



# Breeding 10x Developers

With Developer Productivity Engineering

# Who is this guy?



Java™  
Champions



Warner River  
Bees



codemonkey.fm



Apache Directory



openstack



THE  
APACHE®  
SOFTWARE FOUNDATION

Maven™



Sonatype  
Nexus



**TABS**



**SPACES**





VS



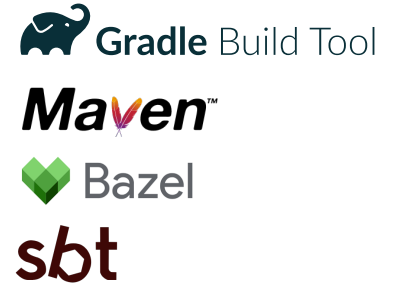
 **Gradle** Build Tool

VS

***Maven***<sup>TM</sup> 



**D**eveloper  
**P**roductivity  
**E**ngineering





10X ENGINEER



# Myth Origin (probably) The Coding War Games








# Peopleware

Productive Projects  
and Teams

THIRD EDITION



Tom DeMarco  
&  
Timothy Lister



The “best” programmers outperformed the worst by **roughly a 10:1 ratio**

There were some interesting “non-factors”:

**Language**

**Years of Experience**

**Number of Defects**

**Salary**

# What Mattered?

- ◆ Paired programmers from the same organization **performed at roughly the same level**
- ◆ The average **difference was only 21%** between paired participants
- ◆ They didn't work together on the task, but they **came from the same organization**
- ◆ **The best organizations performed 11.1x better than the worst**

"While this productivity differential among programmers is understandable, there is also a 10 to 1 difference in productivity among software organizations."

-Harlan D. Mills, Software Productivity



**The best performers are clustering in some organizations while the worst performers are clustering in others.**

**Some companies are doing a lot worse than others.**

**Something about their environment and corporate culture is failing to attract and keep good people or is making it impossible for even good people to work effectively.**

**Average performance of those in the top quarter was 2.6 times better than that of those in the bottom quarter.**

Table 8.3  
**Environments of the Best and Worst Performers  
in the Coding War Games**

Environmental Factor	Those Who Performed in 1st Quartile	Those Who Performed in 4th Quartile
1. How much dedicated work space do you have?	78 sq. ft.	46 sq. ft.
2. Is it acceptably quiet?	57% yes	29% yes
3. Is it acceptably private?	62% yes	19% yes
4. Can you silence your phone?	52% yes	10% yes
5. Can you divert your calls?	76% yes	19% yes
6. Do people often interrupt you needlessly?	38% yes	76% yes

**Though the phrase had not yet been coined, increased productivity came down to developer experience.**




10X ENGINEER



**10x Organizations are Manufactured, Not Born**

# Developer Joy



Developer Productivity Engineering

# ... But Most Organizations Aren't Aligned



In a study dated April 27, 2022, between Microsoft and the University of Victoria in British Columbia, Developers and Managers were surveyed on their interpretation of the SPACE framework



# When surveyed with the following questions, **Developers and Managers answered much differently**

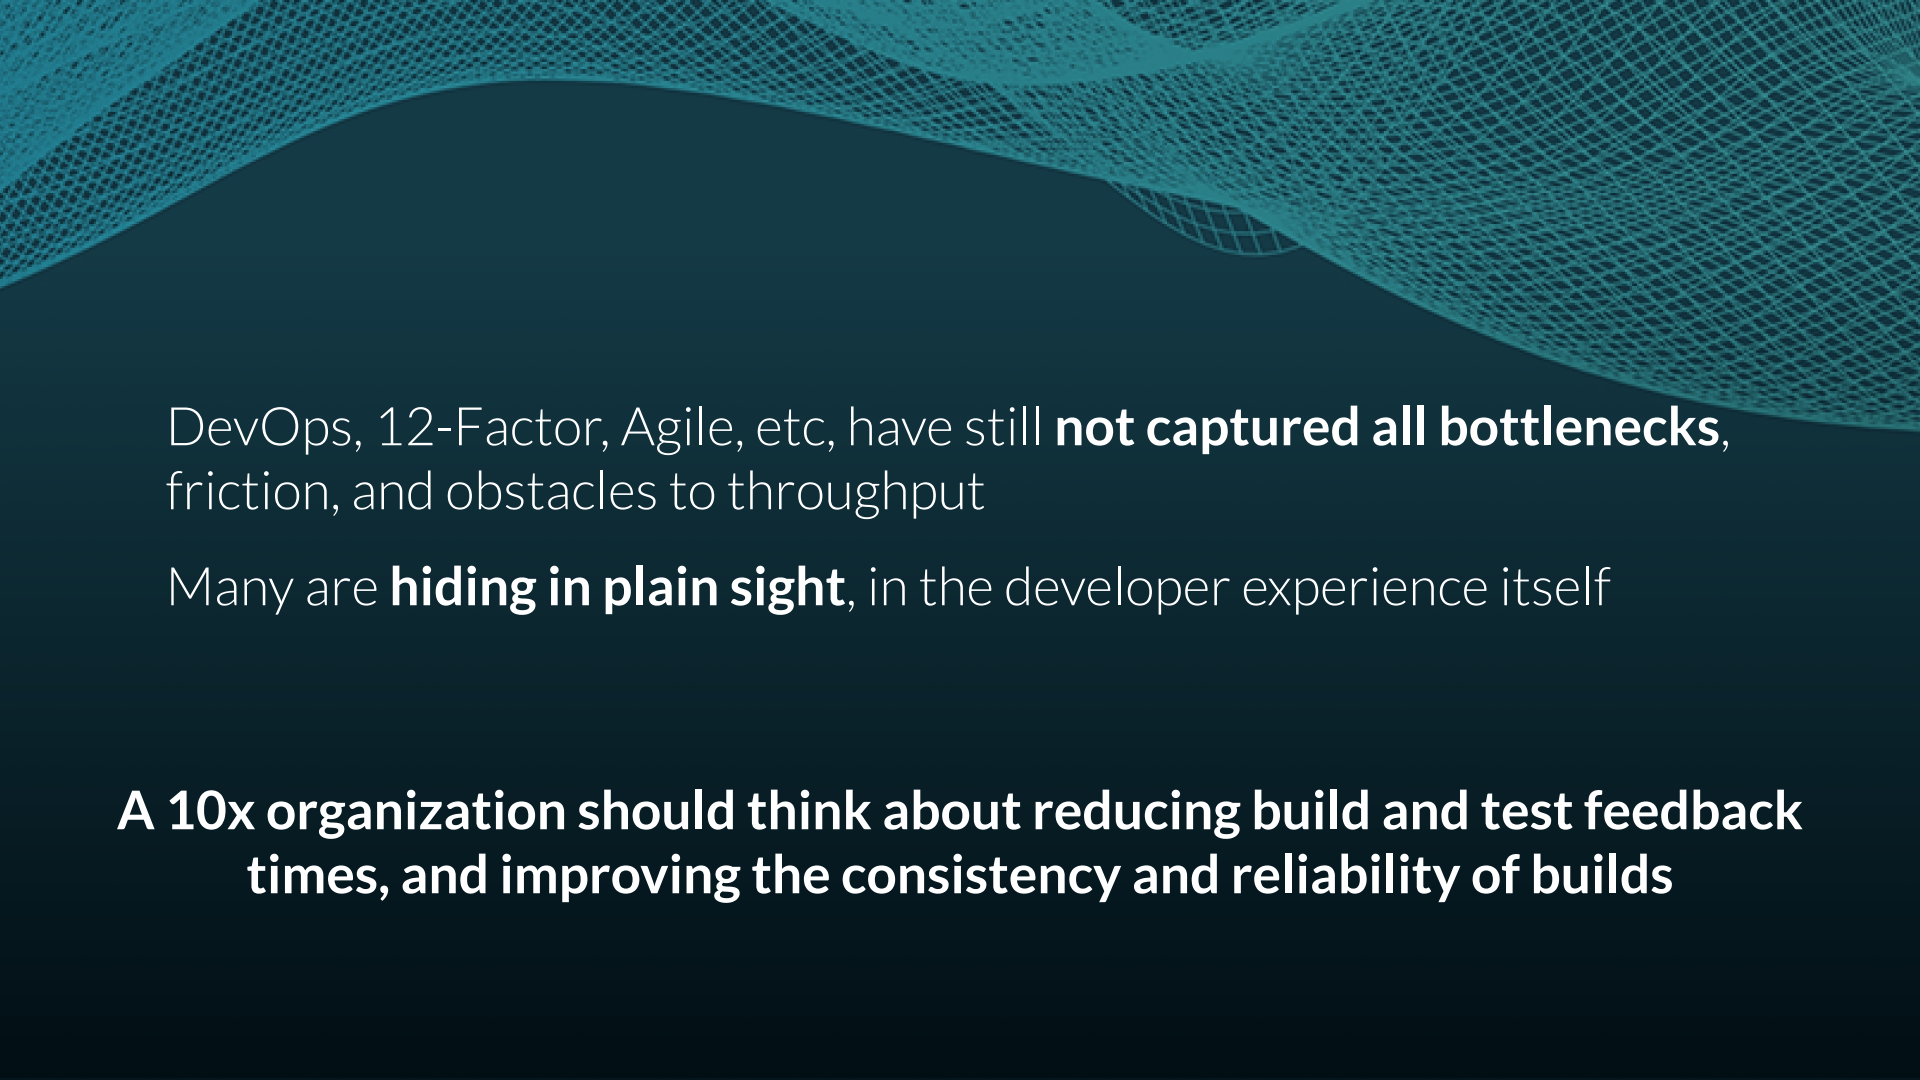
## Developers

When thinking about your work, how do you define productivity?

## Managers

When thinking about your team, how do you define productivity?

	ICs <b>define</b> own productivity	Managers <b>define</b> team's productivity	
S	8%	9%	
P	35%	67% (*)	
A	50%	21% (*)	
C	24%	33%	
E	38%	45%	

A decorative background consisting of a teal wireframe mesh that forms a wavy, undulating shape across the top of the slide. The mesh is composed of small squares and is semi-transparent, allowing the dark blue background to show through.

DevOps, 12-Factor, Agile, etc, have still **not captured all bottlenecks**, friction, and obstacles to throughput

Many are **hiding in plain sight**, in the developer experience itself

**A 10x organization should think about reducing build and test feedback times, and improving the consistency and reliability of builds**



The only initiatives that will positively impact performance are ones which **increase throughput while simultaneously decreasing cost.**

A decorative background featuring a teal wireframe mesh that flows across the top of the image, set against a dark teal gradient background.

# **It's Time for Developer Productivity Engineering**

1970s+

JIT  
Manufacturing

1980s+

Business  
Process  
Reengineering

1990s+

Change  
management

2000s+

Agile, Lean Six  
Sigma

2010s+

DevOps

2020+

DPE



Calendar Today < > August 2022

17

Create

March 2020 < >

S	M	T	W	T	F	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	1	2	3	4
5	6	7	8	9	10	11

Search for people

Other calendars + ^

WED 25

GMT-08

8 AM

9 AM

10 AM

11 AM

12 PM

1 PM

2 PM

3 PM

4 PM

5 PM

6 PM

7 PM

Code

Wait Time for Local Build

Debug Build Failure

Lunch

Code

Wait Time for Local Build

Sprint

Waiting time for CI Build

Investigate/Fix Flaky Tests

THE #1 PROGRAMMER EXCUSE  
FOR LEGITIMATELY SLACKING OFF:

"MY CODE'S COMPILING."

HEY! GET BACK  
TO WORK!

COMPILING!

OH. CARRY ON.





THE #1 PROGRAMMER EXCUSE  
FOR LEGITIMATELY SLACKING OFF:


"MY CODE'S ~~COMPILING.~~"

Testing

HEY! GET BACK  
TO WORK!


COMPILING!

OH. CARRY ON.



THE #1 PROGRAMMER EXCUSE  
FOR LEGITIMATELY SLACKING OFF:

**CHAT GPT IS DOWN**



# What Problems Does DPE Solve?



240 days per year



100s of developers





Productivity



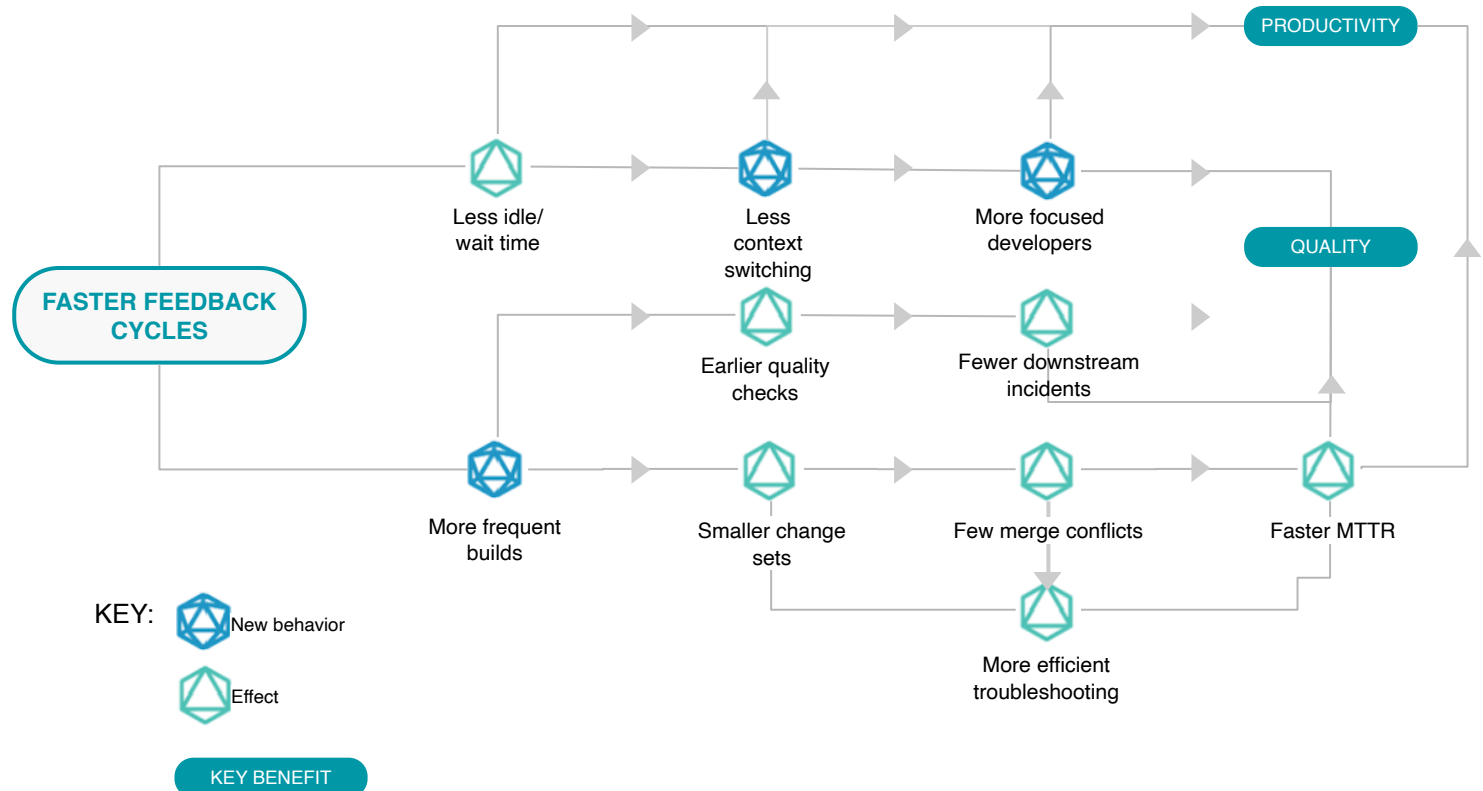
Cost  
10s of millions

 This takes too long!

 This takes too long to fix

 This should have been observable

# The anatomy and importance of fast feedback cycles



A decorative vertical pattern on the left side of the slide, consisting of light blue lines forming a circuit board layout with various geometric shapes like squares, circles, and hexagons.

***Build caching delivers fast build and test  
feedback cycles***



# Build Caching

- Introduced to the Java world by Gradle in 2017
- Maven has an open source build cache too
- **Used by leading technology companies** like Google and Facebook
- Can support both **user local and remote caching** for distributed teams

- Build caches are **complementary to dependency caches**, not mutually exclusive:
  - A dependency cache caches **fully compiled dependencies**
  - A build cache accelerates **building a single source repository**
  - A build cache caches build actions (e.g. Gradle tasks or Maven goals)



# What is a Build Cache?

Inputs



- Gradle Tasks
- Maven Goal Executions

Outputs



When the inputs have not changed, the **output can be reused** from a previous run.



# Cache Key/Value Calculation

The **cacheKey** for Gradle Tasks/Maven Goals is based on the Inputs:

```
cacheKey(javaCompile) = hash(sourceFiles,  
                             jdk version,  
                             classpath,  
                             compiler args)
```

The **cacheEntry** contains the output:

```
cacheEntry[cacheKey(javaCompile)] = fileTree(classFiles)
```

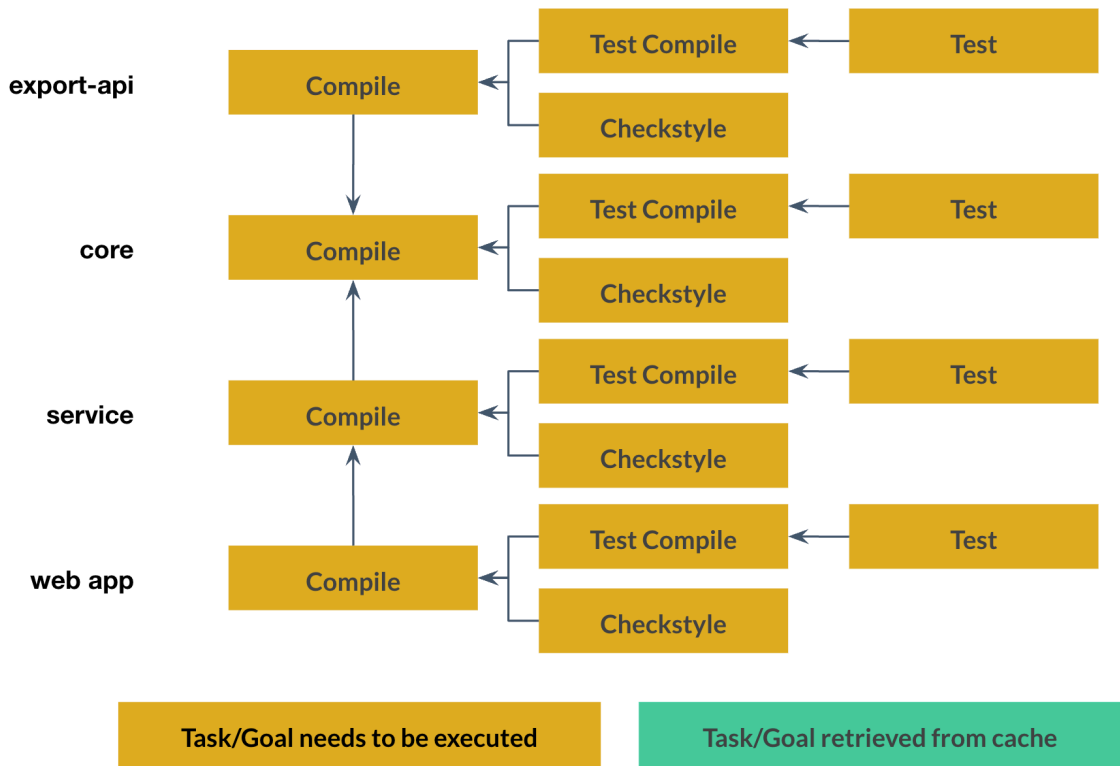
For more information, see:

[https://docs.gradle.org/current/userguide/build\\_cache.html](https://docs.gradle.org/current/userguide/build_cache.html)

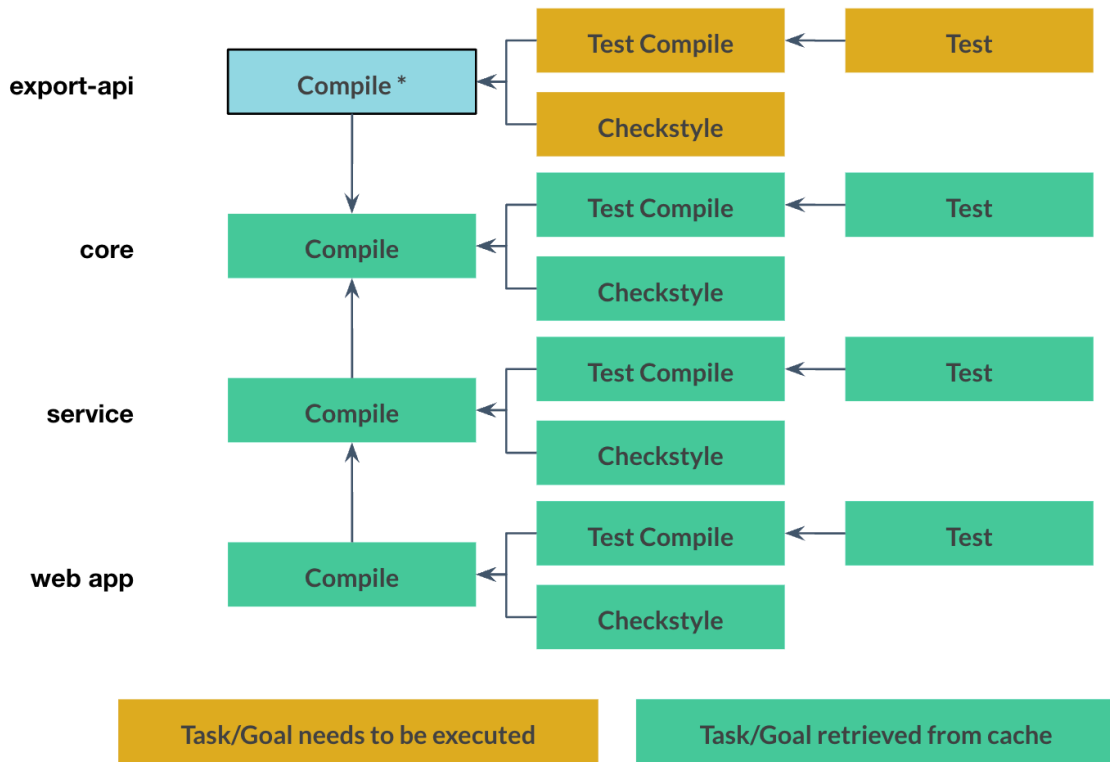




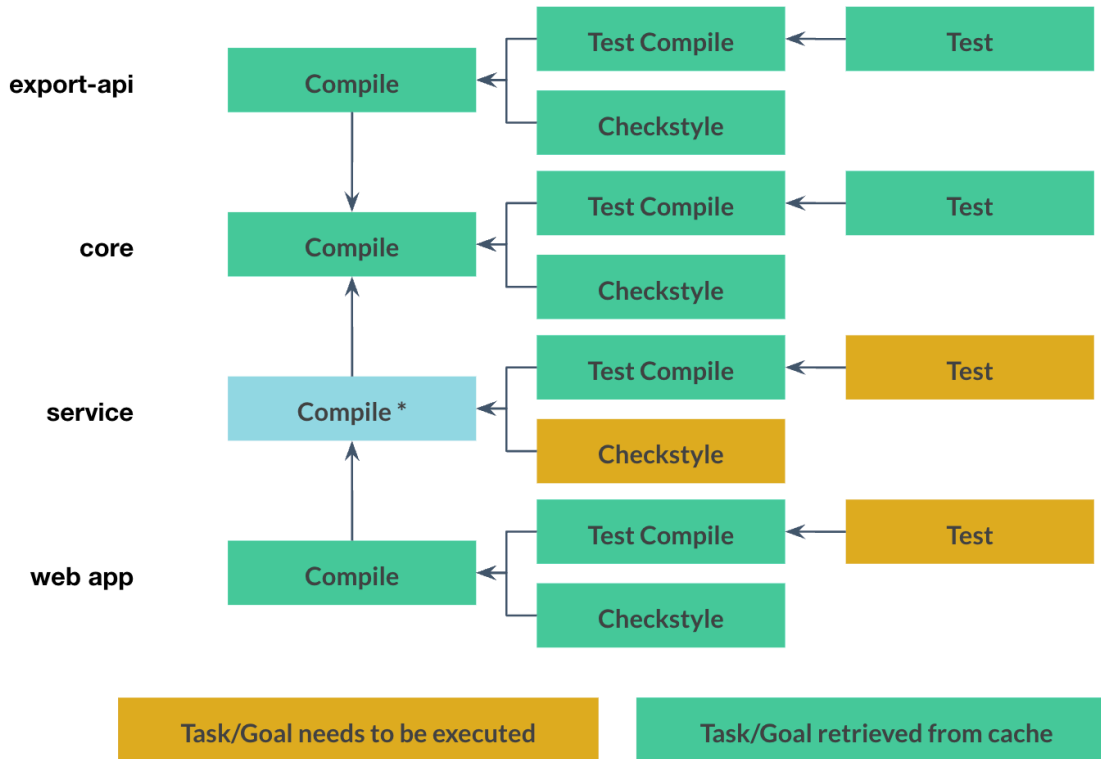
When not using the build cache, with Maven any change will require a **full build**. For Gradle this is the case when doing clean builds and switching between branches.




## Changing an **public** method in the **export-api** module



# Changing an implementation detail of a method in the service module



# Remote Build Cache



- ◆ Shared among different machines
- ◆ Speeds up development for the whole team
- ◆ Reuses build results among CI agents/jobs and individual developers





# Reproducible Builds

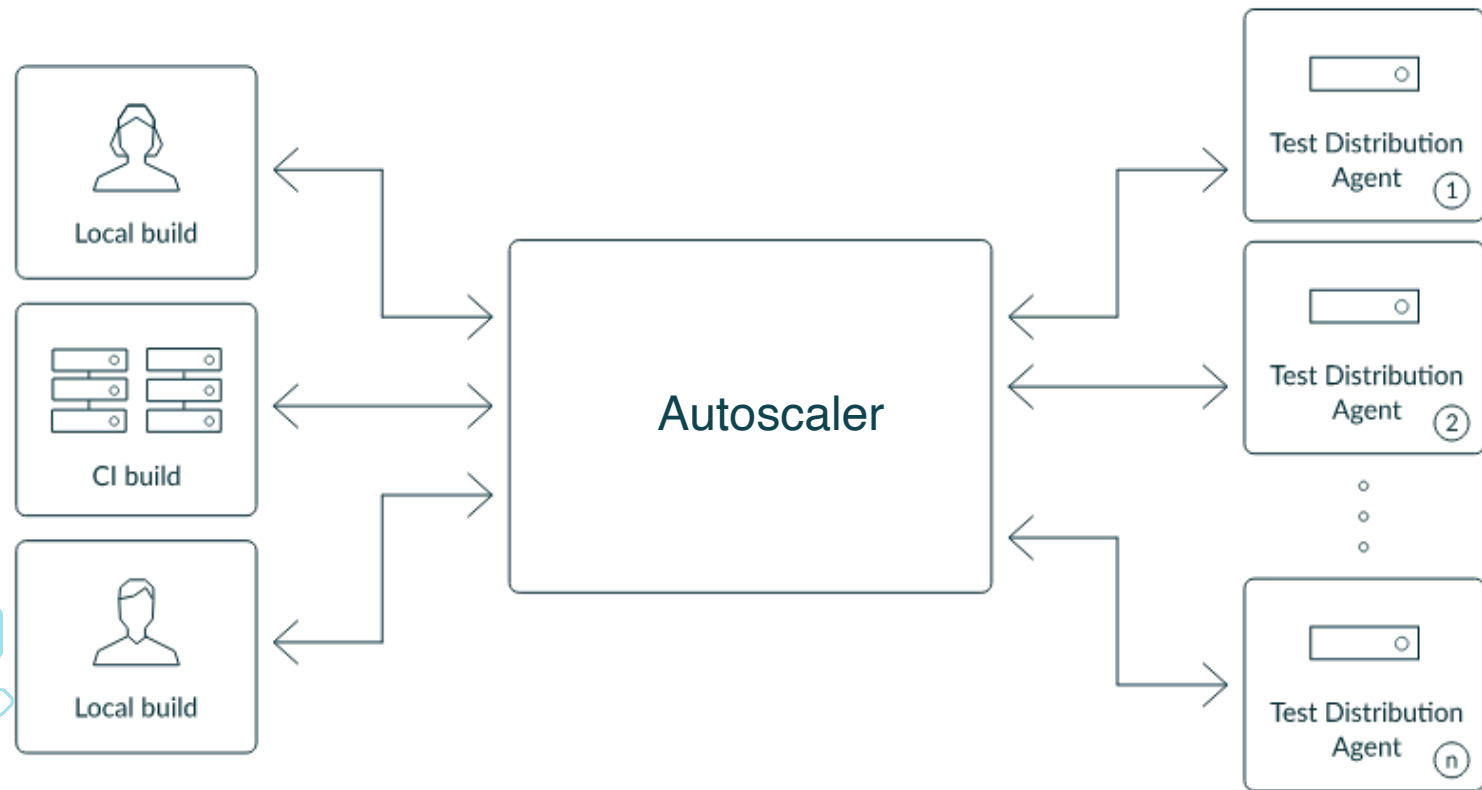
[reproducible-builds.org](https://reproducible-builds.org)



***Test distribution can make tests even faster***



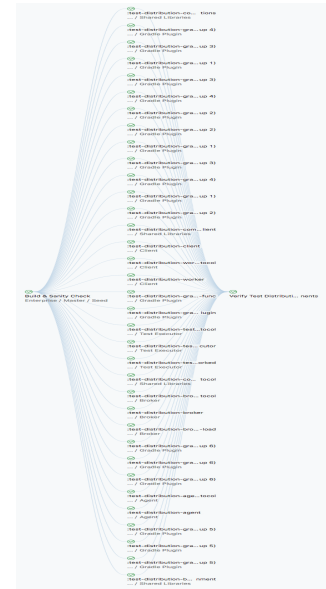
# How it works



# Existing solutions - CI fanout

Test execution is distributed by manually partitioning the test set and then running partitions in parallel on several CI nodes.

```
pipeline {
  stage('compile') { ... }
  parallelStage('test') {
    step {
      sh './gradlew :testGroup1'
    }
    step {
      sh './gradlew :testGroup2'
    }
    step {
      sh './gradlew :testGroup3'
    }
  }
}
```



See <https://builds.gradle.org/project/Gradle> for an example of this strategy







# Assessment of existing solutions

- ◆ **Build Caching** is great in many cases but doesn't help when test inputs have changed.
- ◆ **Single machine parallelism** is limited by that machine's resources.
- ◆ **CI fanout** does not help during local development, is inefficient (in particular on ephemeral CI agents or without build cache), requires manual setup and test partitioning, and result collection/aggregation



# NETFLIX



Netflix reduced a 62-minute test cycle time down to just under 5 minutes!





***Predictive Test Selection leads to greater efficiencies***



## Predictive Test Selection

International Conference on Software Engineering (ICSE)

### Abstract

Change-based testing is a key component of continuous integration at Facebook. However, a large number of tests coupled with a high rate of changes committed to our monolithic repository make it infeasible to run all potentially impacted tests on each change. We propose a new *predictive test selection strategy* which selects a subset of tests to exercise for each change submitted to the continuous integration system. The strategy is *learned* from a large dataset of historical test outcomes using basic machine learning techniques. Deployed in production, the strategy reduces the total infrastructure cost of testing code changes by a factor of two, while guaranteeing that over 95% of individual test failures and over 99.9% of faulty changes are still reported back to developers. The method we present here also accounts for the non-determinism of test outcomes, also known as test flakiness.

 Download Paper

 Copy PDF URL

By: Mateusz Machalica, Alex Samykin, Meredith Porth, Satish Chandra

November 23, 2020


Areas **AR/VR**

Tags **PROBABILITY**


Share   



# Conventional Test Selection Approach



# Predictive Test Selection Approach





***Build Scans speeds up troubleshooting***



Build Scan™ for 'spring-boot-b' x +

ge.spring.io/s/bqt22l4oo7lx4

Gradle Enterprise s... :spring-boot-project:sp... Jul 11 2023 16:09:15 EDT Build Scans Sign in

- Summary
- Console log
- Failure
- Deprecations
- Timeline
- Performance
- Tests
- Projects
- Dependencies
- Build dependencies
- Plugins
- Custom values
- Switches
- Infrastructure

See before and after

Compare Build Scan

JDK-17 dirty Linux Local main

Started today at 16:09:15 EDT, finished today at 16:09:23 EDT

Gradle 7.6.2, Gradle Enterprise plugin 3.12.5

Composite build (1 included build)

[Git commit build scans](#)

[Explore console log](#)

### 1 task failure

The `:spring-boot-project:spring-boot-tools:spring-boot-buildpack-platform:test` task failed. [View task in console log](#)

261 other builds with similar failures in last 7 days [View failure history](#)

There were failing tests. See the report at: `file:///home/sf frederick/Projects/spring/spring-boot/spring-b`

[Explore failure](#)

### 2 build deprecations

Build service 'testResultsOverview' is being used by task ':spring-boot-project:spring-boot-tools:spring-boot-buildpack-p Listener registration using `Gradle.addBuildListener()` has been deprecated.





# Without focus, problems can sneak back in...

- ◆ Infrastructure changes
  - Binary management
  - Caching
  - CI agents
- ◆ New annotation processors or versions of annotation processors
- ◆ Build logic configurations settings
  - Build tool version and plugins
  - Compiler and/or Memory settings
- ◆ Code refactoring
- ◆ New office locations
- ◆ Without observability, it is impossible to have a great and fast developer experience.



A decorative vertical pattern on the left side of the slide, consisting of light blue lines forming a complex circuit or maze-like structure. It includes various geometric shapes like squares, circles, and hexagons, some of which are filled with a light blue color.

"You can observe a lot by just watching"

- Yogi Berra, *Catcher and Philosopher*



# Performance Insights

**Are you tracking local  
build and test times?**



Is the cycle ~~fast~~ enough?


Is the cycle as **fast as it can possibly be?**



GRADLE ENTERPRISE CUSTOMER STATISTIC

## DPE Fosters Developer Joy

84% of surveyed users agree that DPE's impact on their toolchain makes their job more enjoyable.



---

Source: TechValidate survey of 51 users of Gradle Enterprise

✓ Validated

Published: Jul. 2, 2023 TVID: 930-05A-A5F



TechValidate  
by SurveyMonkey


# DPE Organizations Eliminate Avoidable Failures




> There were failing tests. See the report at:  
[file:///\\*/\\*/\\*/\\*/build/reports/tests/\\*/index.html](file:///*/*/*/*/build/reports/tests/*/index.html)

Builds with matching failures


1.23K (2% of 59K total builds)




Affected users



Affected hosts



Top Tags



Failed builds (50 most recent)

Start time	Project	Requested tasks/goals	User	Hostname
today at 1:00:47 AM	gradle	:core:embeddedIntegTest --tests org.grad	jvandort	Justins-MBP.home
	<div style="display: flex; gap: 5px;"> <span>CACHED</span> <span>LOCAL</span> <span>IDEA</span> <span>dirty</span> </div> <p>Execution failed for task ':core:embeddedIntegTest'.</p> <p>&gt; There were failing tests. See the report at: <a href="file:///Users/jvandort/work/gradle/subprojects/core/build/reports/tests/embeddedIntegTest/index.html">file:///Users/jvandort/work/gradle/subprojects/core/build/reports/tests/embeddedIntegTest/index.html</a></p>			

# DPE Organizations Track Failure Rates

# Dealing with Flaky Tests


The test is flaky. What do you do now?

- a. Try it again
  - b. Re-run it
  - c. Re-run it again
  - d. Ignore it and approve PR
  - e. All of the above
-




Builds that executed test class

160 builds



Mean execution time for test class

2 min 47 sec




Tests by flaky count

Name	Outcome trend	Failed	Flaky	Passed	Mean execution time
performs static validation of plugins used by the Gradle build		5 (3%)	144 (90%)	11 (7%)	2 min 47 sec

# DPE Organizations Analyze Flaky Tests

# All Of This Will Improve CI

## Distributed Agent Availability - Main Branch





**DPE Will Become Standard Practice  
Because the World Should Foster Developer Joy**

# Questions?

 **BrianDemers**

 **bdemers**



Learn more & get free swag



The background features a series of undulating, wave-like shapes rendered in a teal wireframe mesh. These shapes create a sense of depth and movement against a dark teal gradient background. The mesh lines are thin and light, giving the overall appearance a digital or architectural feel.

**THANKS!!**