BUILDING RESILIENCE WITH CHAOS ENGINEERING

prime video

WHAT WE WILL COVER TODAY

COMMON QUESTIONS — FAQ

EPL TECHNICAL PROGRAM APPROACH — BUSINESS RESULTS

GOALS AND EXAMPLES — PRACTITIONER EXAMPLE

LOOKING AHEAD — DAY ONE

Q&A

prime video

01

LET'S START WITH QUESTIONS

THREE FREQUENT QUESTIONS

1.WHY INVEST IN CHAOS ENGINEERING?

2. WHERE DO I START?

3. HOW DO I MAKE MEANINGFUL PROGRESS?

prime video

DECISION

JOURNEY

prime video

INPUT TO DECISION:
START HERE

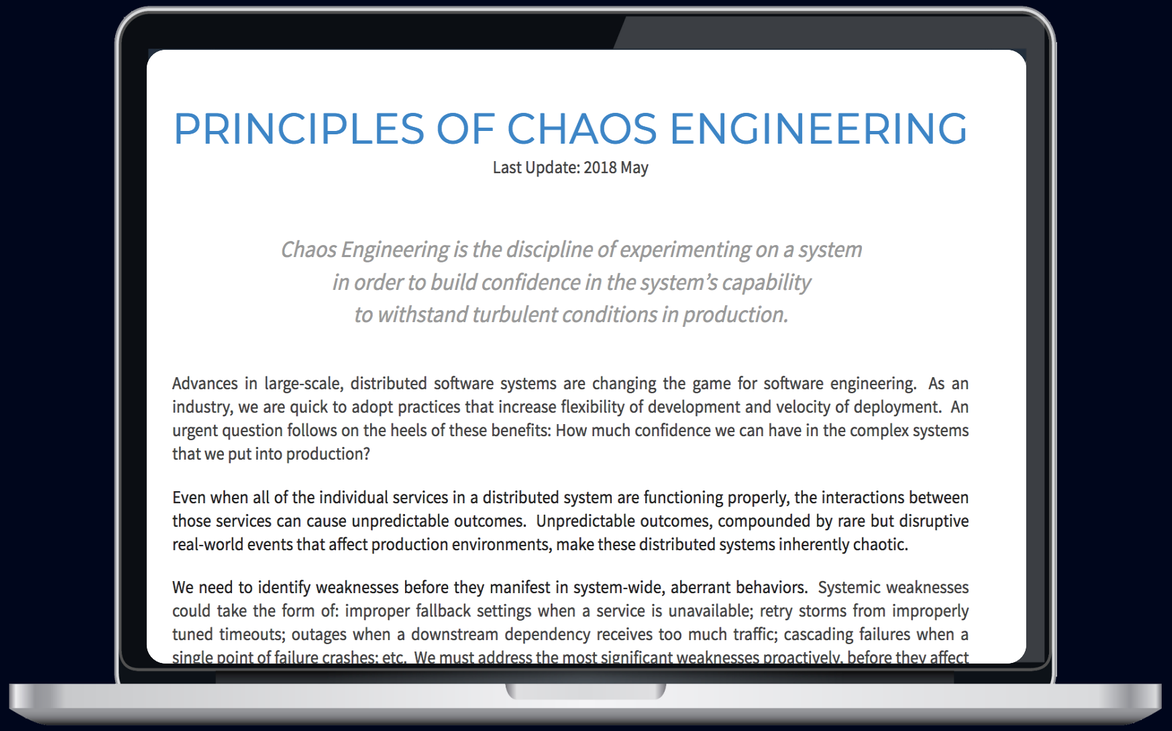

Chaos Engineering Principals

principlesofchaos.org

Chaos engineering is about breaking things in a controlled environment and through well-planned experiments in order to build confidence in your application to withstand turbulent conditions.

INPUT TO DECISION:
DOES THIS LOOK FAMILIAR?

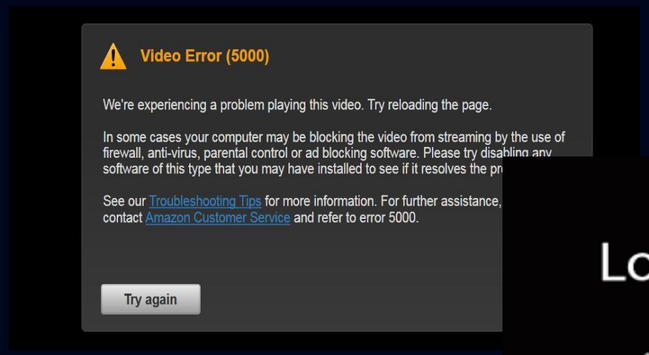

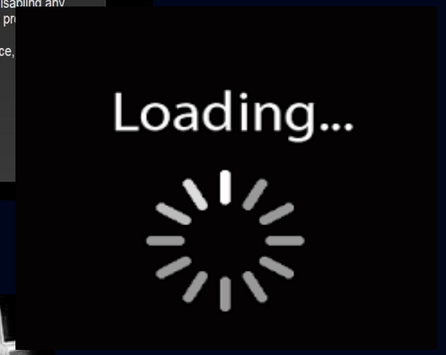

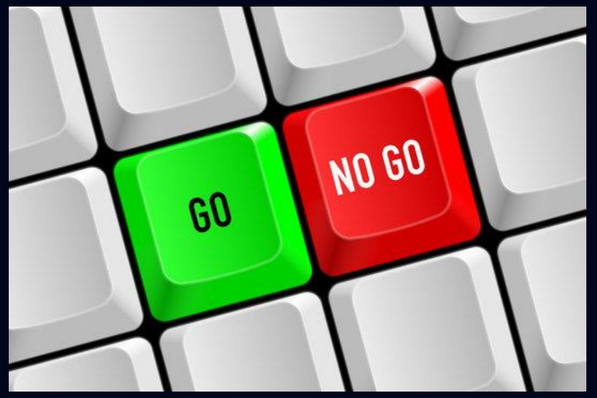

INPUT TO DECISION:
DISTRIBUTED SYSTEMS ARE HARD

- Errors happen anytime , often in combination with other errors.
- Results of network operations can be unknown *(succeeded, failed, or received but not processed).*
- Problems occur at all logical levels.
- Problems get worse at higher levels of the system, due to recursion.
- Bugs often show up long after they are deployed to a system.
- Bugs can spread across an entire system.

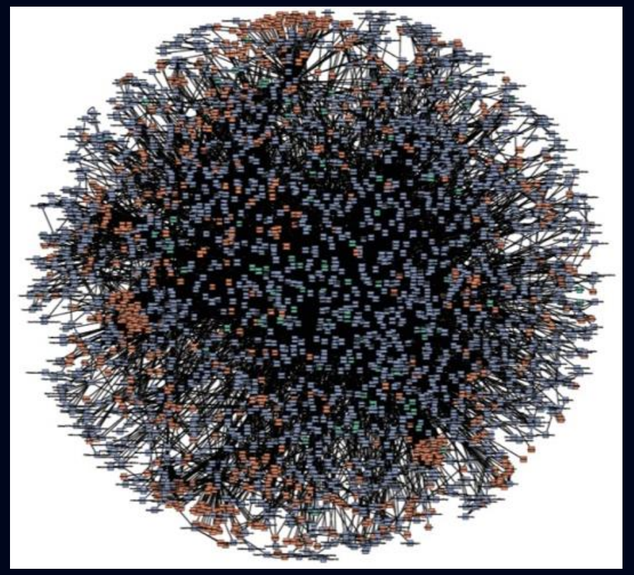

Amazon

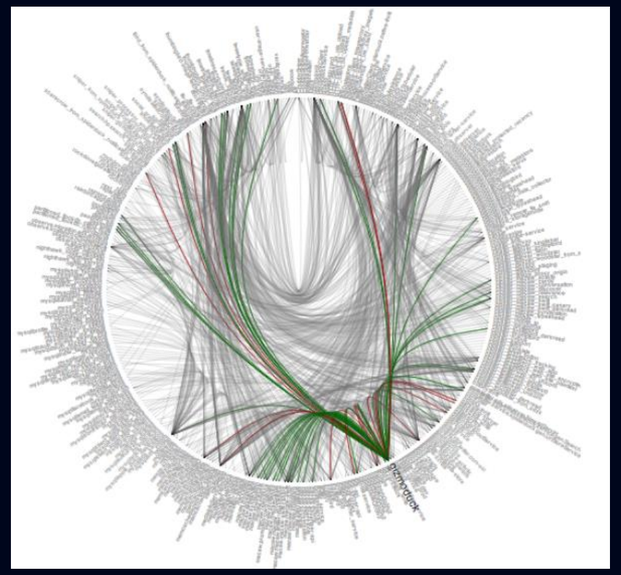

Twitter

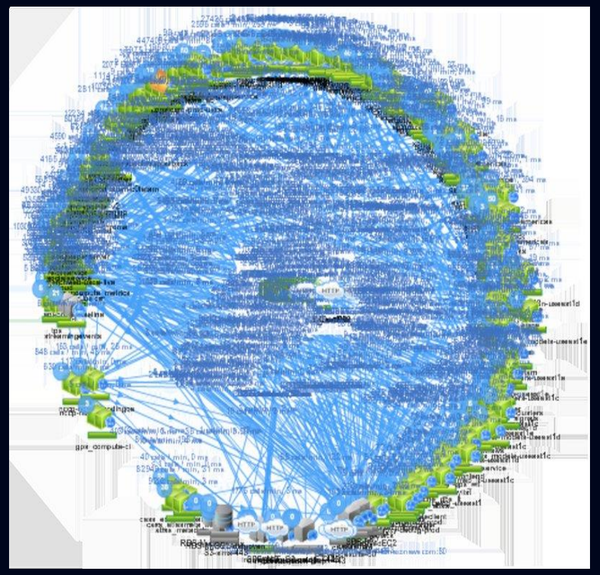

Netflix

INPUT TO DECISION:
ASK IF TRADITIONAL TESTING IS ENOUGH?

Testing: verifying a **KNOWN** condition:

e.g. assert(A = B) ?

Unit testing of components:

- Tested in isolation to ensure function meets expectations.

Functional testing of integrations:

- Each execution path tested to assure expected results.

“Failures are a given and everything will eventually fail over time”.

Werner Vogels
CTO – Amazon.com

DECISION MADE!
BUILDING CHAOS JOURNEY

02

WHERE TO START

"A JOURNEY OF THOUSAND MILES BEGINS WITH A SINGLE STEP." – LAO TSU

INPUT TO JOURNEY:
HIRE THE RIGHT TALENT

Jesse Robbins, “Master of Disaster”

In 2001, a volunteer firefighter applied for two jobs: a bus driver and a backup systems engineer at Amazon.com.

In the decade that followed, he transformed the way Web companies design and manage complex networks.

- Created **GameDay** to purposefully create regular major failures.
- Founded **Chef**, the Velocity Web Performance & Operations Conference.

INPUT TO JOURNEY:
UNDERSTAND YOUR INITIAL OBJECTIVES

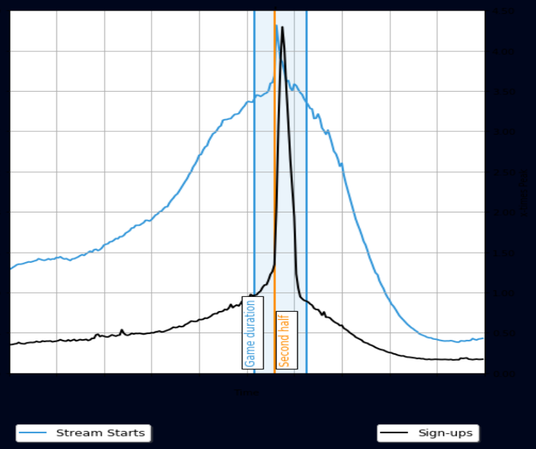

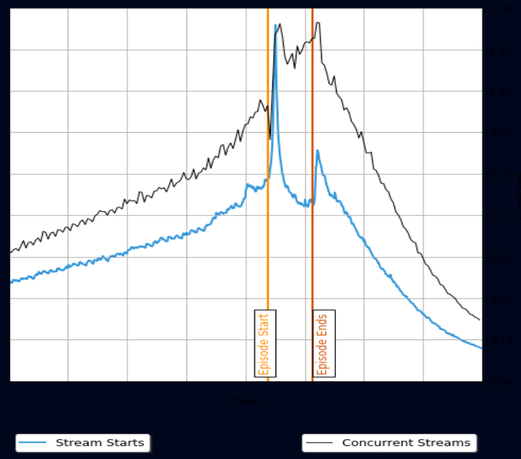

1. BUSINESS CASE
   - A specific event vs capabilities launch

2. TRENDS
   - What is your systems profile, how much has been developed within last 3 years?
   - What is the #1 source of failures? What is the depth of dependencies?

3. CULTURE
   - How strong is your technical operations culture?
   - Is this about evergreen, always ready, state of tech?

GOALS PLAYBOOK

1. PRIORITIZE BUSINESS CASE

2. MEASURE TRENDS, SET KPIS AND IMPROVEMENT GOALS

3. AIM TO CHANGE CULTURE

prime video

PRIME VIDEO PEAK EVENTS AND LAUNCHES

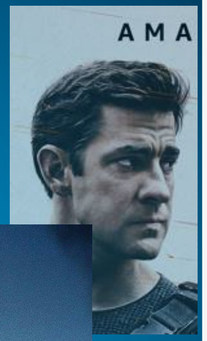

NFL NETWORK
THURSDAY NIGHT FOOTBALL
FOX
prime video

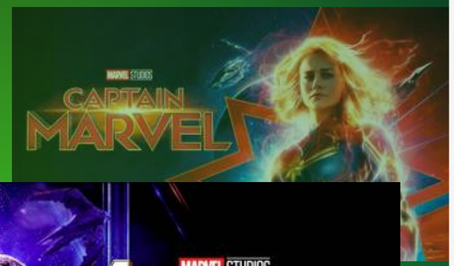

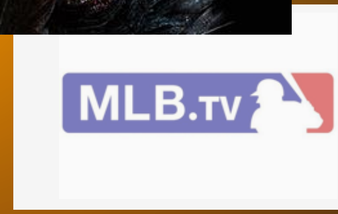


SCALING

INPUT TO JOURNEY:
UNDERSTAND YOUR BUSINESS CASE

Number of Customers In System = Arrival Rate x Mean Time In System

$$L = \lambda x W$$

Arrival rate matters greatly for live/and or exclusive video.

Customer arrival rate

Discover

Select

Transact

λ

Session starts

Concurrency

Play

L W

Concurrent sessions

PRIME VIDEO JOURNEY:
FROM MANUAL CHAOS TO AUTOMATED RESILIENCE

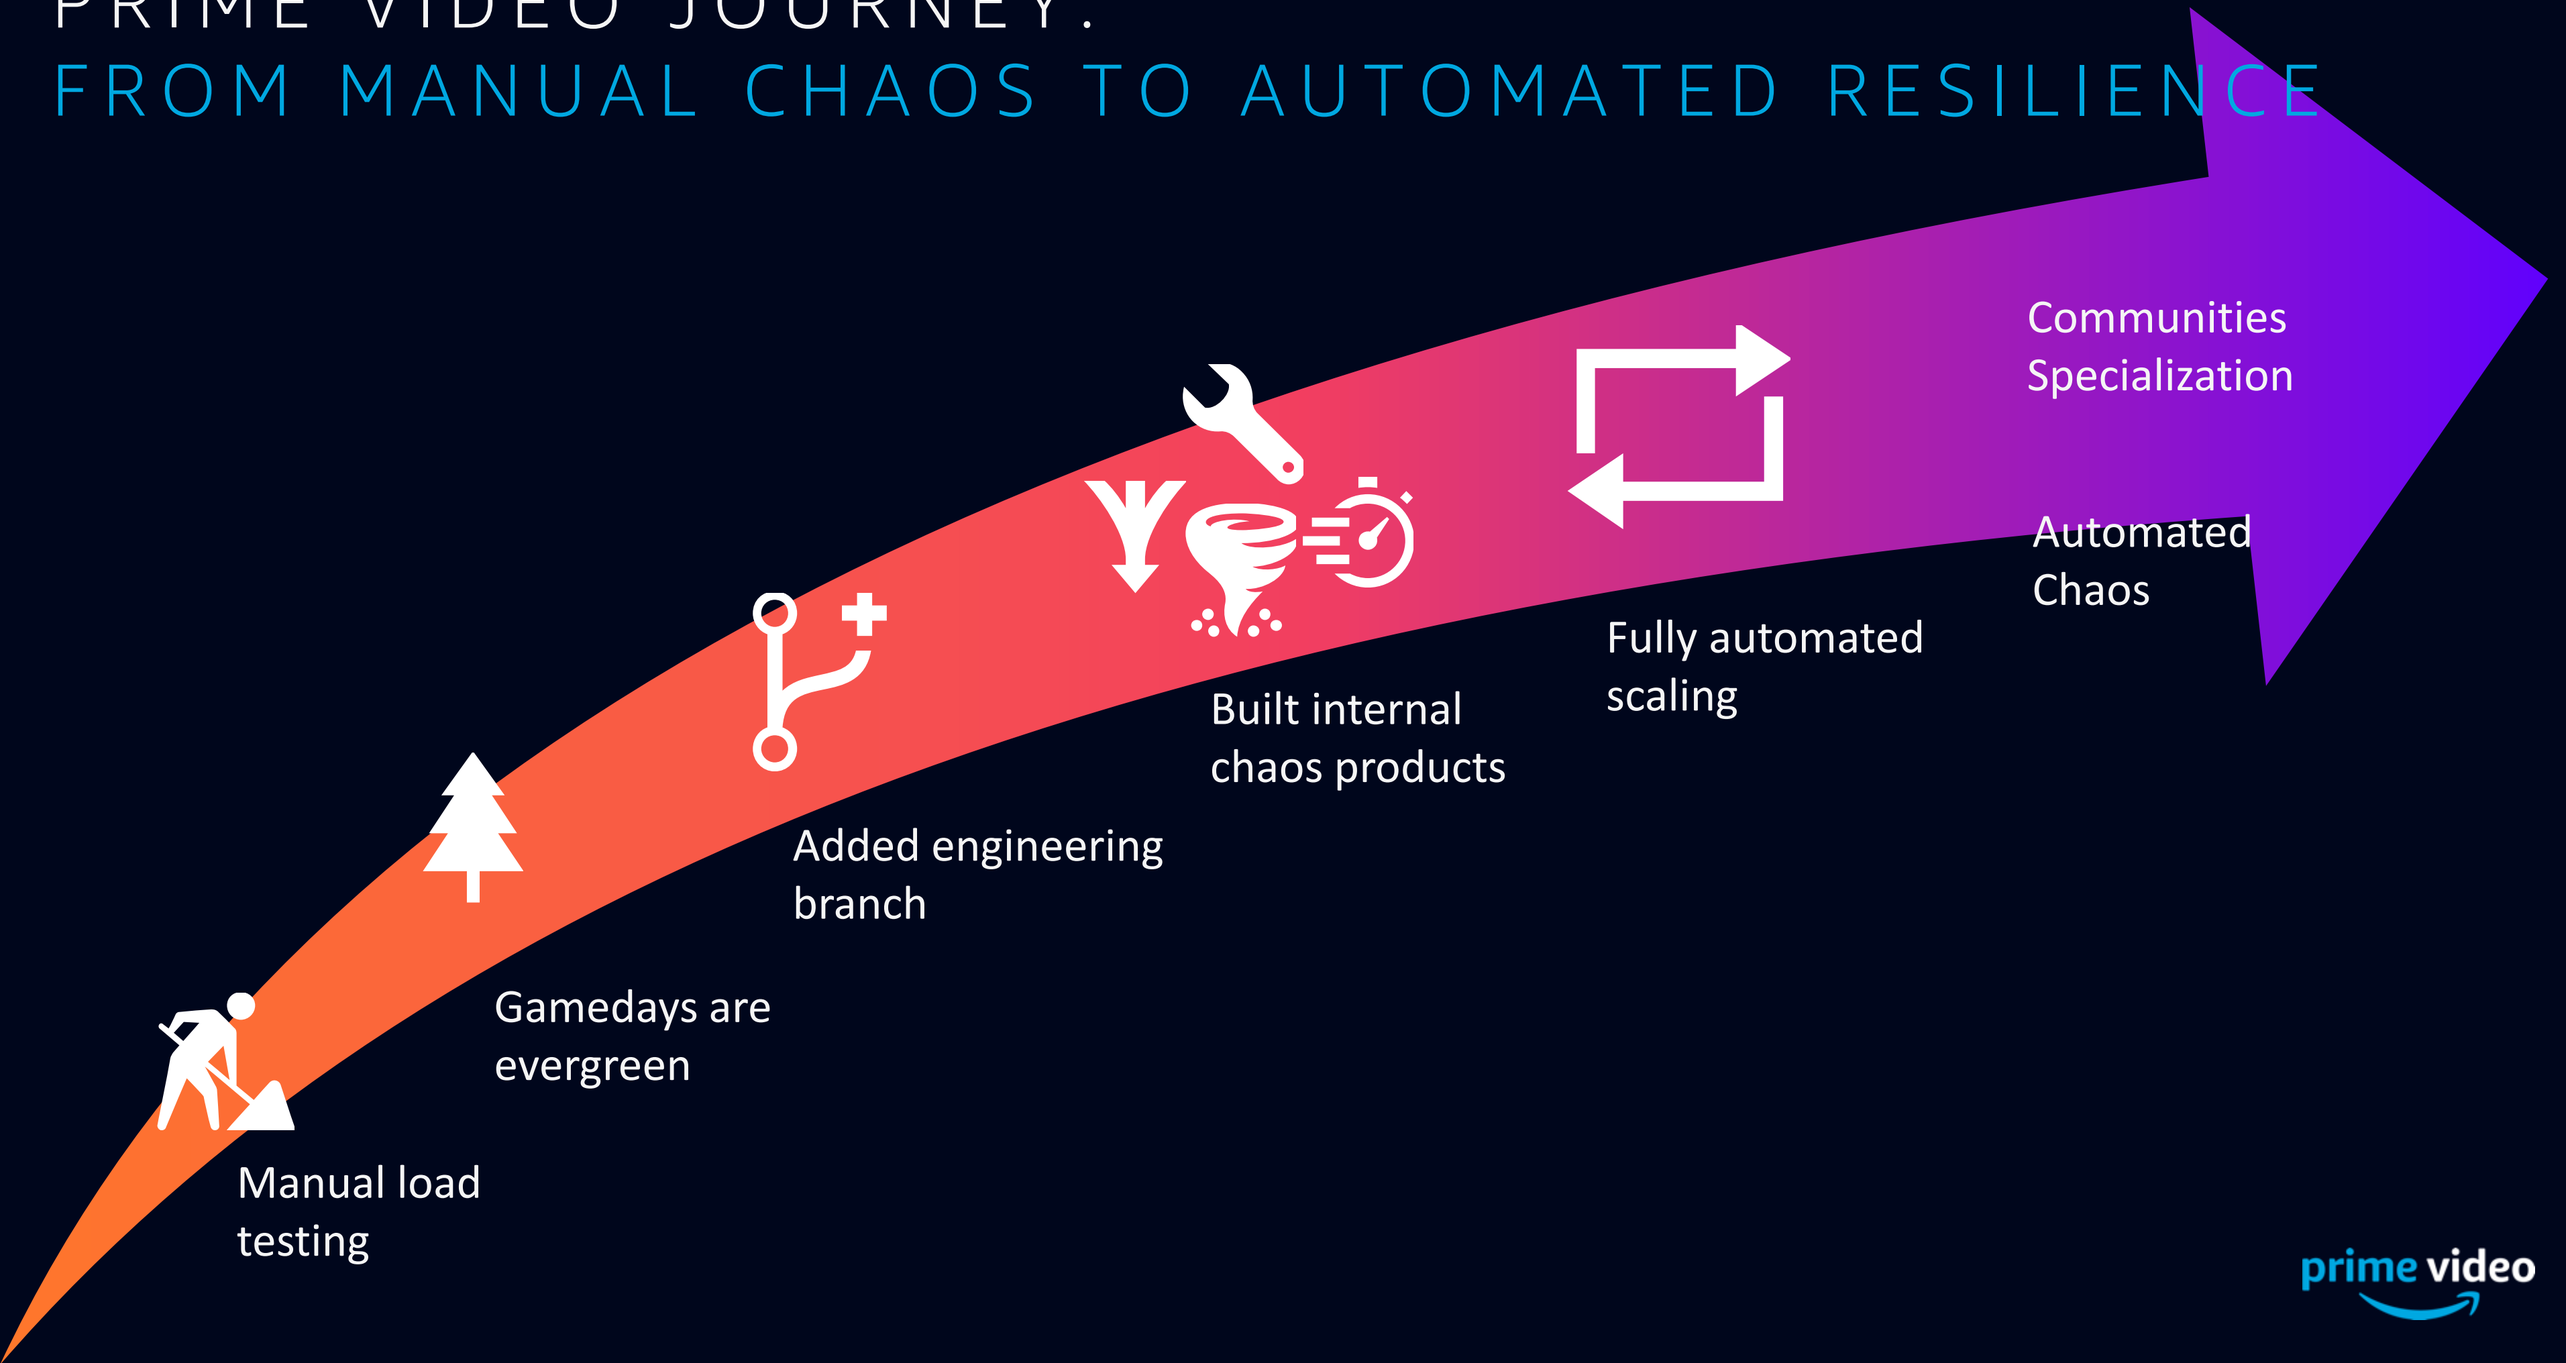

prime video

Evergreen programs

- Availability
- Scaling
- Resilience
- Peak Readiness
- Performance

Program structure

- Central engineer team
- Infrastructure TPMs team
- Infrastructure efficiency and optimization
- Opt-in SPOCs in each team (SRE equivalents)

INPUT TO JOURNEY:
CHALLENGES AND PERCEPTIONS

- Starting is perceived as hard!
- No time or flexibility to simulate disasters.
- Teams already spending all of its time fixing things.
- Can be political.
- Leads to deep conversations.
- Deeply invested in a specific technical roadmap (micro-services)

- Start with a specific business case
- Start with a process simulation, bring automation later
- Bring fun & depth into “doing things”
- Engage your business partners, take joint goals
- Trade Offs are ok!
- One tech bar, multiple mechanisms to achieve it

03

HOW TO MAKE MEANINGFUL PROGRESS?

START WITH GOALS, NOT RESOURCES!

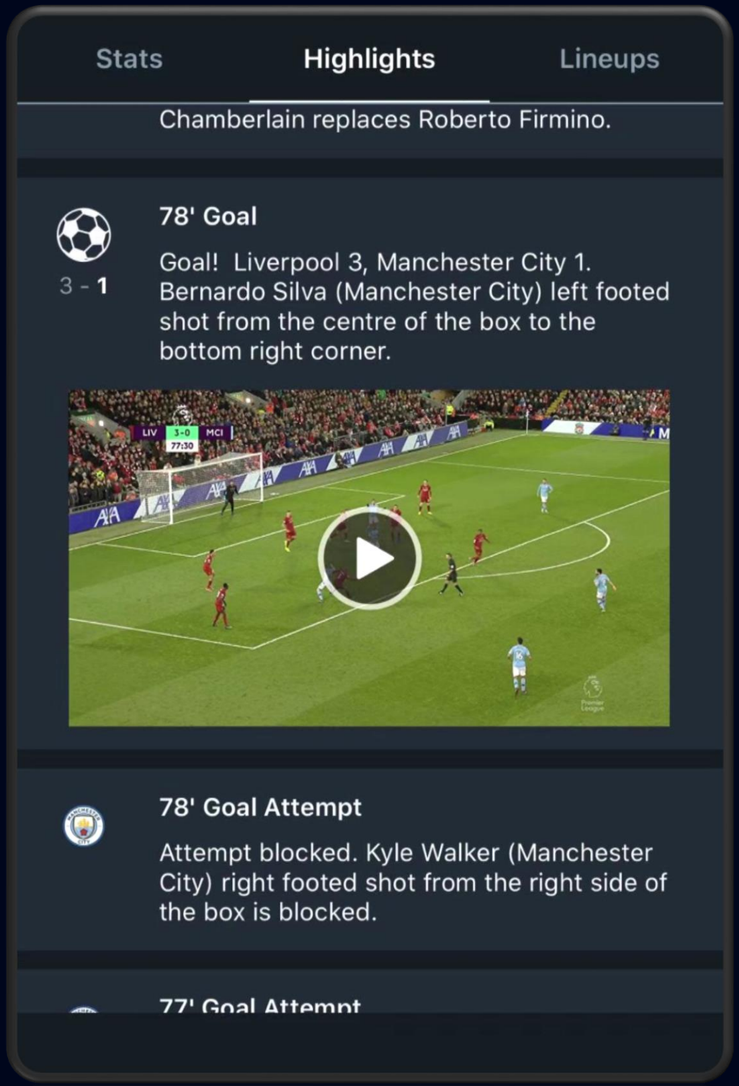

prime video

GOALS STRATEGIES

BREADTH VS DEPTH

SPECIFIC VS GENERIC E.G. LAUNCH OR CAPABILITIES

SINGLE OWNER VS. PARTNERSHIP

COVERAGE VS EXPERIMENTATION

prime video

FAILURE INJECTION TEST (FIT) GOALS

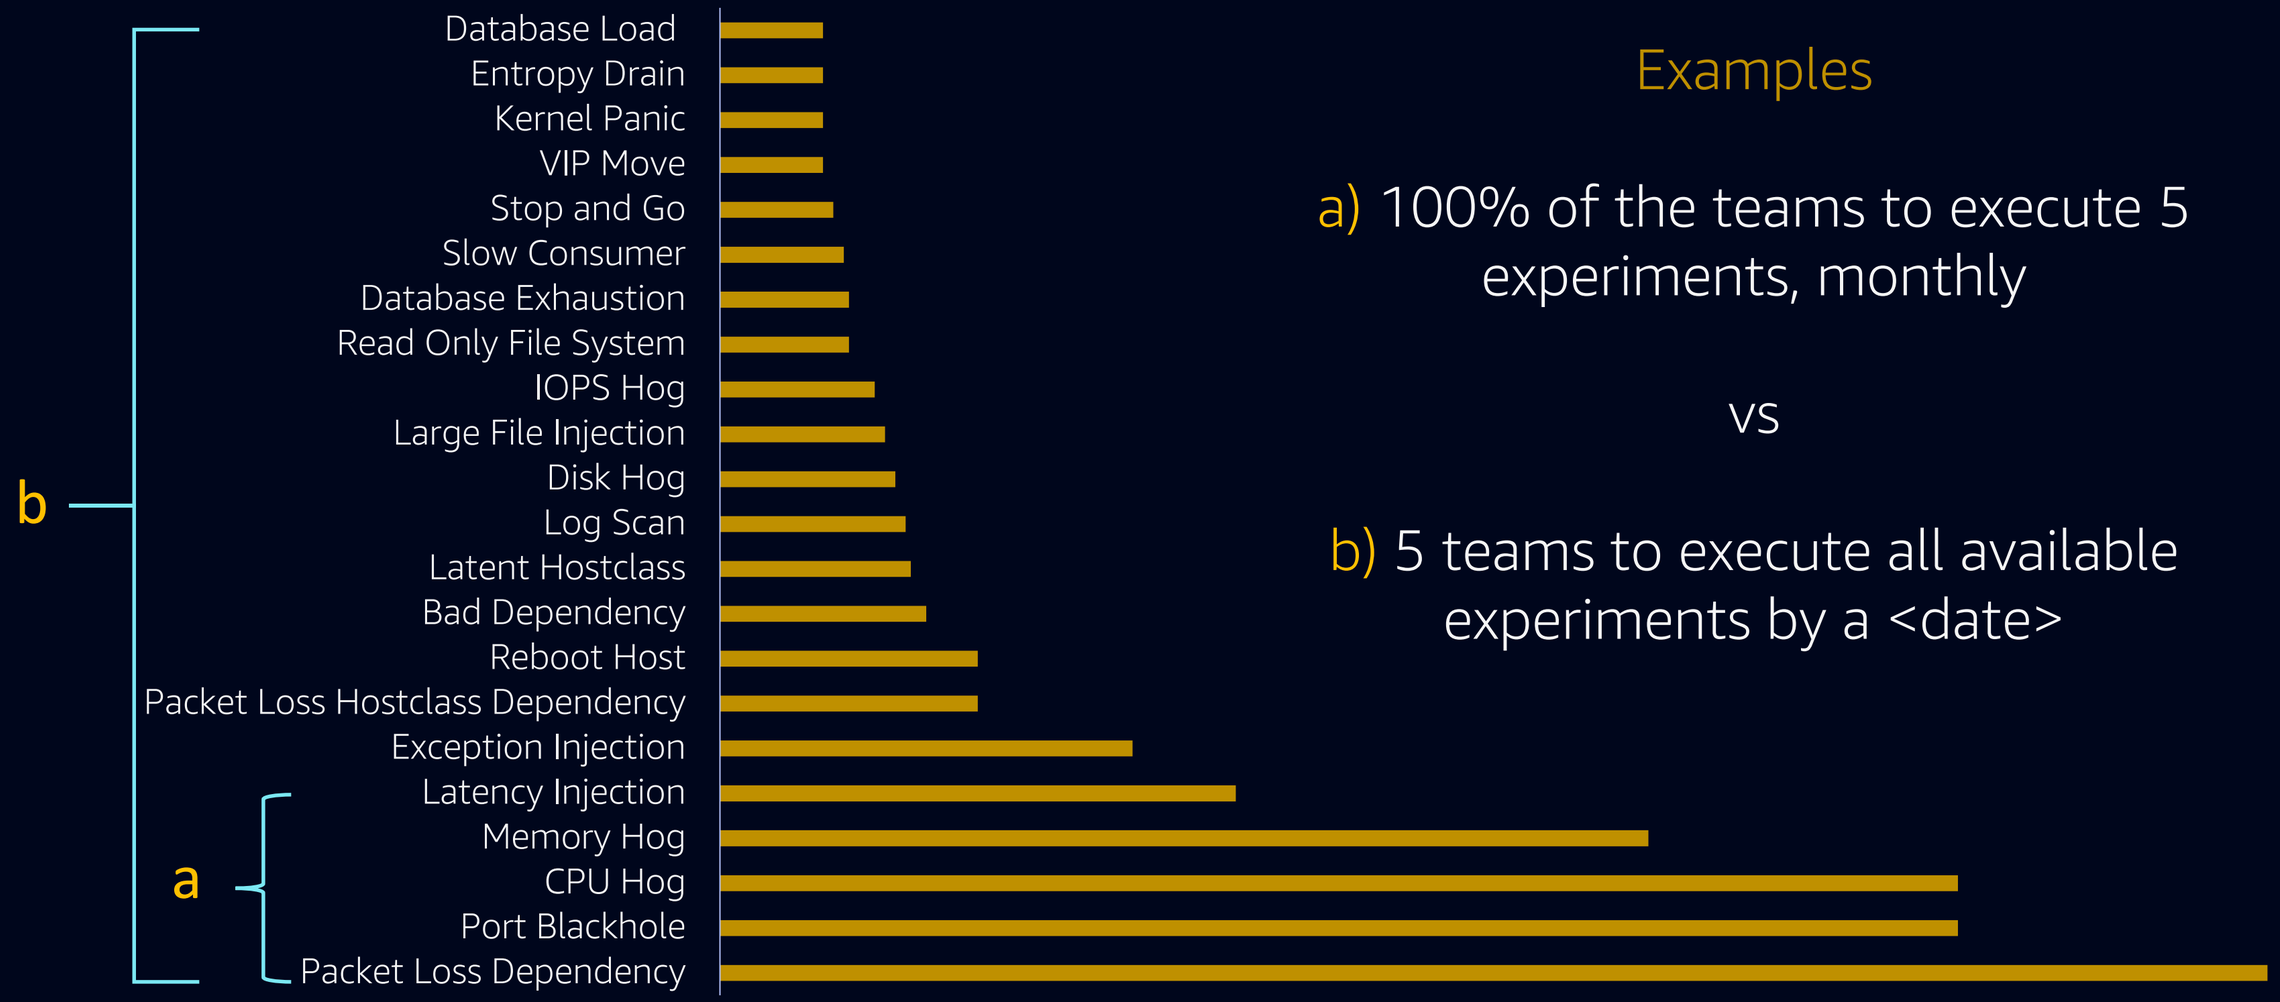

Examples

a) 100% of the teams to execute 5 experiments, monthly

vs

b) 5 teams to execute all available experiments by a <date>

prime video

FAILURE INJECTION TEST (FIT) PREVENTIVE EXAMPLE

AWS Open Source Blog, Aug 18, 2020

Building resilient services at Prime Video with chaos engineering

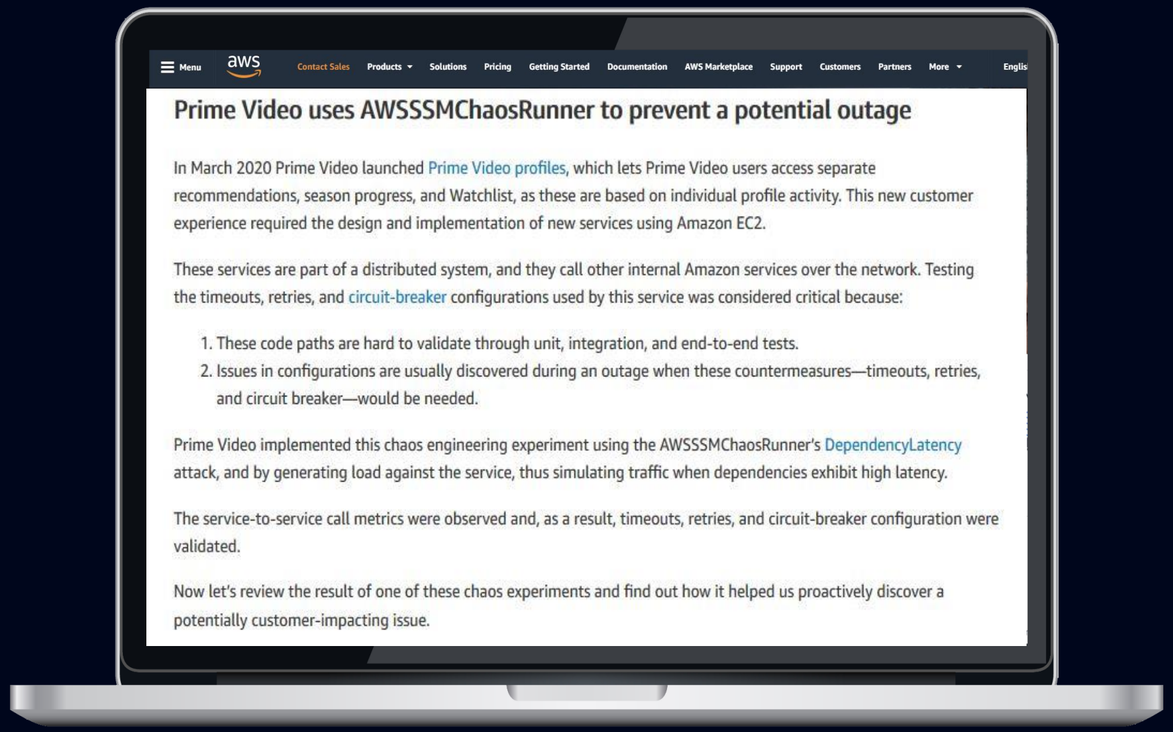

- Goal: confirm that APIs timeouts do not exceed agreed upon limit
- Hypothesis: the timeout is set at 40 milliseconds
- Experiment: generated 1000 requests against the service, injected 2s latency
- Learning: the timeout reached 6s
- Action: find/fix a bug in the configuration, rerun the experiment

prime video

FAILURE INJECTION TEST (FIT) REACTIVE EXAMPLE

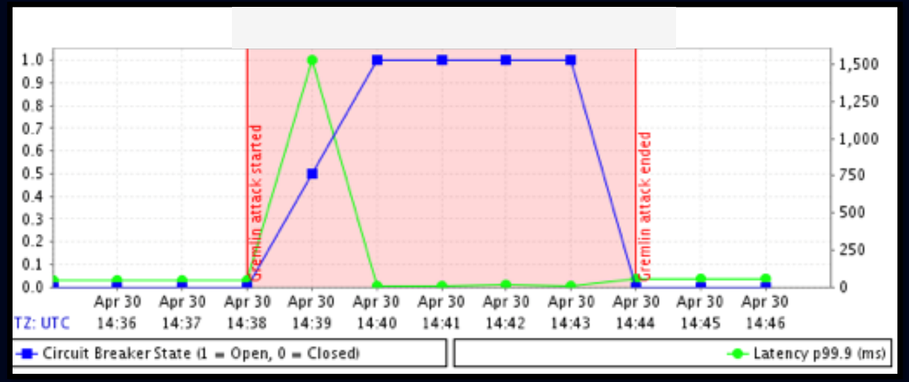

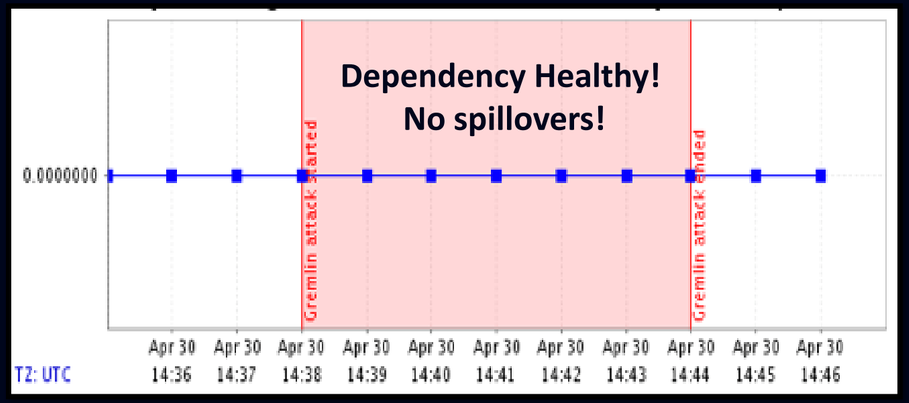

- Problem: Downstream dependency throttled traffic & caused availability drop
- Solution: Implemented fast-fail circuit-breaker
- Experiment: simulate dependency throttling, inject latency
- Learning: Confirmed circuit break performance, retries and latency impacts

prime video

EPL LAUNCH GOALS

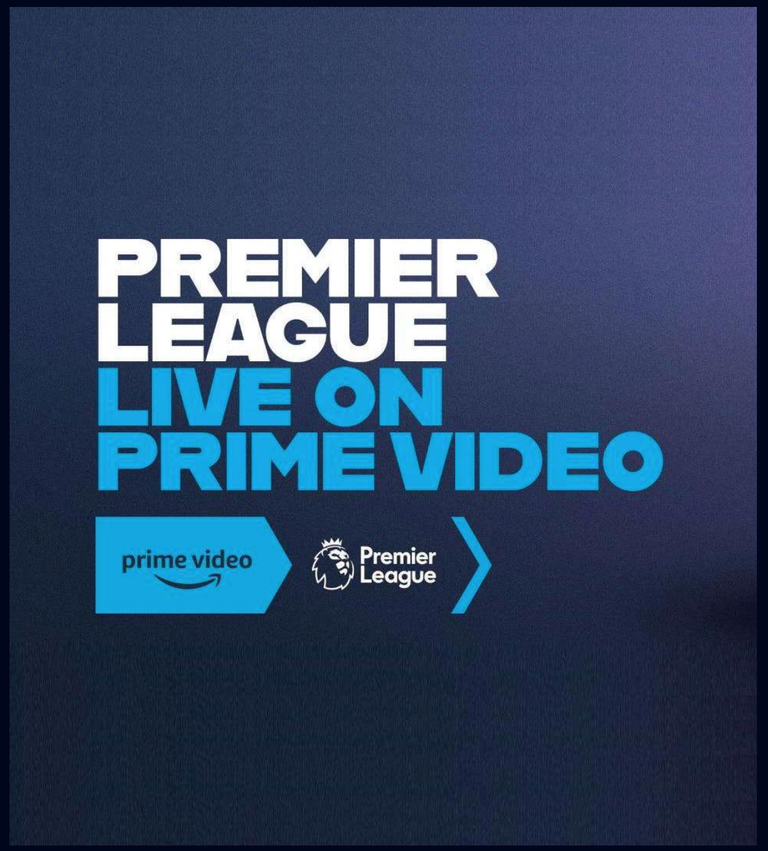

Generic

FEATURES GOAL

- Avoid a large feature wave, flatten
- Release to Beta early

Specific

SCALE AND RESILIENCE GOAL

- 100% Availability
- Gamedays and Simulations
- Chaos and Resilience as mechanisms

prime video

Specific

LAUNCH GOAL: TO BE ABLE TO DETECT AND RESOLVE ANY ISSUE REAL TIME.

Generic

SIMULATION GOAL: FOR ALL TEAMS AND PARTICIPANTS TO UNDERSTAND THEIR ROLES AND RESPONSIBILITIES IN VARIOUS SCENARIOS, AND ENSURE THAT ALL RESOURCES ARE IN PLACE TO FULFILL THESE ROLES AND RESPONSIBILITIES.

SIMULATIONS

- Attended by all teams supporting EPL, majority of which flew to London, over two days.
- Each day mimicked our live coverage, using real EPL matches broadcasted to beta customers over eight hours
- Simulations were injected to audit overall readiness, and pressure test triage strategies.

SIMULATIONS

Success Criteria

1. How quickly were we able to recognize the issue?
2. How quickly did we mitigate it?
3. Are we satisfied with our containment of the overall blast radius given the issue at hand (i.e. did it meet our success criteria)?

Execution

- Generated a list of possible issues across three categories:
    1) Recent high severity events,
    2) Scenarios that could be a risk, and
    3) Scale related issues we haven't seen before.
- We didn't make individual teams aware of the specific scenarios that we will test.
- Operated multiple scenarios at once.

SIMULATIONS

LIST OF SIMULATIONS

1. Host failures (~10 services)
2. Operator errors (use of “double agents”)
3. Critical tooling outages (Chime, video conferencing)
4. Misconfigured links
5. Live signal loss
6. Managing noise (red herring issues)
7. CDN/ISP failures

GOALS OWNERSHIP MODEL

SINGLE THREADED OWNER	YOU CONTROL INPUTS
JOINT OWNERSHIP	YOU DEPEND ON INPUTS FROM OTHERS
DISTRIBUTED OWNERSHIP	YOU SUCCEED WHEN MANY SUCCEED

prime video

ADOPTION VS EXPERIMENT GOAL

X% OF OUR SYSTEMS WILL ADOPT “Y” TECH BY “Z” DATE

WE WILL DEVELOP A PILOT BY “X” DATE AND CREATE A PLAN FOR ROLL OUT BY “Y” MILESTONE

prime video

ALL GOALS LEAD ... TO RESILIENCE!

THE CAPACITY TO RECOVER QUICKLY FROM DIFFICULTIES; TOUGHNESS.

RESILIENCE BUILDS ON SCALING & AVAILABILITY BRINGS IN INFRASTRUCTURE INTO FOLD

Dr Emmy E. Werner, 40 years of resilience study

1. Process of adapting well in the face of adversity, trauma, tragedy, threats or significant sources of stress.
2. Ordinary, not extraordinary.
3. Is not a trait that we either have or do not have.
4. It involves behaviors, thoughts and actions that can be learned and developed in anyone.

JOURNEY TAKEAWAYS

1. Start with a specific goal.
2. The first step is important. Next, keep a programmatic focus. Iterate on goals, use goals playbook and strategies. Decide on ownership. Repeat.
3. Keep learning about customers, create a feedback loop between your learnings and customers.

prime video

04

LOOKING AHEAD

WHAT DOES THE FUTURE HOLD...

"Always Day 1"

1. Continue to roll out new events and features that delight our customers, at scale
2. Distribute goals ownership
3. Strike the right balance

Chaos Engineering won’t make your system more robust,
People will.


OU THANK YOU T
K YOU THANK YO
THANK YOU THA
OU THANK YOU
K YOU THANK YO
THANK YOU THA
prime video

"AMAZON PRIME HAS STOLEN THE PREMIER LEAGUE'S VIEWER BASE THIS WEEK WITH SOME MAGNIFICENT FOOTBALL COVERAGE"

- DAILY STAR

"OPINION: AMAZON PRIME IS THE FUTURE OF FOOTBALL"

- THE MANCUNIAN