Real Talk

What We Think We Know – That Just Ain't So

~~Lies we don't know we're telling ourselves~~

~~Mass self-delusion so we can feel ok, or feel *something*~~

~~*(wink-wink, nudge-nudge, say no more)*~~

~~How dare you describe the glaringly obvious~~

John Allspaw

Adaptive Capacity Labs

```
2230 /*
2231  * If the new process paused because it was
2232  * swapped out, set the stack level to the last call
3333  * to savu(u_ssav).  This means that the return
2235  * actually returns from the last routine which did
2236  * the savu.
2237  *
2238  * You are not expected to understand this.
2239  */
```

BUT I'M GENERALLY OPTIMISTIC

“It ain’t what you *don’t know* that gets you into trouble.

It’s what you know for sure…that just ain’t so.”

Mark Twain

or maybe Josh Billings?
or Artemus Ward?
or Kin Hubbard?
or Will Rogers?
or Edwin Howard Armstrong?

The Great Wall of China is **NOT** visible from space.

A penny dropped from the Empire State building will **NOT** kill someone if it hits them in the head.

“We cannot call it a scientific field unless we can admit we’ve gotten things wrong in the past.”

– David Woods


Today
DAVID WOODS
Ohio State University
Ergonomics Professor
C-SPAN 2

Our Team

John Allspaw

John Allspaw has worked in software systems engineering and operations for over twenty years in many different

Beth Adele Long

Beth Adele Long is a writer and software engineer with over twenty years of experience building, maintaining, and repairing

Dr. David Alderson

David L. Alderson teaches at the Naval Postgraduate School where he is Professor in the Operations Research Department and

Dr. David Woods

Dr. David Woods founded Resilience Engineering as an approach to safety in complex systems in 2000-2003 as part of the response

“We cannot call it a scientific field unless we can admit we’ve gotten things wrong in the past.”

– David Woods

ADAPTIVE TACTICS IN SOFTWARE OPERATIONS: HOW PEOPLE MANAGE COMPLEXITY BY ASKING FOR HELP

Michael Wettick | LUND UNIVERSITY

“The greatest obstacle to discovery is not ignorance - it is the illusion of knowledge.”

Daniel J. Boorstin

Bottom Line, Up Front

1. On engineering productivity
2. On changes
3. On shallow incident data
4. On sequences
5. On “repeat” incidents
6. On incident response

1. On engineering productivity

“The use of lines of code metrics for productivity and quality studies [is] to be regarded as professional malpractice starting in 1995.”

Capers Jones (1994). “Assessment and Control of Software Risks”, Prentice Hall

2. On “Changes”

“Change is one of the leading causes that induce incidents.”

Wu, Y., Chai, B., Li, Y., Liu, B., Li, J., Yang, Y., & Jiang, W. (2023, May). An empirical study on change-induced incidents of online service systems. In *2023 IEEE/ACM 45th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)* (pp. 234-245). IEEE.

Gartner Research

Causal Analysis Makes Availability and Performance Data Actionable

Published: 07 October 2015

Summary

There are five distinct types of causal analysis that make availability and performance data

“85% of Performance Incidents Can be Traced to Changes”

2. On “Changes”

Change is one of the leading causes of *resolving* incidents.

Change is one of the leading causes of incidents that do NOT happen.

3. On Shallow Data

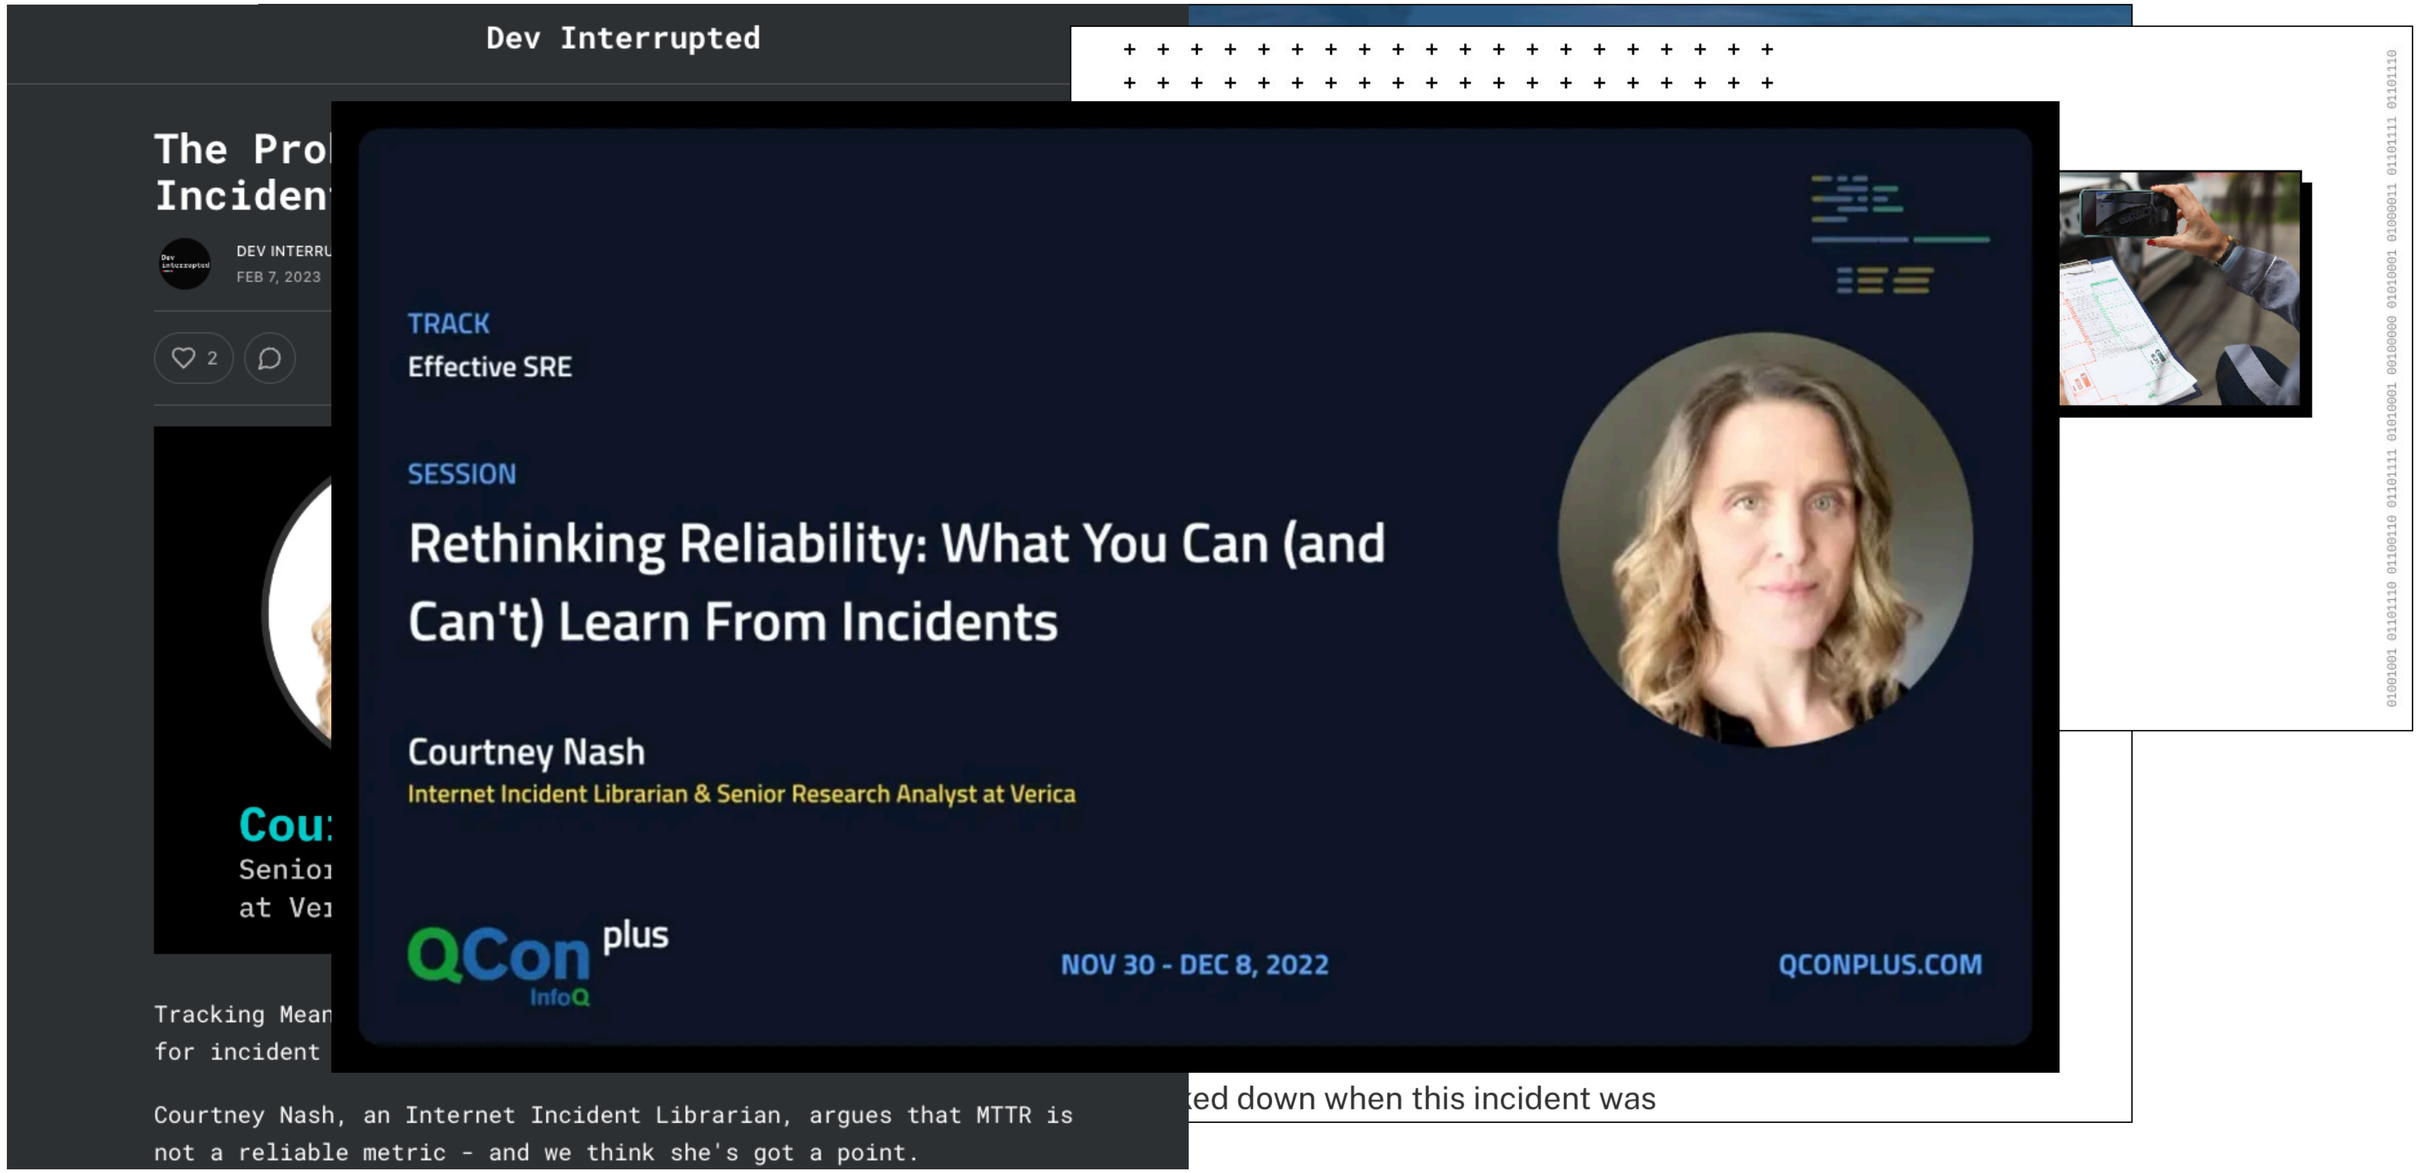

4. On “The Sequence™”

It’s really quite simple!

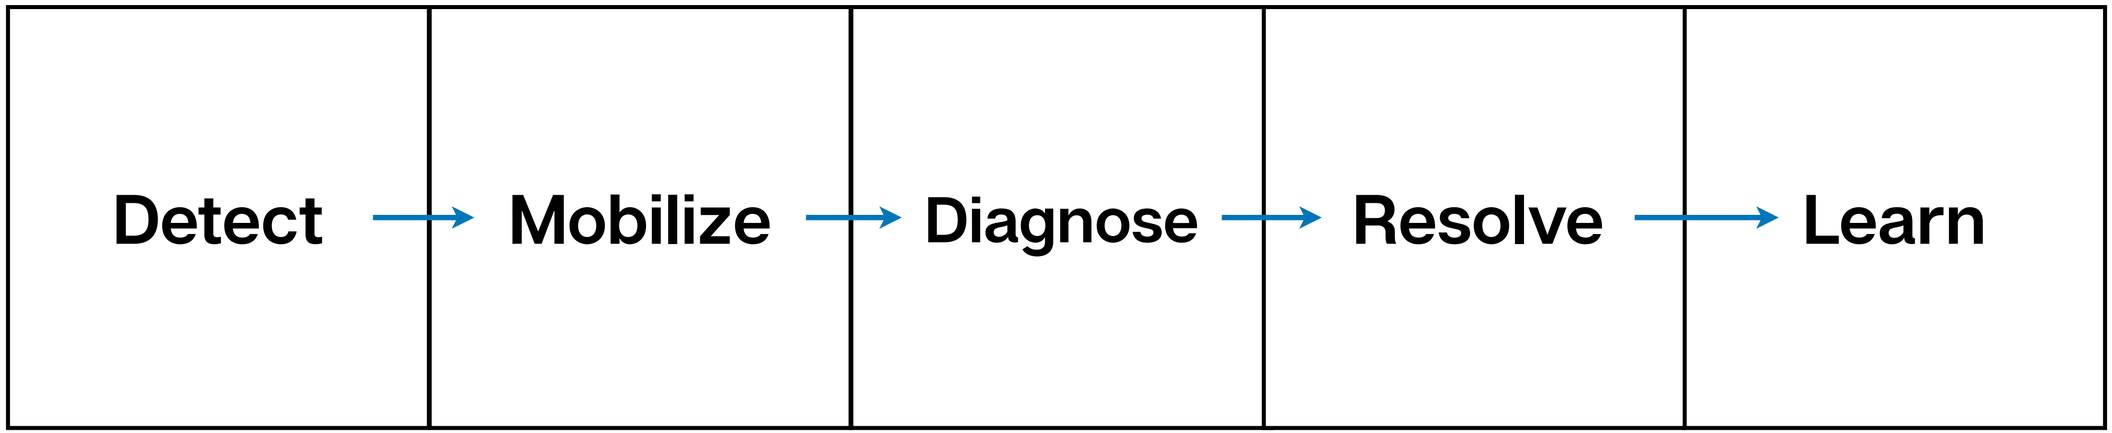

4. On “The Sequence™”

Here’s a curly version of the same thing!

4. On “The Sequence™”

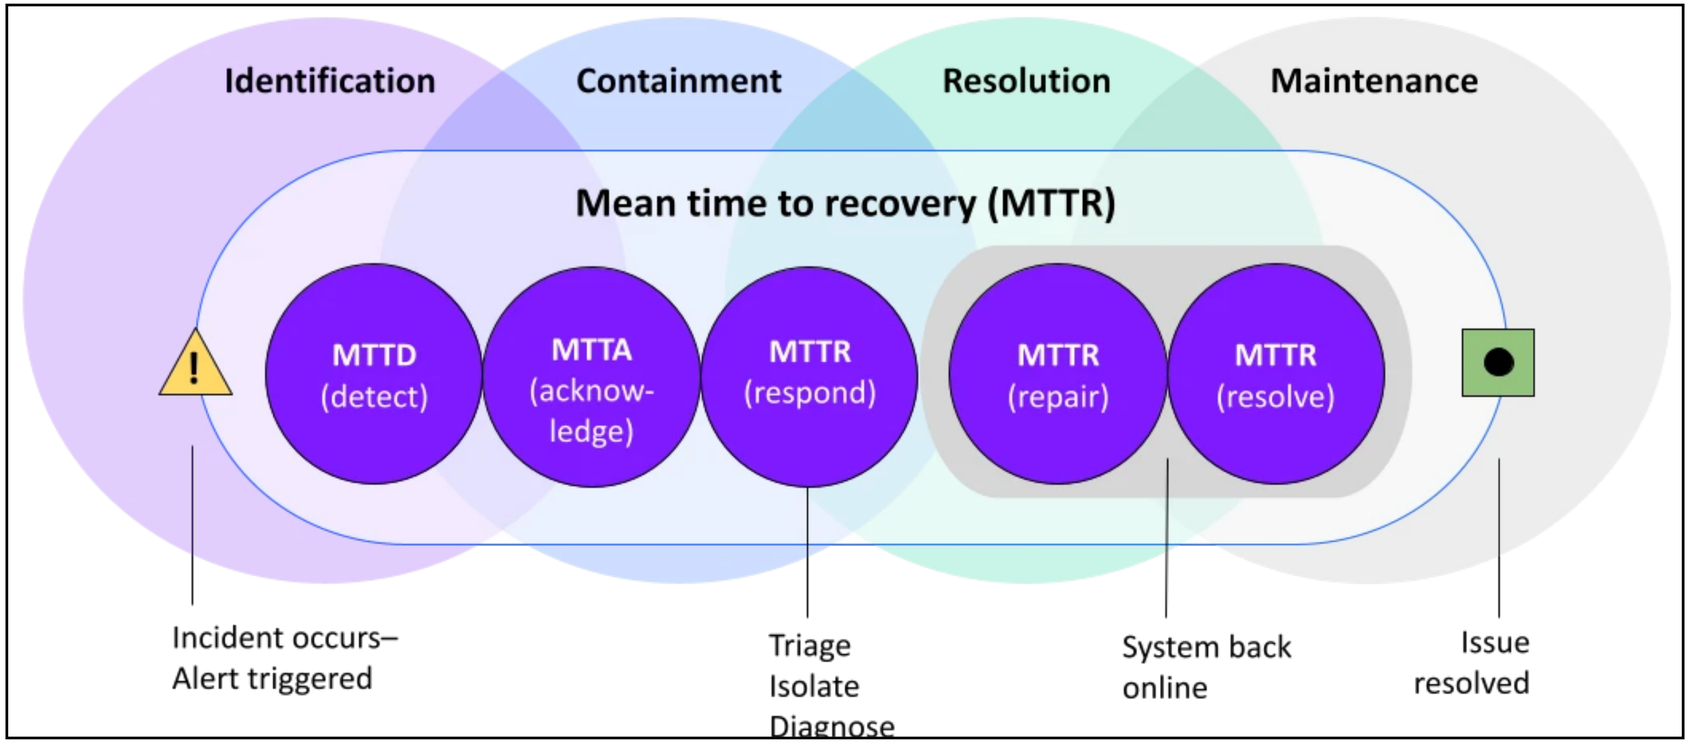

Here’s another version that has some sweet 50% opacity circles.

4. On “The Sequence™”

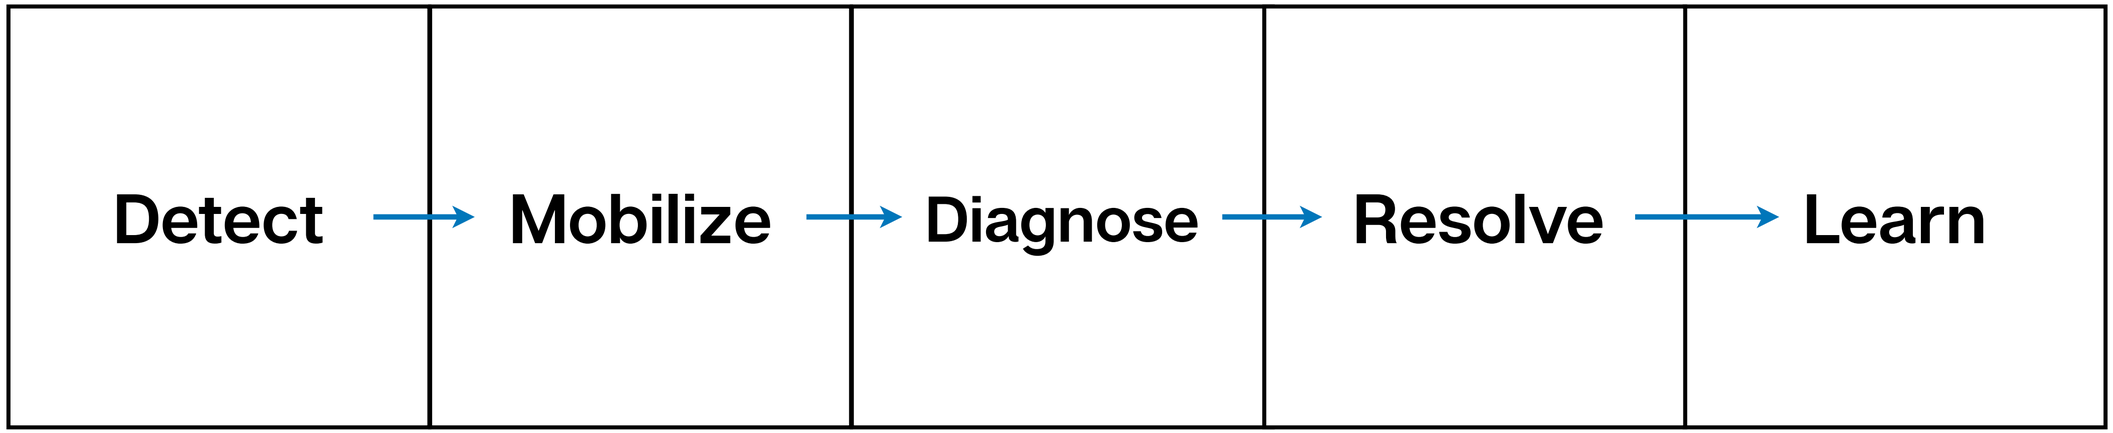

Is this how all incidents play out?

Is this how *most* incidents play out?

Can we assume *enough* incidents play out this way...*ish*...that we can ignore or dismiss those that don’t?

4. On ce™”

Physicists are like:
“To simplify calculations, we will assume the cat is cubical.”

Detect → Learn

Is this how all ir

Is this how *mos*

Can we assum hat we can ignore or dismiss thos

https://www.itsonlyamodel.com/

"It's Only a Model"

| 'mädl | *A schematic description or representation of something, especially a system or phenomenon, that accounts for its properties and is used to study its characteristics.*

- **"The map is not the territory."** (A. Korzybski, 1933) A model is not the thing or things it is intended to represent. This means that every model will be an imperfect representation.
- **"All models are wrong, but some are useful."** (G. Box, 1976) Models can be used for specific or general purposes and are inherently context-sensitive to those purposes. For example, reusing a model intended for teaching a concept may or may not be helpful as one intended to demonstrate an example of a concept.
- **"It is only when a breakdown occurs that we become aware of the fact that 'things' in our world exist"** (T. Winograd & F. Flores, 1987) A model is usable until the assumptions made during its creation turn out to be inadequate, often as part of a fundamental surprise. There is little incentive to update an otherwise working model.

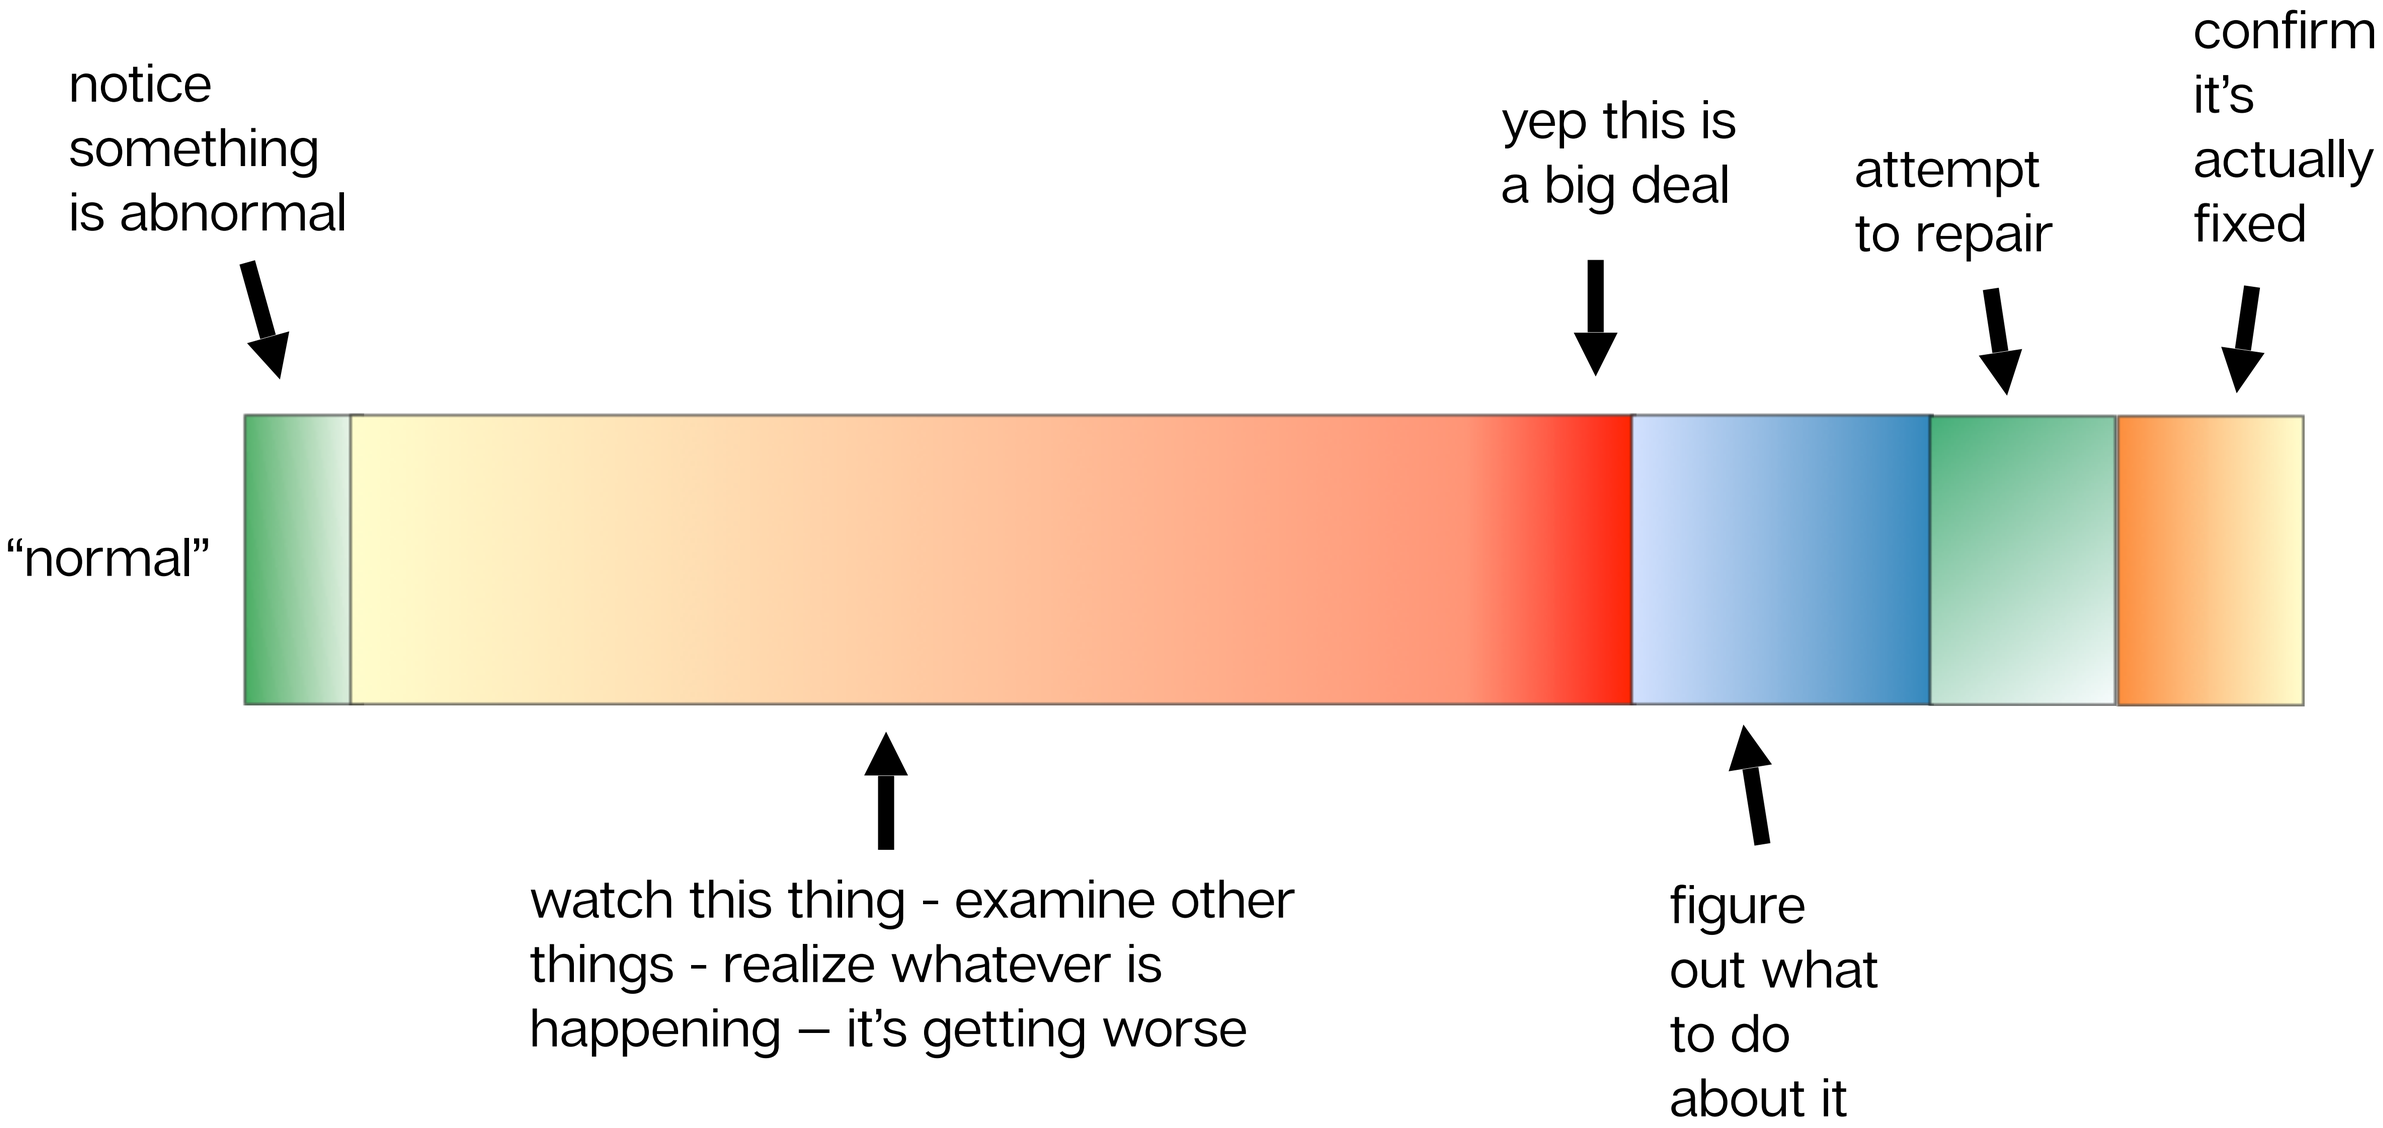

(real incident)

4. On “The Sequence™”

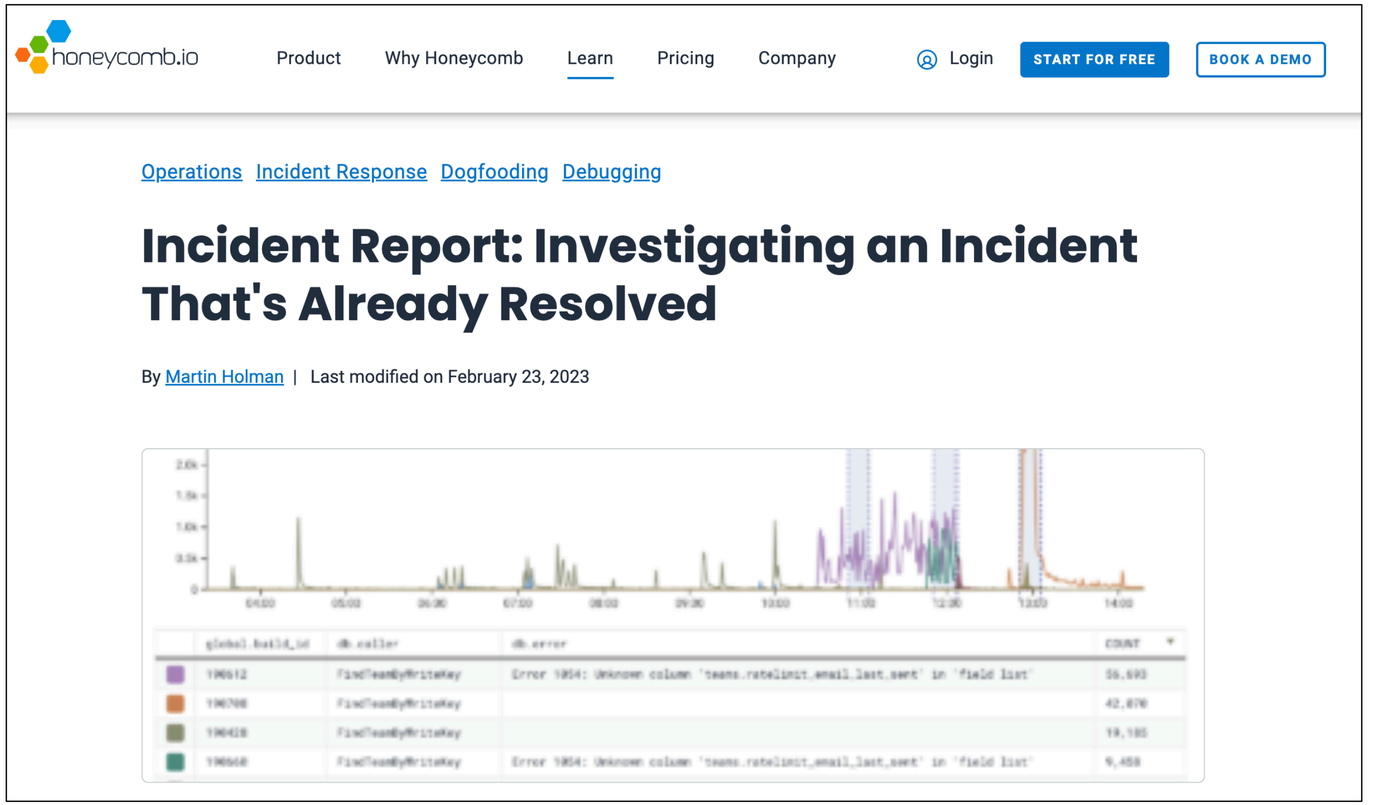

“But if we include this...it has a negative Time-To-Resolve...?”

4. On “The Sequence™”

4. On “The Sequence™”

Resolve

Detect

Diagnose

4. On “The Sequence™”

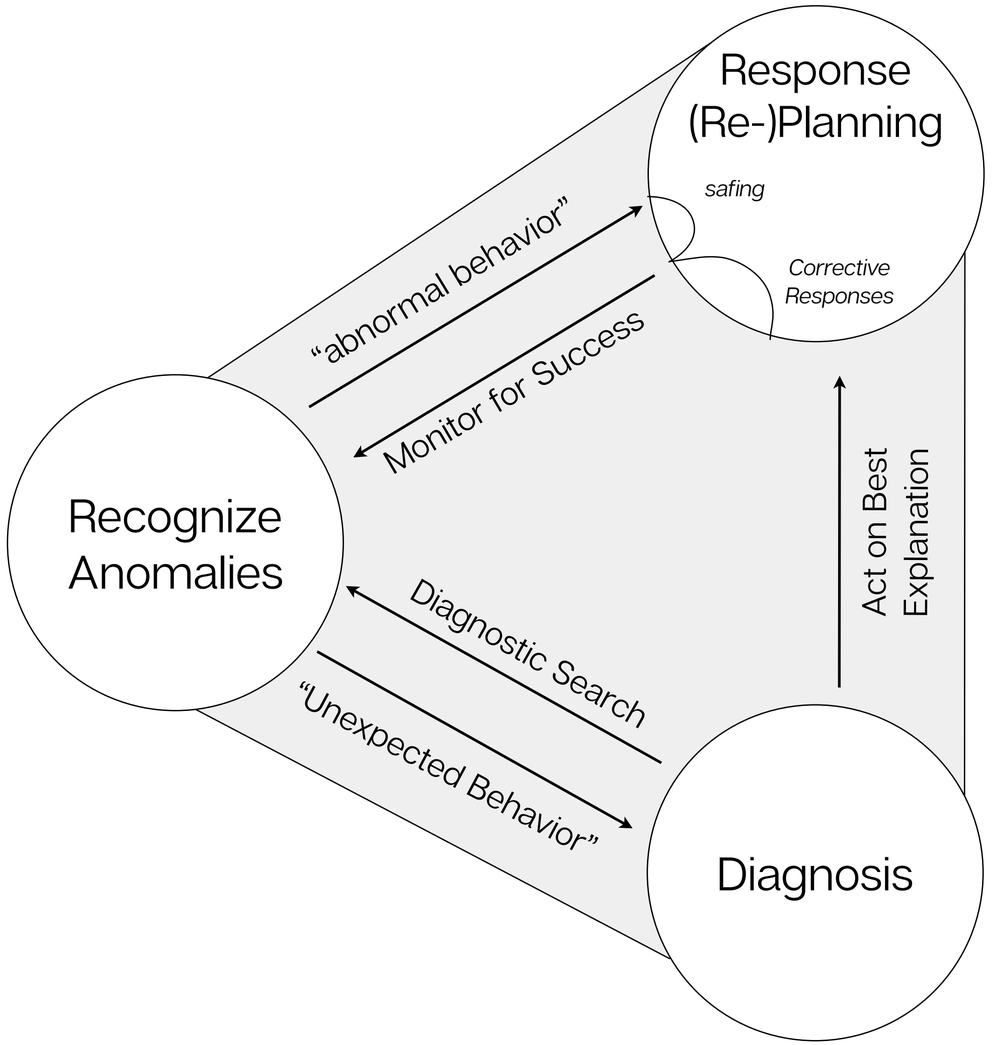

4. On “The Sequence™”

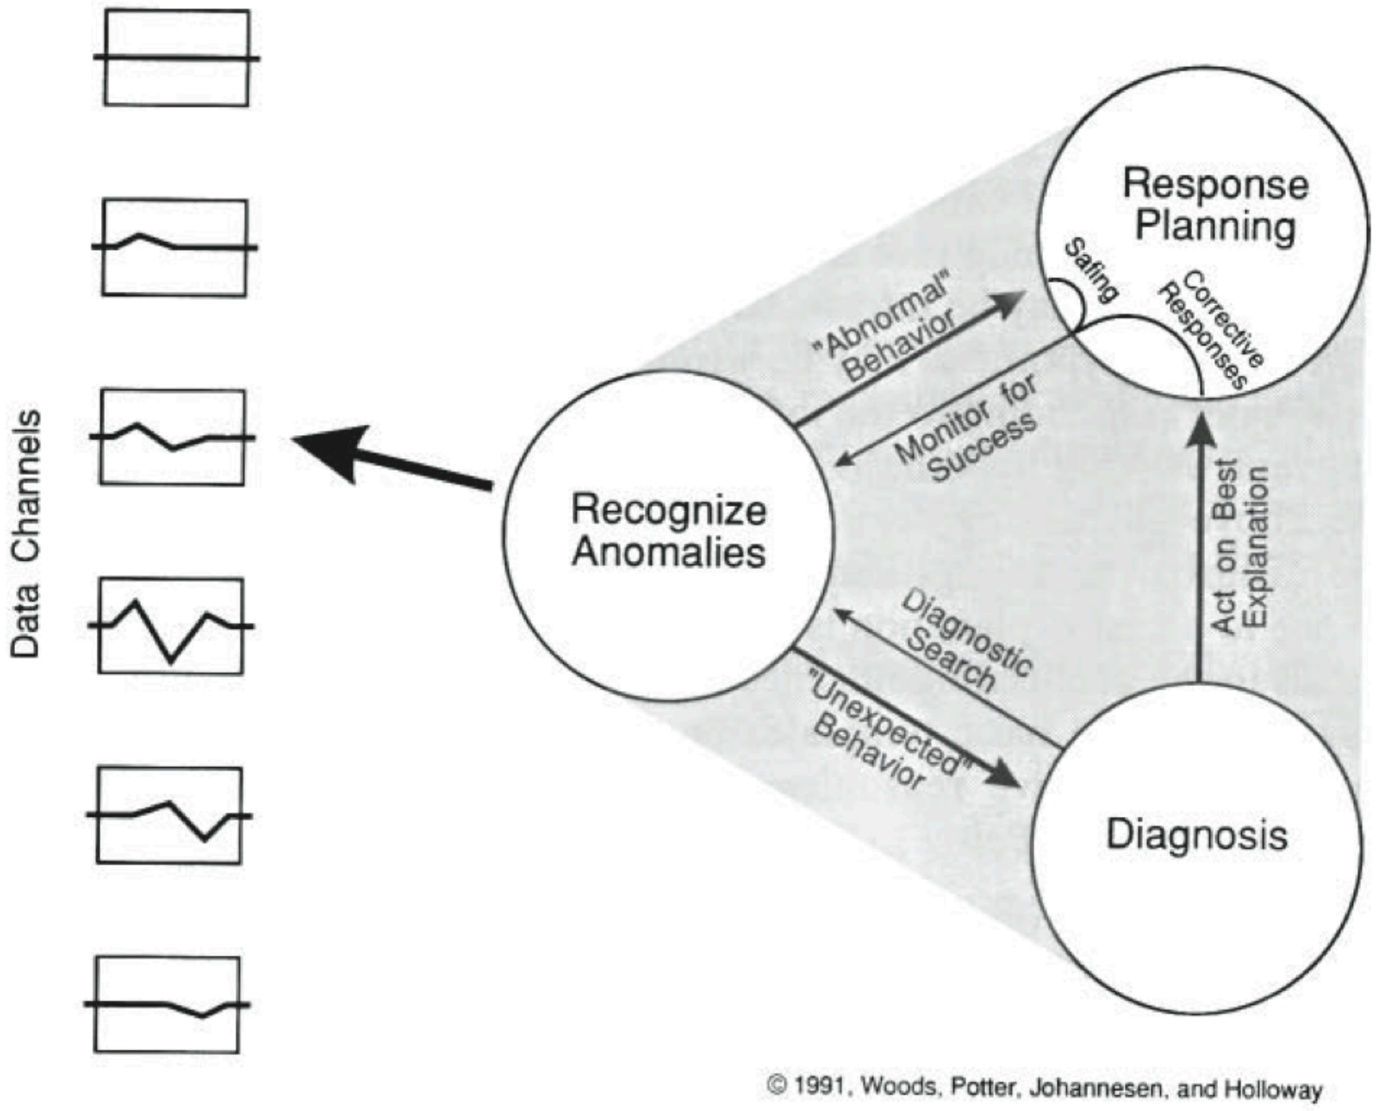

5. On “repeat” incidents

The criteria used in ***labeling*** an incident a “repeat” matters more than the “repeat” happening.

Who gets to label an incident as a “repeat” can matter a great deal.

5. On “repeat” incidents

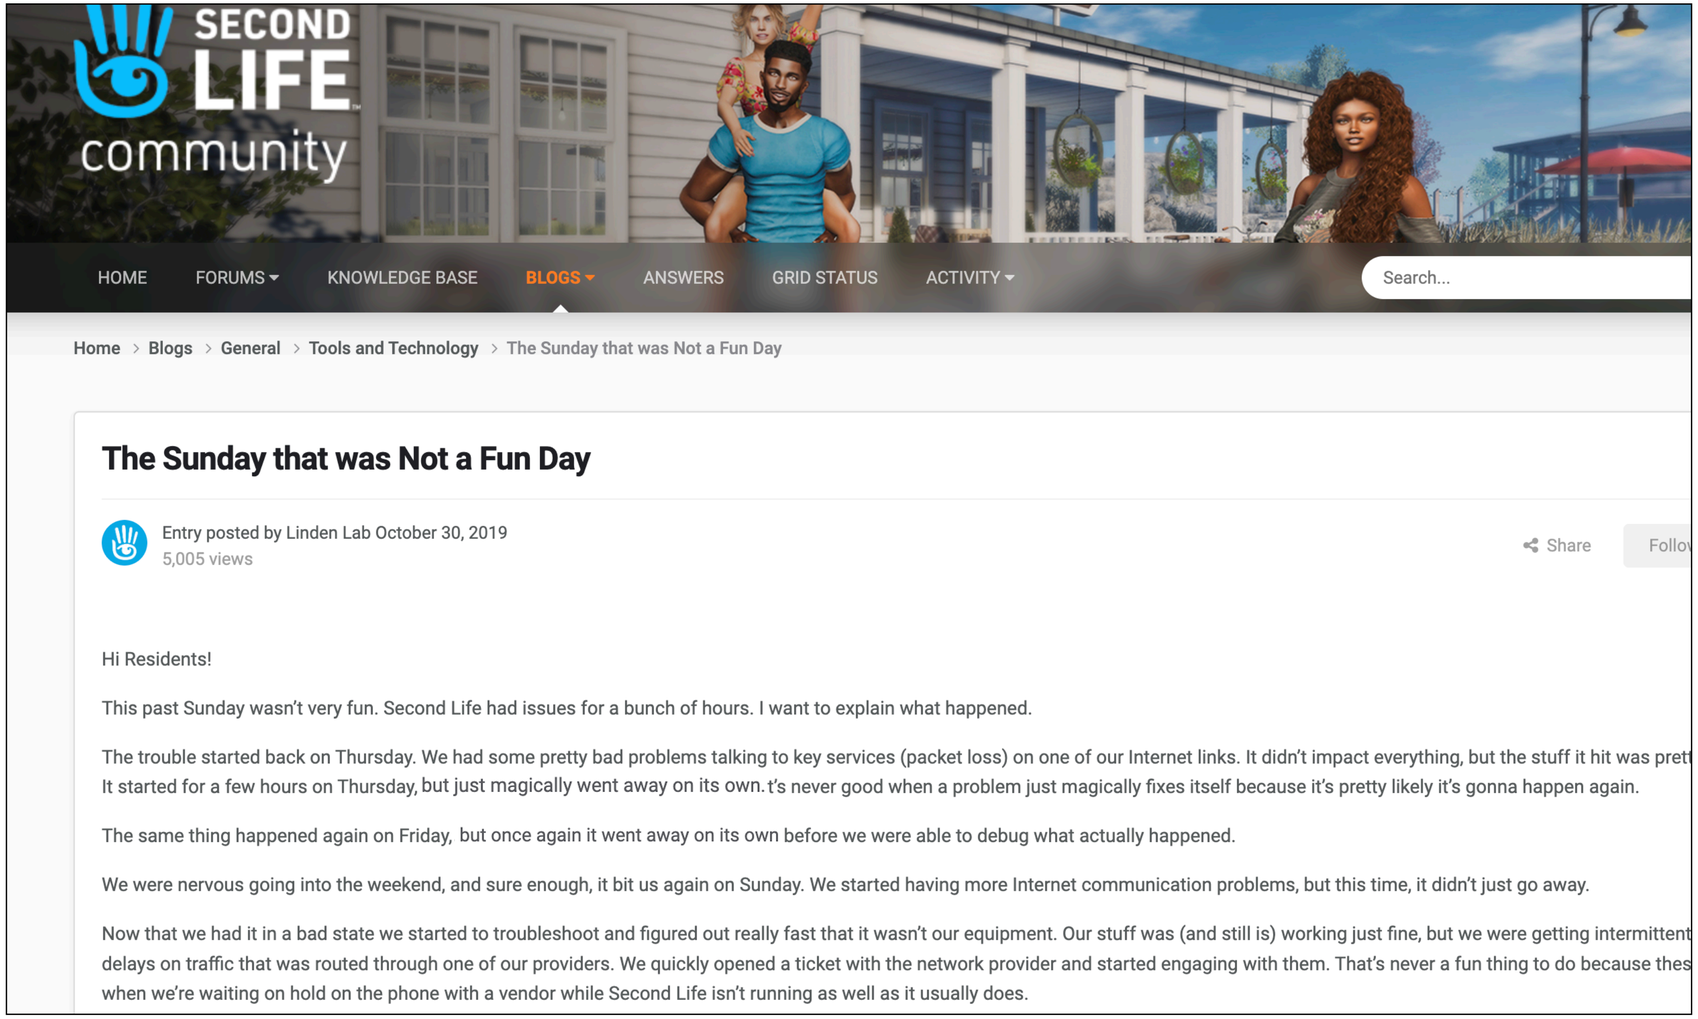

6. On incident response

An organization can be the most skilled and efficient at keeping stakeholders up to date about ongoing incidents and still be terrible at learning from them or responding to them

6. On incident response

The best scenario when it comes to responding to an incident is:

a. the people responding can recognize immediately what is happening, and
b. know exactly what to do about it.

Anything that can bolster people's expertise in support of those two things is paramount.

Everything else is secondary.

And when this happens…the event is often not even labeled as an "incident."

- Dunbar's Number is wrong
- The "blameless" stance is specifically for *accidents and mistakes*
- There is no objective *start time* of an incident
- There is no objective *end time* of an incident
- Severity is always a negotiable classification
- The absence of incidents is not evidence learning from incidents is happening
- The presence of incidents is not evidence learning from incidents is *not* happening
- Incident responders follow "red herrings" in incidents because they are confident they are *not* red herrings
- Psychological safety is necessary…but not sufficient!
- The more people in a company assert their post-incident activities are "blameless", the less likely it is true
- Categorizing incidents depends critically on who gets to make the categories…who are often never those who respond to them
- An organization can be the most skilled and efficient at keeping stakeholders up to date about ongoing incidents and still be terrible at learning from them or responding to them
- Timelines mandated to complete post-incident work that are inversely proportional to their severity, impact, or difficulty is backwards and nonsensical
- "…to ensure it never happens again" is the most disingenuous remark found in public incident articles
- Template-based post-incident write-ups are akin to coloring books
- Conway's Law tends to be cited when it serves to forward the agenda of those citing it, and rarely otherwise
- Stories about incidents told by those who responded to them will always have notable details that official write-ups don't include

Amazing and inspiring closer slide

Thanks!