

# ENGINEERED CHAOS

BREAKING PROD AND GETTING AWAY WITH IT

# Oliver Leaver-Smith

- Devops/Platform/Site Reliability Engineer/Sysadmin
- Core Tribe, Sky Betting and Gaming
- Login, SSO, registration, verification, payments, account services, responsible gambling tools
- When we're down, everyone's down









Chaos Engineering is the discipline of experimenting on a distributed system in order to build confidence in the system's capability to withstand turbulent conditions in production

<http://principlesofchaos.org>

&lt;&gt; Code

Issues 0

Pull requests 0

Projects 0

Wiki

Insights

Chaos Monkey is a resiliency tool that helps applications tolerate random instance failures.

112 commits

10 branches

3 releases

14 contributors

Apache-2.0

Branch: master

New pull request

Create new file

Upload files

Find file

Clone or download



lorin Merge pull request #40 from Branislav1993/patch-1

Latest commit f5e27ed 11 days ago

cal	fix typo	a year ago
clock	initial commit	a year ago
cmd/chaosmonkey	Add support for pluggable constrainer	a month ago
command	Add support for pluggable constrainer	a month ago



Failures happen, and resilience drills help organizations prepare for them.

<http://queue.acm.org/detail.cfm?id=2371516>



# Global computer outage grounds British Airways flights


Travellers report being trapped in Heathrow airport after major IT system failure




BA-owner IAG has announced plans for a new airline offering cheap transatlantic services

**Naomi Rovnick** and **Tanya Powley**

MAY 27, 2017

 223


A man with glasses and a beard, wearing a dark blue sweater and jeans, is sitting in a red office chair in a server room. He is using a laptop and has his feet propped up on a server rack. The room is filled with rows of server racks on both sides, with various cables and equipment visible. The floor has metal grates for ventilation.

**BA IT systems failure: Uninterruptible  
Power Supply was interrupted**  
Potentially by a panicking contractor, if reports are  
to be believed









# The Preamble

- Day of the week
- Time of day
- Events on now or starting soon
- The environment in which to work
- Slack channel



slack




# The People

- **Orchestrator**
  - Breaks the thing and plays the role of Service Desk
- **Platform Engineer and 1 or 2 Developers**
  - One of these will be “primary” as is the case for real on-call scenarios
- **Incident Commander (IC)**
  - Organises the Engineers working on the incident, allowing them to focus on the fix. Translates tech to “normal” language
- **Service Level Analyst/Manager (SLA/SLM)**
  - Knows the priorities and business impact of services being unavailable. Manages the relationship between the response team and business stakeholders. Translates updates into business speak
- **Third parties**
  - Includes other tribes, service desk, directors, service providers, etc. Often played by the orchestrator



















## Commit Security Check



Warning! This code fails 23 tests.

☒ Roll out as high priority to All circles ▾

You have administrative rights. Bypass test flags?

- ☐ via trusted users
- ☒ immediately

Rolling out this code could damage Google's infrastructure

☒ I understand

Cancel

Save











# Retro

- Post-its and pens
  - What went well
  - What didn't go so well
  - Actions
- Stick them on the wall
- Discuss



It's not easy when you're regulated

It's not always about fixing the incident

It's not always planned

# The Future

- Automate more
- More diverse scenarios
- Really break LIVE

# @muggahtee

- Twitter
- GitHub
- ... most things

Any questions?