

# Expect the Unexpected

Preparing SRE Teams for Responding to Novel Failures

John Arthorne  
@jarthorne

Shopify Production Engineering



# Help > About

John Arthorne

- Developer/Manager/SRE at Shopify

Shopify

- Software for Commerce
- 30 → 120 SREs
- Average \$1700 GMV / second

# What We'll Cover

- Why should I care about incident prep?
- Survey of incident preparation practices
- Preparing for truly unexpected failures

SRE practices lead us to a  
**high failure novelty rate**




- Engineers love orderly and predictable
- Deliberate practice makes incidents more comfortable
- How do we practice the unpredictable??



# Transparent Response

- Make realtime shadowing easy
- All decision making and data behind it shared widely
- Senior staff leading by example. Not knowing or guessing wrong is ok.





## Incident Simulation


- aka “Wheel of Misfortune”
- Enables taking on incident response role in a safe practice environment
- Caveat: only as good as existing human understanding of the system

# Game Days

- Start from hypothesis of system behavior in failure state
- Induce real production failure
- Observe, recover, adapt








## Turn Rusty Knobs

- Exercise failure recovery practices and software in production
- Builds confidence that it will work when needed
- Weather events are a good forcing function


# Automated Failure Tests

- Focus on most routine failures: timeouts, connection failures
- Many open source tools available to help
- Caveat: Can't validate full system behavior



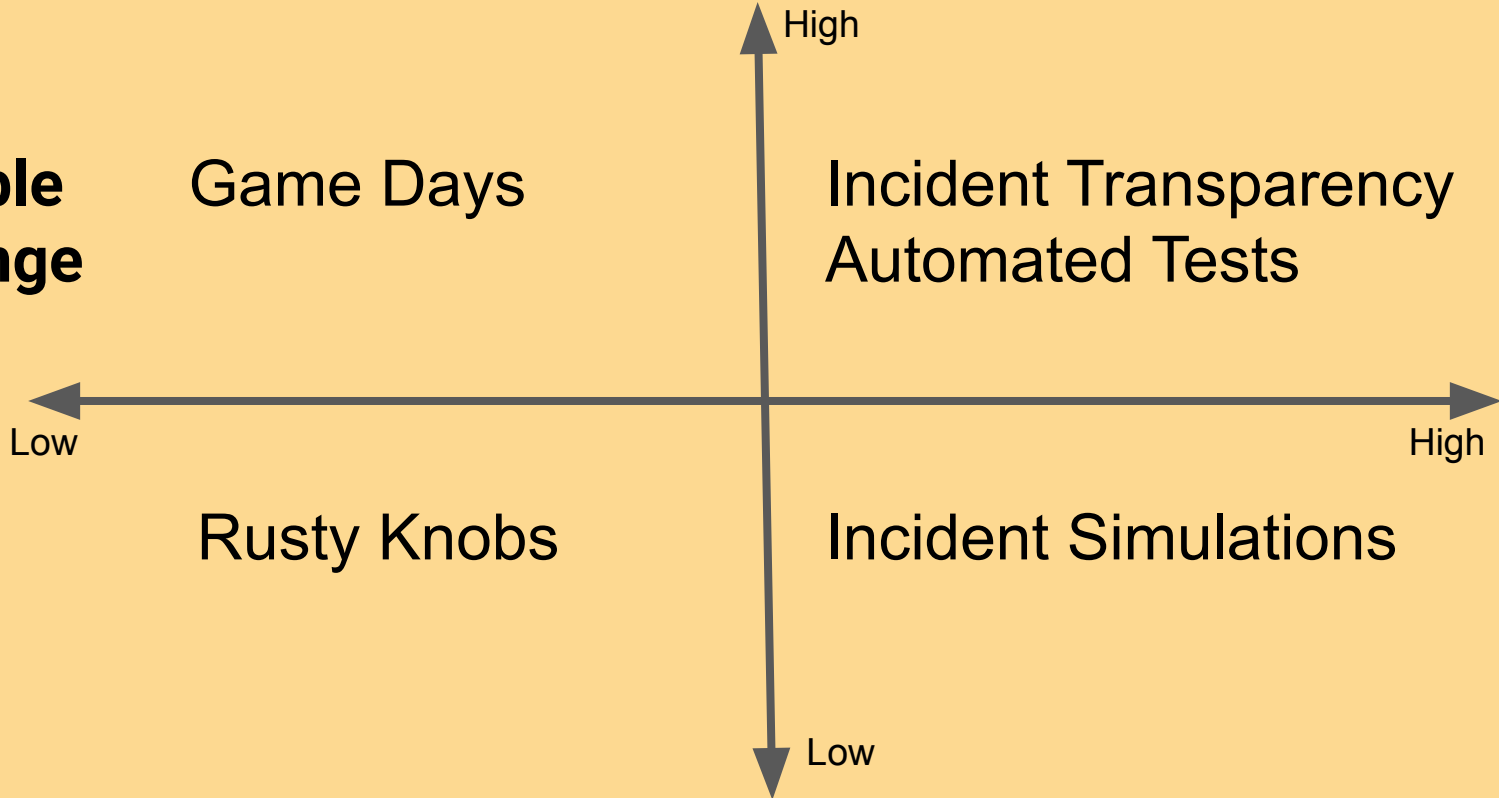
# Software Change Rate

**People  
Change  
Rate**



# Software Change Rate

**People  
Change  
Rate**





Does this really help with  
**novel failures?**

# Magic Recipe for Novel Failures

High trust teams

Fast consensus

Communication

Calculated risk-taking

Creativity

Empathy

# Training Exercise Formula

- Very clear or intuitive end goal
- Little or no added constraints
- Progressive disclosure of new information
- Randomized teams
- Fun!!



LEGO CREATOR  
16+  
LEGO CREATOR  
16+  
LEGO CREATOR  
16+

LEGO IDEAS  
16+

LEGO CREATOR  
16+

LEGO CREATOR  
16+  
10080  
2480 pcs/parts

LEGO  
70618  
Masters of Spinjitzu

LEGO  
NINJAGO  
Masters of Spinjitzu  
Ages/Age des  
9-14  
70655  
Dragon Pit  
1660 pcs/parts

LEGO  
STAR WARS  
75190  
Disney







# Case 5: Part Allocation (b)

Participant: Examine the bill on the breaking event

Scenario: It turns out this was a complete system outage that took out a whole lot of structures. Many of the parts that were affected were created and we don't know

What should we do? Who else should we talk to? How do we get you

1:59 PM



Kyle Dematt







## Summing Up

Successful SRE teams see a continual drift toward novel failures

Many practices available to prepare the people, systems, and teams for responding to failures

Don't overlook the importance of a strong team culture emphasizing trust and creativity



# Thank You!

[github.com/jarthorn/lego-incident-response](https://github.com/jarthorn/lego-incident-response)

John Arthorne  
@jarthorne

Shopify Production Engineering

