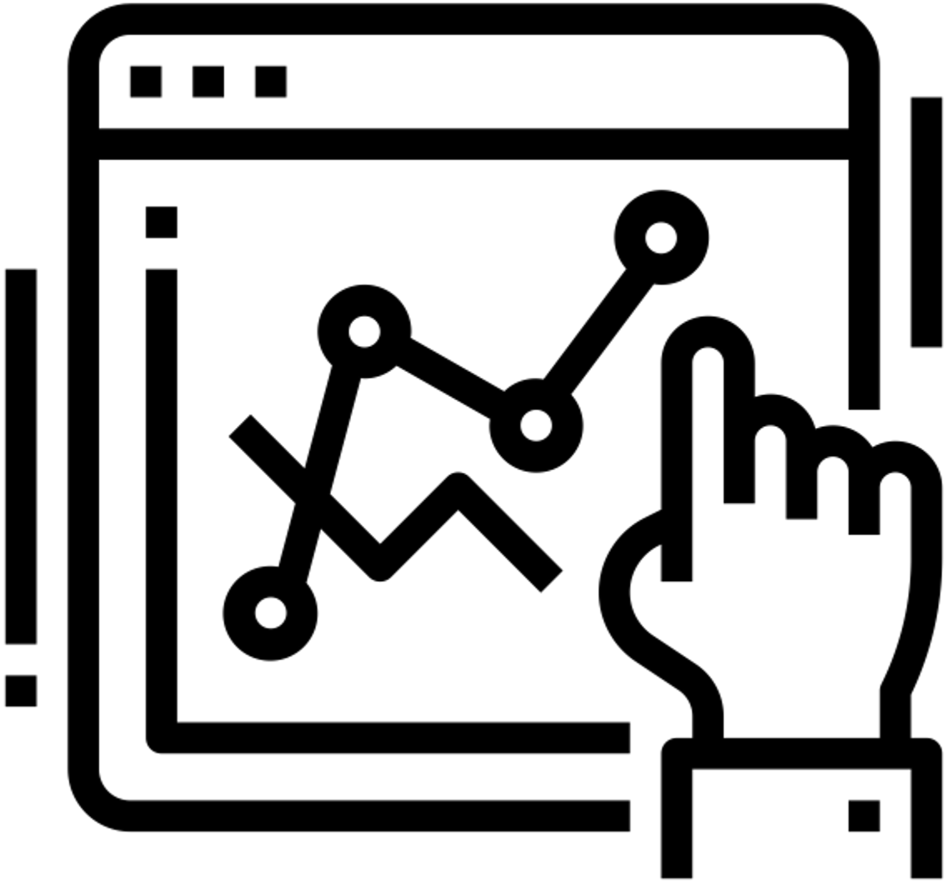

SLI

A **S**ervice **L**evel **I**ndicator is a quantitative measurement that expresses an aspect of the service.

Image credit: Indicator by Nhor from the Noun Project

DATADOG

Service Level [A, E, I, **O**, U, and definitely Y]

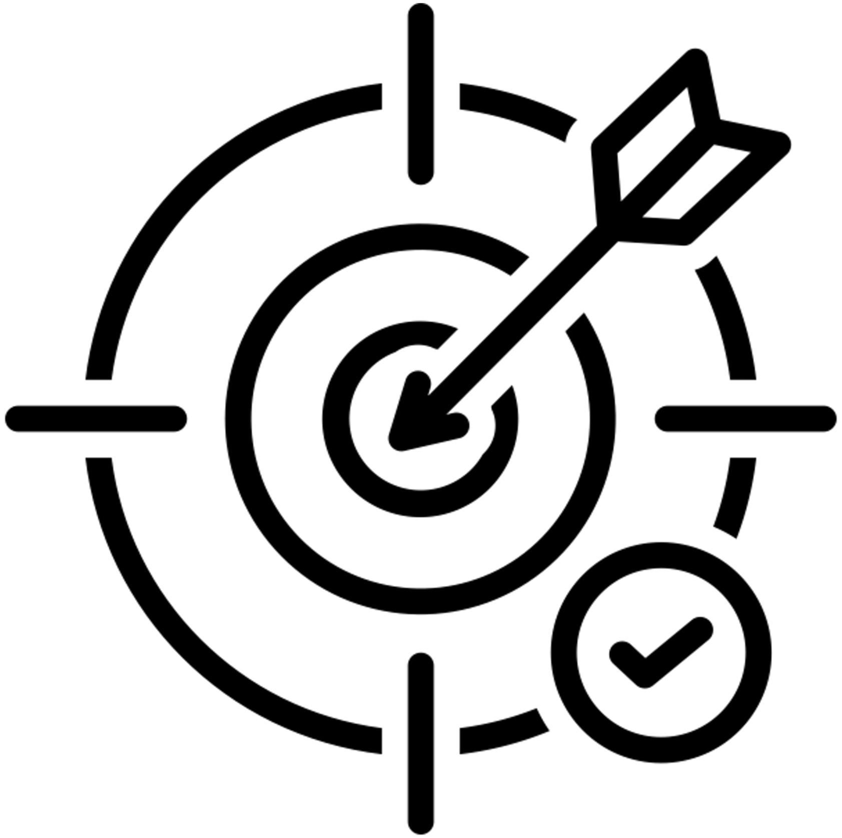

SLO

A **S**ervice **L**evel **O**bjective is a target value for a service, as measured via an SLI, over a specified time window.

Image credit: Objective by priyanka from the Noun Project

DATADOG

SLOs are applied SLIs

Requests are 99.95% successful in the last 24 hours.

90% of requests complete under 500ms over the past 30 days.

Error Budgets

“Move fast and *fix* things!”

- Failure is unavoidable.
- Balance innovation with reliability and stability.

Building an Error Budget

An SLO is identified by the product owner.

The actual objective is measured by a neutral party (hint: a monitoring system).

The difference between the actual measurement and the objective is the error budget.

Created by Teewara soontorn from Noun Project

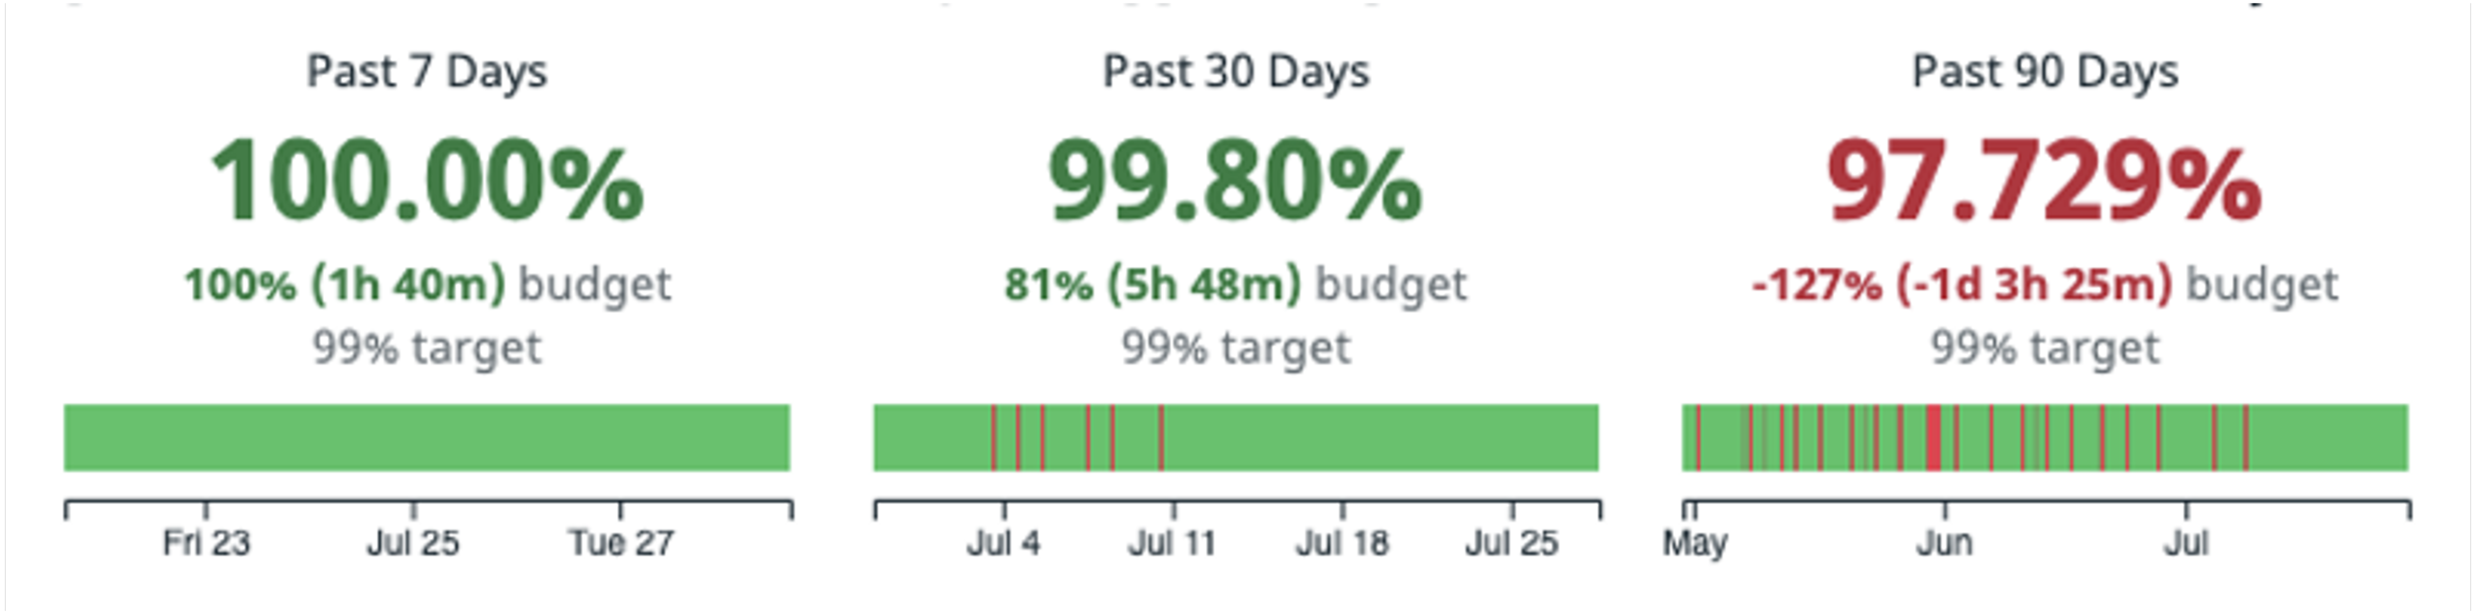

Past 7 Days

100.00%

100% (1h 40m) budget

99% target

Past 30 Days

99.80%

81% (5h 48m) budget

99% target

Past 90 Days

97.729%

-127% (-1d 3h 25m) budget

99% target

Build Budget

If the budget is zero or less, you should focus on improving that.

- Freeze new features.
- Improve o11y state.
- Prioritize tech debt.

Spend Budget

If SLO is being met, you have room to innovate.

- Add new features!
- Experiment!
- Resiliency Tests!

Service Level [A, **E**, I, O, U, and definitely Y**]**

S.L. Environments

Service Level [A, E, I, O, **U**, and definitely Y]

S.L. Updates

Service Level [A, E, I, O, **U**, and definitely Y]

S.L. Updates

Service Level [A, E, I, O, U, and **definitely Y**]

S.L. “Y”

Service Level [A, E, I, O, U, and **definitely Y**]

S.L. “Why”

S.L. “Why”

Service Level [A, E, I, O, U, and definitely Y]

S.L. “Why”

Building a SLO Culture

When should we look at our SLO status?

DATADOG

SLOs help prioritize feature dev or reliability dev

- Sprint Planning
- Starting a new task
- Alerted to a fault

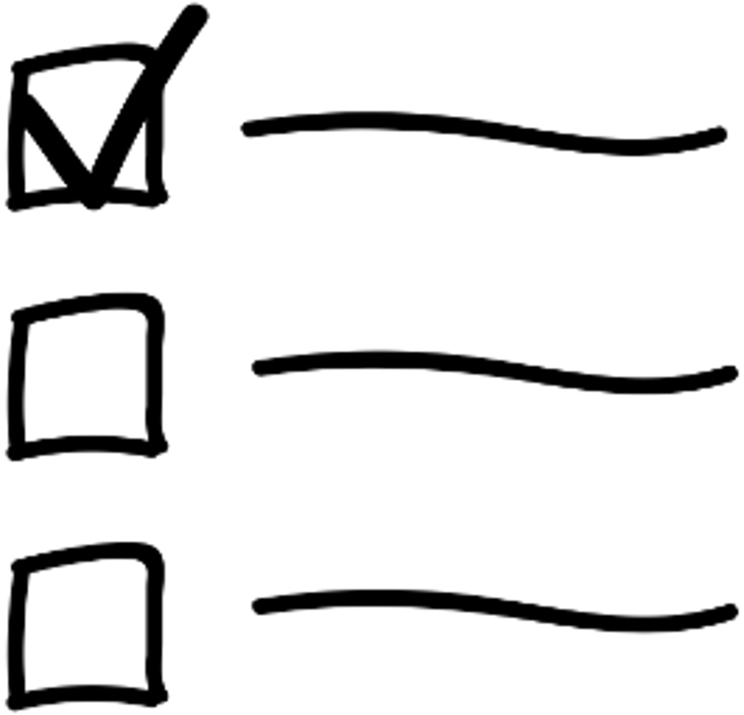

Created by W. X. Chee
from Noun Project

DATADOG

What problems can we expect?

DATADOG

- Choosing the wrong Metric
- Choosing the wrong availability
- Not understanding what the client wants
- Not using Error Budget
- Panic when a team is using it’s Error Budget

In summary

DATADOG

SLOs are Simple

But I won’t promise that they’re Easy.

“People are a problem.”
-- Douglas Adams

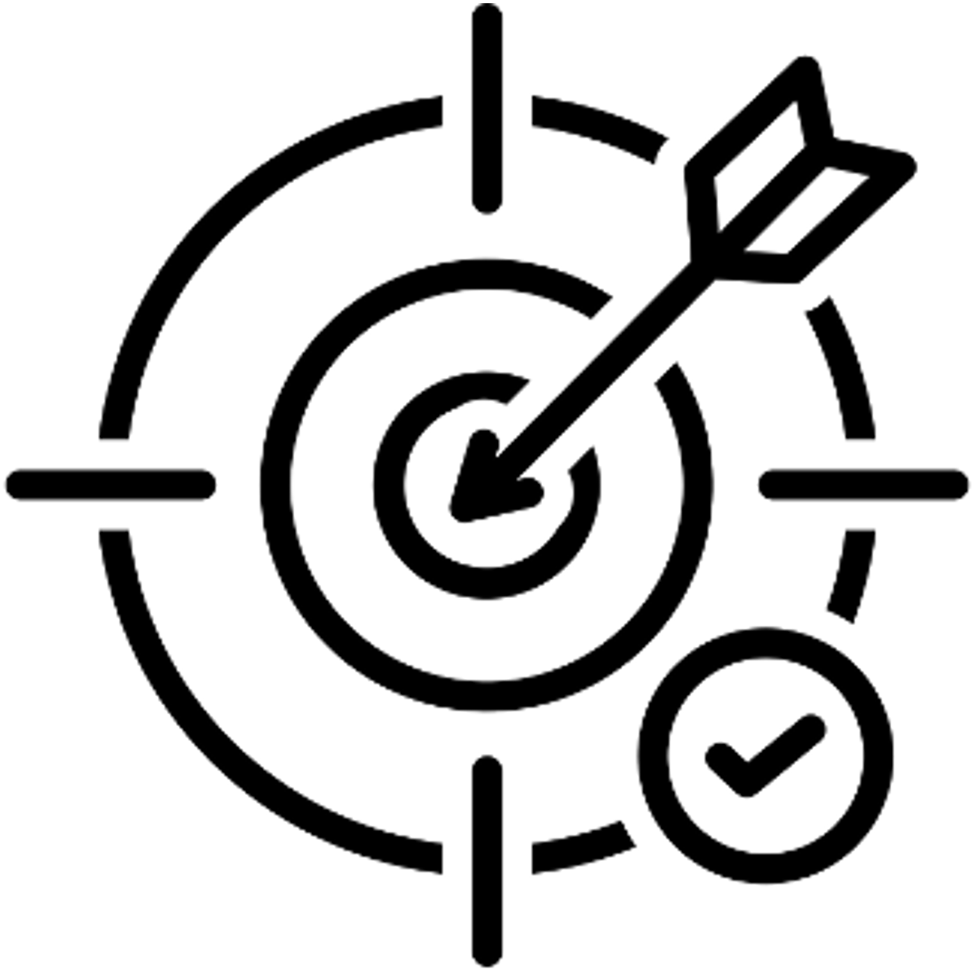

DATADOG

nCloudsSM

Poll

www.nclouds.com

aws

Putting SRE into Practice for DevOps Teams

Gagan Goswami

DevOps Engineer

www.nclouds.com

Key Metrics

nClouds

SLA

Promise__

- Promises by service provider for availability and performance of a service.

SLO

Goal__

- A target value or a range of values for a service measured using SLI.

SLI

How to measure?

- A quantitative measure of some aspect of the level of service.

Error Budget

Downtime tolerance.

- How much time you are willing to allow your systems to be down.

SLO Best Practices & Example

- Set **realistic** targets.
- Don’t have too many SLIs and SLOs.
- Set **SLI & SLO** for customer exposing components.
- Documentation & stakeholders’ agreements.
- Continuous feedback and continuous **improvements**.

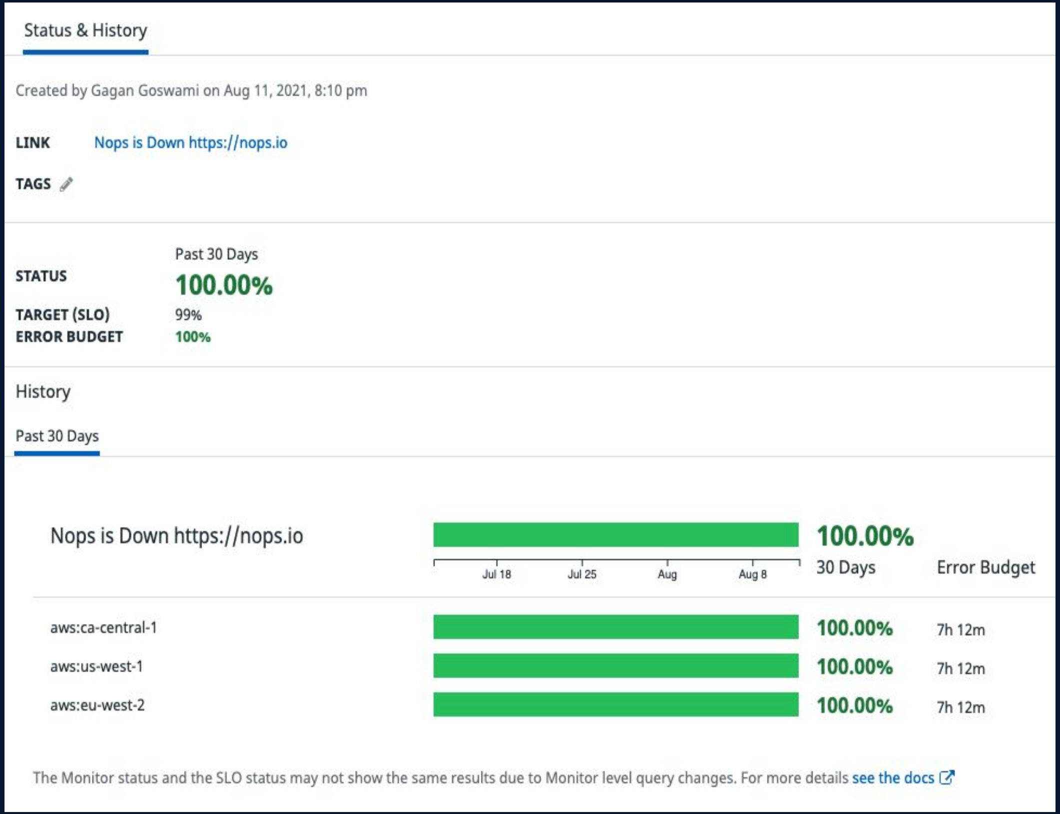

Status & History

Created by Gagan Goswami on Aug 11, 2021, 8:10 pm

LINK Nops is Down https://nops.io

TAGS

	Past 30 Days
STATUS	**100.00%**
TARGET (SLO)	99%
ERROR BUDGET	100%

History

Past 30 Days

		30 Days	Error Budget
Nops is Down https://nops.io	🟩	100.00%	
aws:ca-central-1	🟩	100.00%	7h 12m
aws:us-west-1	🟩	100.00%	7h 12m
aws:eu-west-2	🟩	100.00%	7h 12m

The Monitor status and the SLO status may not show the same results due to Monitor level query changes. For more details see the docs

https://app.datadoghq.com/slo

Demo

Blameless Postmortem

nClouds℠

The root cause of the incident.

Actions taken to mitigate/ resolve issue.

The business impact.

Actions taken to prevent recurrence.

Eliminating Toil

Ask yourself:

- Doing repetitive tasks?
- Using reactive approach?
- Not able to focus on innovation?
- Can be automated?

Chaos Engineering

Injecting harm like a vaccine, to build immunity.

- Injecting chaos.
- Stressing out systems.
- Recreate past outages.
- Implement fixes.

Run**book**

A **document** that contains all the information regarding day-to-day tasks and **remediation** steps for issues and tickets.

Plays an important role in reducing **MTTR** and resolving issues to mitigate failures and **downtime**.

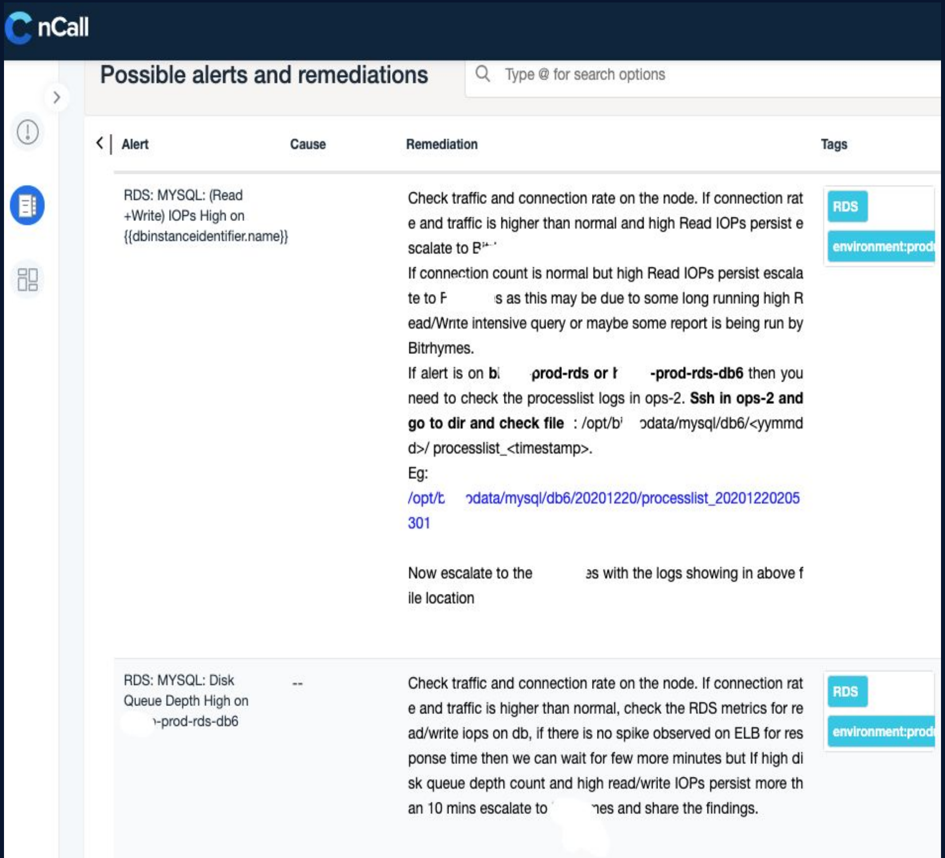

Possible alerts and remediations

Alert	Cause	Remediation	Tags
RDS: MYSQL: (Read +Write) IOPs High on {{dbinstanceidentifier.name}}		Check traffic and connection rate on the node. If connection rate and traffic is higher than normal and high Read IOPs persist escalate to B[redacted] If connection count is normal but high Read IOPs persist escalate to F[redacted]s as this may be due to some long running high Read/Write intensive query or maybe some report is being run by Bitrhymes. If alert is on **b[redacted]prod-rds or [redacted]-prod-rds-db6** then you need to check the processlist logs in ops-2. **Ssh in ops-2 and go to dir and check file** : /opt/b[redacted]data/mysql/db6/<yymmdd>/ processlist_<timestamp>. Eg: /opt/b[redacted]data/mysql/db6/20201220/processlist_20201220205301 Now escalate to the [redacted]es with the logs showing in above file location	RDS environment:prod
RDS: MYSQL: Disk Queue Depth High on [redacted]-prod-rds-db6	--	Check traffic and connection rate on the node. If connection rate and traffic is higher than normal, check the RDS metrics for read/write iops on db, if there is no spike observed on ELB for response time then we can wait for few more minutes but If high disk queue depth count and high read/write IOPs persist more than 10 mins escalate to [redacted]nes and share the findings.	RDS environment:prod

Demo

Q&A

How DevOps Teams Use SRE to Innovate Faster with Reliability

PRESENTERS

Mark Solomon
VP, DevOps Practice Lead

nClouds℠

Waldo Grunenwald
Tech Evangelist

DATADOG

Gagan Goswami
DevOps Engineer

nClouds℠

Special OFFERS

nClouds℠

Free SRE Assessment

for all eligible attendees

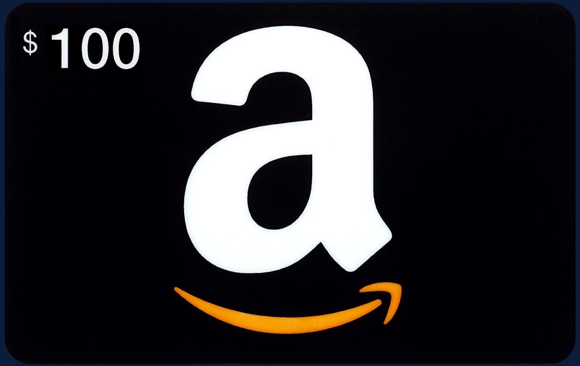

nClouds℠

Thank You!