




Have You Tried Turning It Off and On Again?

David N. Blank-Edelman
Senior Cloud Ops Advocate







@otterbook

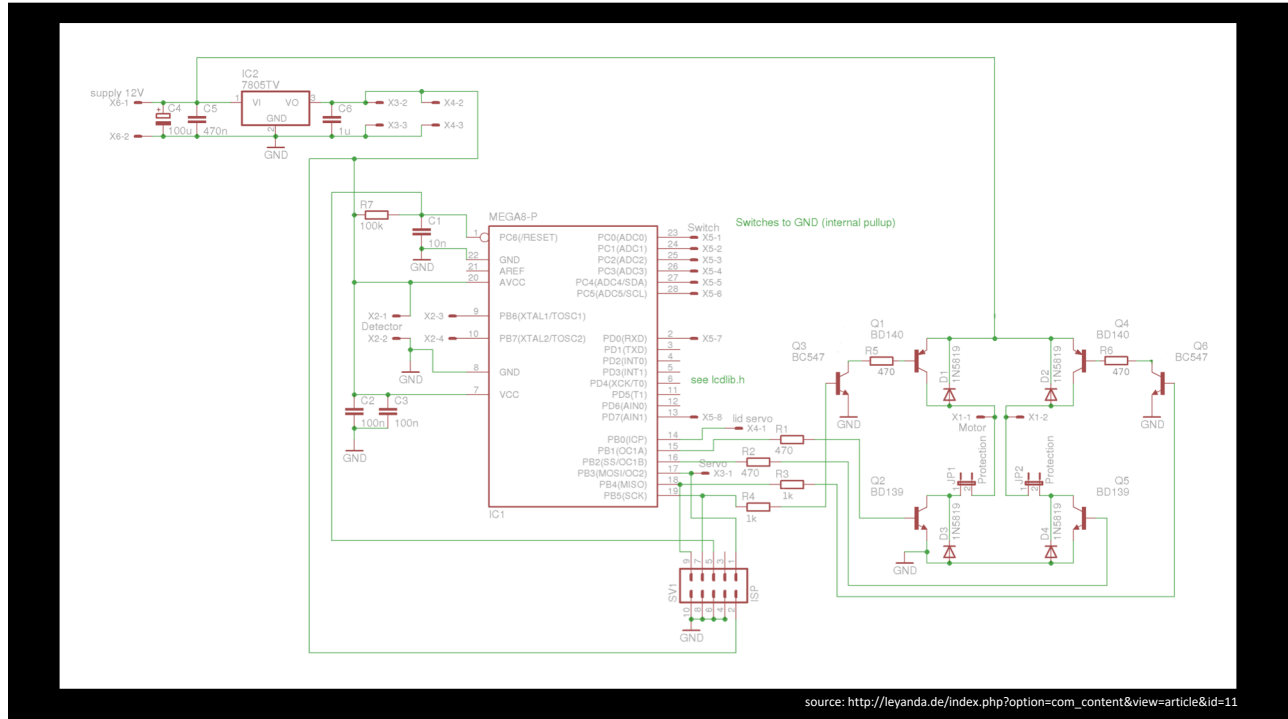


This is Production.



This is Production.

@otterbook





Q&A

@otterbook

Volunteers?


@otterbook

Rules

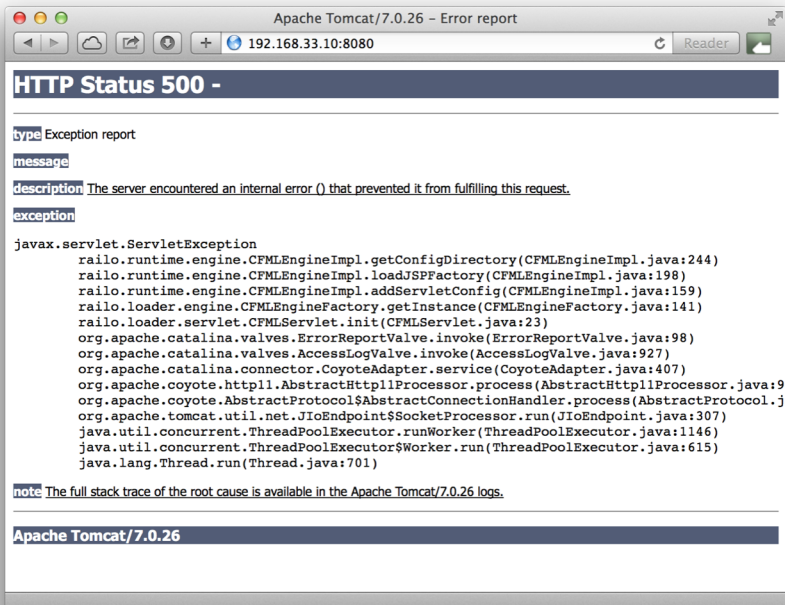

@otterbook

Level Set: SRE

@otterbook



@otterbook



Apache Tomcat/7.0.26 - Error report

192.168.33.10:8080

HTTP Status 500 -

type Exception report

message


description The server encountered an internal error () that prevented it from fulfilling this request.

exception


```
javax.servlet.ServletException
    railo.runtime.engine.CFMLServletImpl.getConfigDirectory(CFMLServletImpl.java:244)
    railo.runtime.engine.CFMLServletImpl.loadJSPFactory(CFMLServletImpl.java:198)
    railo.runtime.engine.CFMLServletImpl.addServletConfig(CFMLServletImpl.java:159)
    railo.loader.engine.CFMLServletFactory.getInstance(CFMLServletFactory.java:141)
    railo.loader.servlet.CFMLServlet.init(CFMLServlet.java:23)
    org.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:98)
    org.apache.catalina.valves.AccessLogValve.invoke(AccessLogValve.java:927)
    org.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:407)
    org.apache.coyote.http11.AbstractHttp11Processor.process(AbstractHttp11Processor.java:9)
    org.apache.coyote.AbstractProtocol$AbstractConnectionHandler.process(AbstractProtocol.j
    org.apache.tomcat.util.net.JIoEndpoint$SocketProcessor.run(JIoEndpoint.java:307)
    java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1146)
    java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)
    java.lang.Thread.run(Thread.java:701)
```

note The full stack trace of the root cause is available in the Apache Tomcat/7.0.26 logs.


Apache Tomcat/7.0.26



Edited by Betsy Beyer, Chris Jones, Jennifer Petoff & Niall Richard Murphy




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




Edited by David N. Blank-Edelman

@otterbook



- Airbnb
- Amazon
- Apple
- Baidu
- Dropbox
- Etsy
- Facebook
- GitHub
- LinkedIn
- Microsoft
- Netflix
- Pinterest
- Spotify
- Stack Exchange
- Twitter
- Uber
- Yahoo!
- Yelp



SREcon14

05.30.14 | SANTA CLARA, CA

THE ADVANCED
COMPUTING SYSTEMS
ASSOCIATION

SOLD
OUT!

CONNECT WITH US

- Twitter
- Facebook
- LinkedIn
- Google+
- YouTube

SPONSORS

Silver Sponsor

Keys to SRE

Ben Treynor

Open Access Content

USENIX is committed to Open Access to the research presented at our events. Papers and proceedings are freely available to everyone once the event begins. Any video, audio, and/or slides that are posted after the event are also free and open to everyone. [Support USENIX](#) and our commitment to Open Access.

Abstract:

Ben Treynor joined Google as Site Reliability Tsar in 2003. He is the founder of Google's Site Reliability team and grew it organically from an original core of 7 "production" engineers to its current ranks of >1200 software engineers. SRE is responsible for everything from Google's internal software infrastructure, to user services like Search, Gmail, Ads, and to the burgeoning Cloud Platform. Additionally, Ben is responsible for Google's worldwide internal and external network (since 2004), its data centers and hardware operations (since 2009), and is part of the Google Cloud Platform management team (since late 2013).

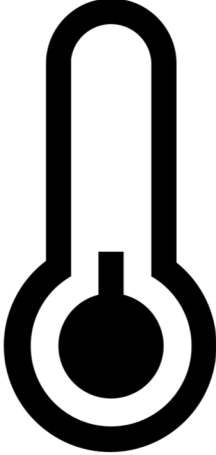
CONNECT WITH US

What Makes SRE, SRE (dramatic recreation)

- hire only coders
- have an SLA for your service
- measure and report performance against SLA
- Use Error Budgets and gate launches on them
- Common staffing pool for SRE and DEV
- Excess Ops work overflows to DEV team
- Cap SRE operational load at 50%
- Share 5% of ops work with DEV team
- Oncall teams at least 8 people, or 6x2
- Maximum of 2 events per oncall shift
- Post mortem for every event
- Post mortems are blameless and focus on process and technology, not people

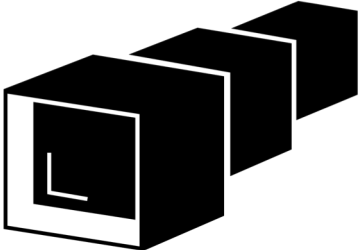

What Makes SRE, SRE (dramatic recreation)

- hire only coders
- **have an SLA for your service**
- **measure and report performance against SLA**
- **Use Error Budgets and gate launches on them**
- Common staffing pool for SRE and DEV
- Excess Ops work overflows to DEV team
- Cap SRE operational load at 50%
- Share 5% of ops work with DEV team
- Oncall teams at least 8 people, or 6x2
- Maximum of 2 events per oncall shift
- Post mortem for every event
- Post mortems are blameless and focus on process and technology, not people



SLO

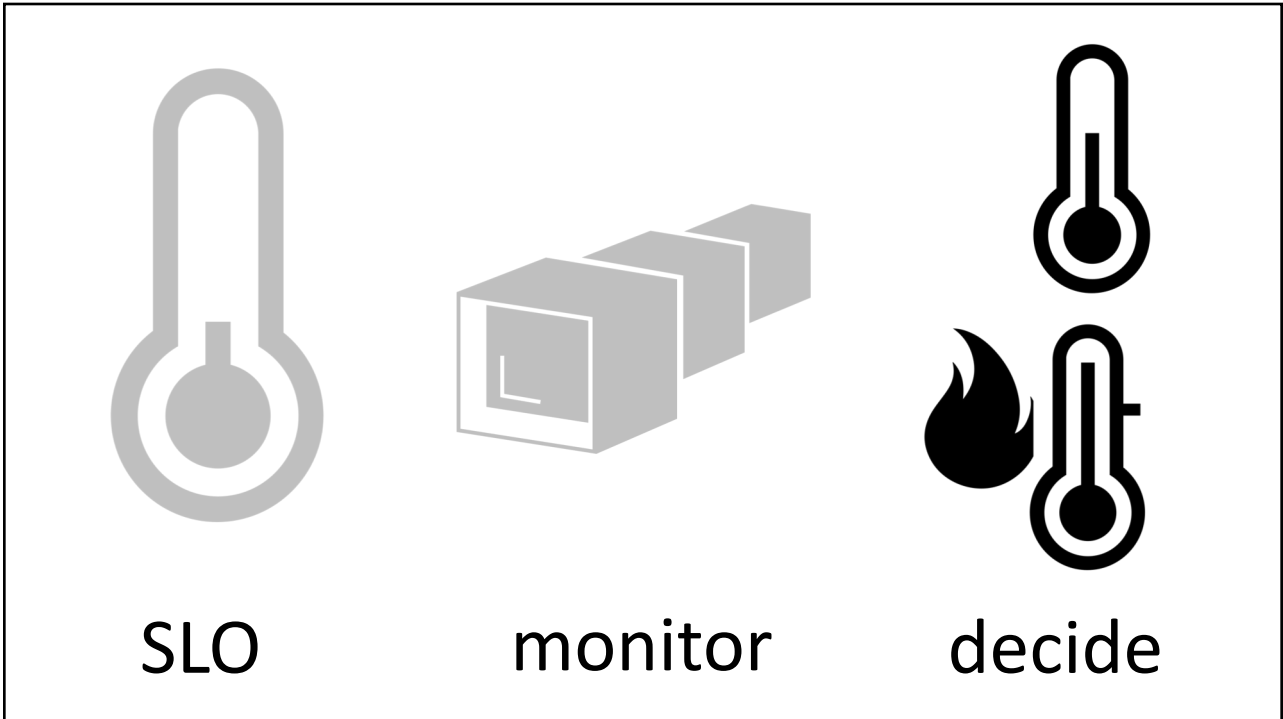
@otterbook



SLO

monitor

@otterbook



Observation #1:

Create virtuous and reinforcing feedback loops

The illustration shows a man with short dark hair and a goatee, wearing a blue suit jacket, a white shirt, and a red tie. He is holding a brown clipboard in his left hand and pointing with his right index finger towards the text on the left. The background behind him is a stylized green and blue pattern with white dots.


What Makes SRE, SRE (dramatic recreation)

- hire only coders
- have an SLA for your service
- measure and report performance against SLA
- Use Error Budgets and gate launches on them
- Common staffing pool for SRE and DEV
- Excess Ops work overflows to DEV team
- Cap SRE operational load at 50%
- Share 5% of ops work with DEV team
- Oncall teams at least 8 people, or 6x2
- Maximum of 2 events per oncall shift
- **Post mortem for every event**
- **Post mortems blameless and focus on process and technology, not people**

Observation #2:


You can't fire your way to reliable.






Q: What are the characteristics of an operations practice that actively influence a system towards greater resiliency?

Q: What are ***some of*** the characteristics of an operations practice that actively influence a system towards greater resiliency?




The Nature of the Work

@otterbook




Interfaces

@otterbook




Data

@otterbook



Errors

@otterbook




Ambiguity

@otterbook

“...I would like to beg you, dear Sir, as well as I can, to have patience with everything unresolved in your heart and to try to love the questions themselves as if they were locked rooms or books written in a very foreign language. Don't search for the answers, which could not be given to you now, because you would not be able to live them. And the point is, to live everything. Live the questions now. Perhaps then, someday far in the future, you will gradually, without even noticing it, live your way into the answer.”

—Rainer Maria Rilke, *Letters to a Young Poet* (#4)



Q: What are ***some more of*** the characteristics of an operations practice that actively influence a system towards greater resiliency?

@otterbook

(More) Characteristics of an Operations Practice

1

2 3 4 5

@otterbook

Check In

@otterbook



**David N.
Blank-Edelman**

Senior Cloud Ops Advocate

@otterbook

dnb@microsoft.com

/in/dnblankedelman

