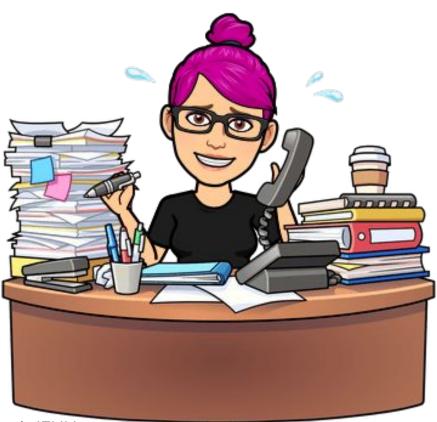
The Path to Build Happiness

Jenn Strater, Developer Advocate at Gradle





Jenn Strater



- Developer Advocate at Gradle
- Previously a software engineer at various companies
- Berlin <-> Minneapolis, MN, USA



Gradle Inc



Open Source
Build Automation



Commercial Product (on-premises)

Development Productivity





Outline

What is Developer Productivity Engineering?

Current state of Developer Productivity

Achieving Build Happiness

Measure

Analyze

Optimize

Iterate



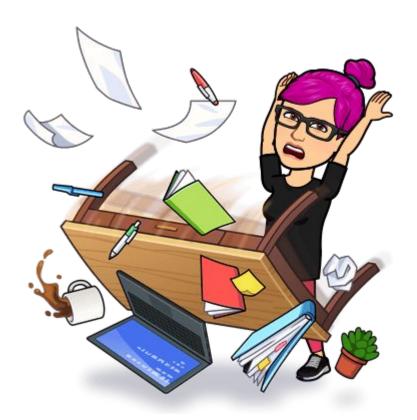








The Paradox of Success



Without intervention, as the

- Lines of Code
- No. of developers
- No. of repositories
- No. of dependencies
- No. of support tech stacks

increase, so does the frustration!







How do we get back to bliss?









Quality of Creative Flow + Collaborative Effectiveness

Team Productivity







Developer Productivity Engineering is a culture where the whole organization commits to an effort to maximize developer productivity.





Developer Productivity Engineering KPIs

- Degree of automation
- Speed of feedback cycles
- Correctness of the feedback

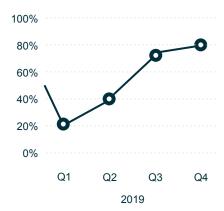


Developer Productivity Engineering

Priorities and success criteria primarily based on data that comes from a fully instrumented toolchain.











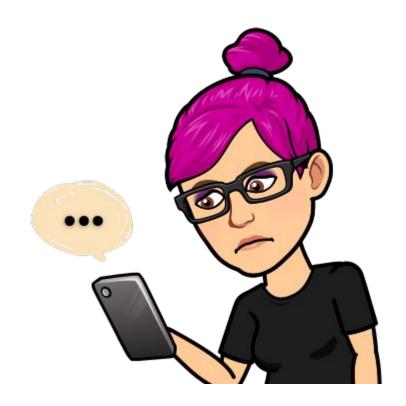


Current State of Developer Productivity Engineering





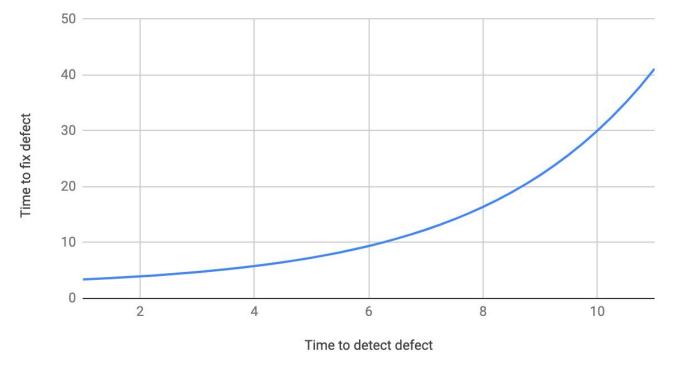
The longer the build the harder to debug







Fix time grows exponentially over detection time

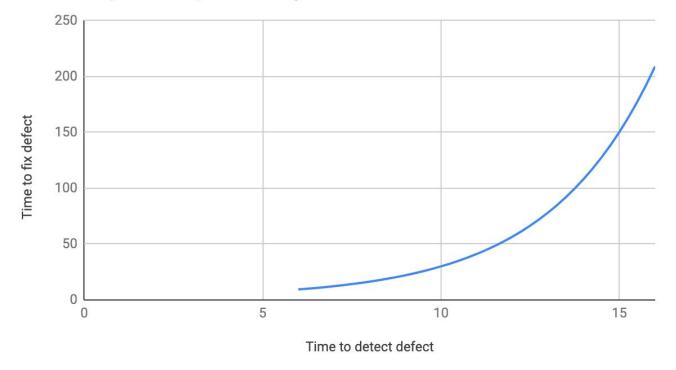


• In the case of a failure, the time fixing the failure is growing exponentially with the time it takes to detect it.





Fix time grows exponentially over detection time



- Because of growing build times, test and builds are pushed to a later point in the life cycle.
- The exponential costs for debugging is increased by that.
- It also increases the change set size as it becomes inconvenient to get feedback.



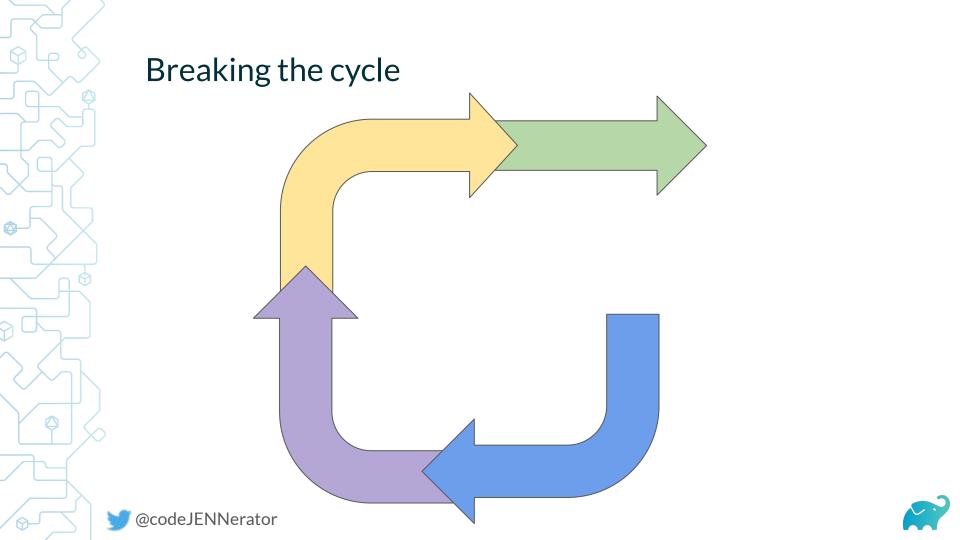




Low Developer Productivity is blocking innovation







Step O. Focus on Developer Productivity





Prioritization without data

- Later
- Tomorrow
- Today



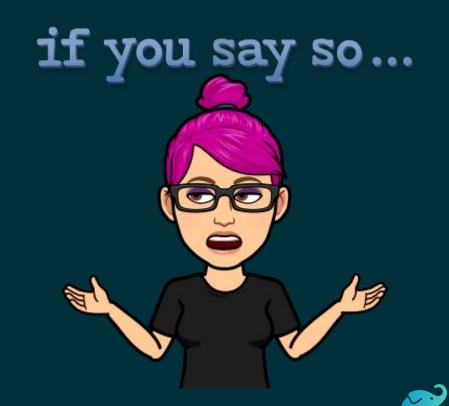
- Troubleshooting sessions beginwith a game of 20 questions
- Person asking for help often doesn't know what context is important
- Helpers can burn out on helping
- Root cause analysis often impossible without the helper reproducing the problem
- Impossible to determine the impact of the issue





Results in anecdote-driven prioritization

- Non-verification failure masks as verification failure (flakey test)
- Verification failure masks as non-verification failure (snapshot dependency issue)
- Non-verification failure might be caused by bug in a plugin or user mis-configuration
- Many issues are flakey and hard to reproduce









Reliability







Flaky builds and tests are **maddening**





Frustrated Developers Leave

Most developers want to work in an environment that enables them to work at their full potential.

Organizations that can not provide such an environment will lose talent.

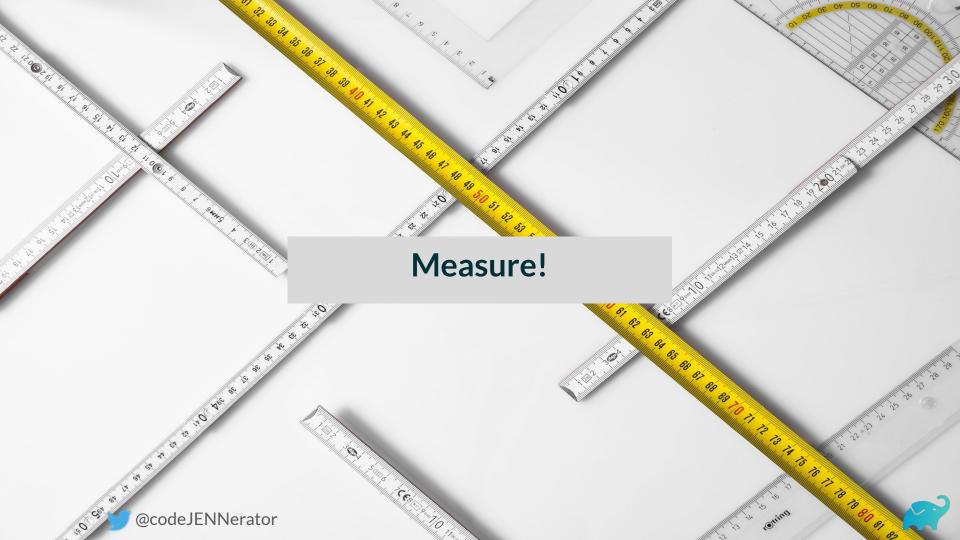












Faster builds improve the creative flow

	Team 1	Team 2
No. of Devs	11	6
Build Time	4 mins	1 mins
No. of local builds	850	1010

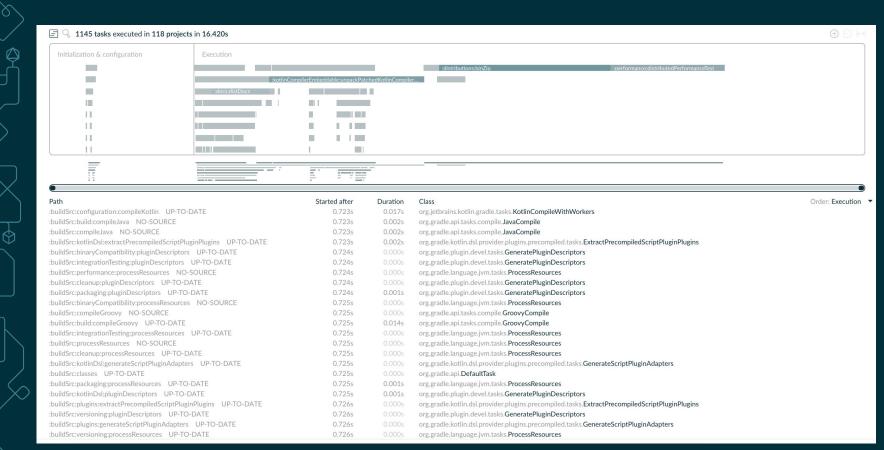
- The faster the feedback is, the more often devs ask for feedback
- The more often they ask for feedback, the more fine grained they can work.





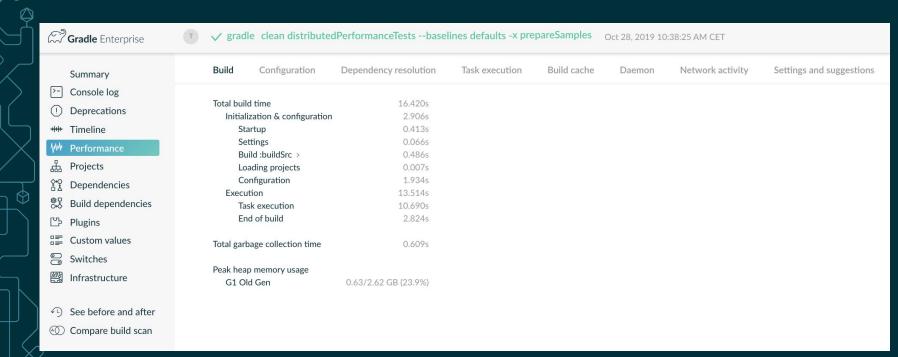












https://e.grdev.net/s/vpui4db7vx6xo/performance



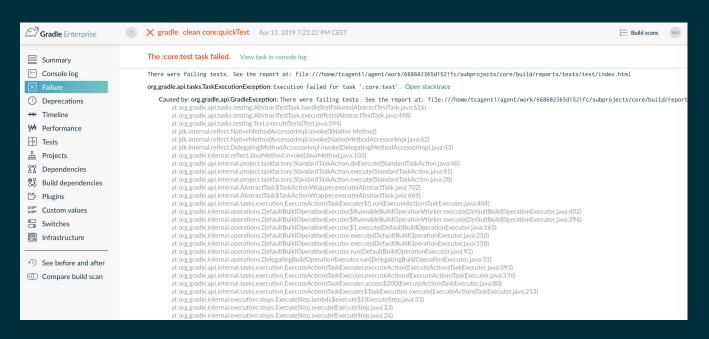


Gradle Enterprise	✓ gradle clean distribute	edPerformanceTe	stsbaselir	es defaults -x pre	pareSamples	Oct 28, 2019 10	:38:25 AM CET	
Summary	Build Configuration	Dependency res	solution	Task execution	Build cache	Daemon	Network activity	Settings and suggestions
>- Console log								
	Time spent executing tasks		10.923s					
Deprecations								
+++ Timeline	All tasks	1145	26.133s					
₩ Performance	Tasks avoided	447 (62.3%)	10.430s					
	From cache	262 (36.5%)	9.676s					
	Up-to-date Tasks executed	185 (25.8%) 270 (37.7%)	0.754s 15.526s					
Dependencies	Cacheable	0 (00.0%)	0.000s					
Build dependencies	Not cacheable	270 (37.7%)	15.526s					
	Lifecycle	176	0.012s					
Plugins	No source	249	0.160s					
Custom values	Skipped	3	0.005s					
Switches								
Infrastructure	Snapshotting task inputs	966	7.726s					
	₹ 1145 tasks in 118 projects							
	:distributions:binZip				3.454s			
See before and after	:kotlinCompilerEmbeddable:unpackPatchedKotlinCompilerEmbeddable			ddable	3.118s			
Compare build scan	:performance:distributedPerformanceTest FROM-CACHE			2.2925				
	:docs:distDocs FROM-CACHE				1.153s			
	:kotlinCompilerEmbeddable:cle	ean			0.739s			
	:toolingApi:toolingApiShadedJ	ar			0.706s			
	:javascript:processResources				0.686s			
	:kotlinCompilerEmbeddable:pa	tchKotlinCompilerEn	mbeddable FF	OM-CACHE	0.664s			
	:processServices:jar				0.589s			
	:performance:intTestImage				0.533s			
	:resourcesGcs:classpathManife	est FROM-CACHE			0.482s			
	:dependencyManagement:jar				0.350s			
	:internalIntegTesting:prepareVe	ersionsInfo FROM-C	ACHE		0.327s			
	:core:jar				0.320s			
	:internalIntegTesting:jar	THE			0.257s			
	:core:compileJava FROM-CAC :kotlinDsl:jar	ПС			0.201s 0.199s			
	:dependencyManagement:com	nile lava EROM-CA	CHE		0.162s			
	:internalTesting:jar	ipiicava i Kolvi CA	CITE		0.134s			
	:modelCore:compileJava FRO	M-CACHE			0.125s			
	:ide:jar	A 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1			0.116s			
	:modelCore:iar				0.116s			



Capture data from every build run on local AND CI

Comprehensive data allows for root cause analysis without reproducing locally

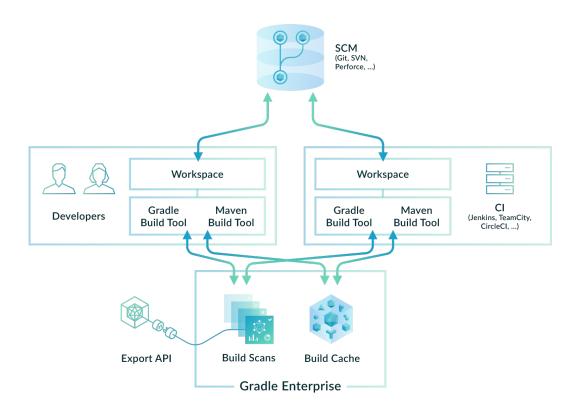








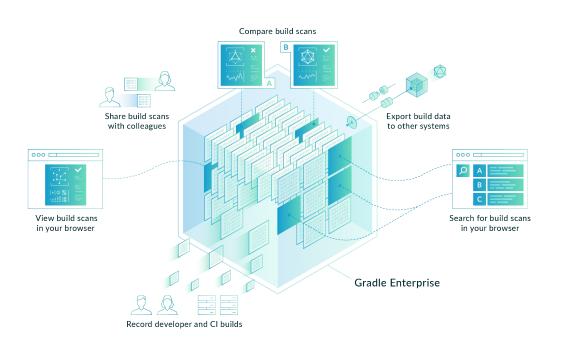
What is Gradle Enterprise?







Gradle Enterprise is a data platform



- Collect team data all the data about every build across the team creates unique dataset and insights.
- Build Performance Management only with representative, actionable data will builds get and stay fast and reliable.
- Debugging acceleration only with comprehensive, deep data is it possible to quickly discover the root cause for build failures.





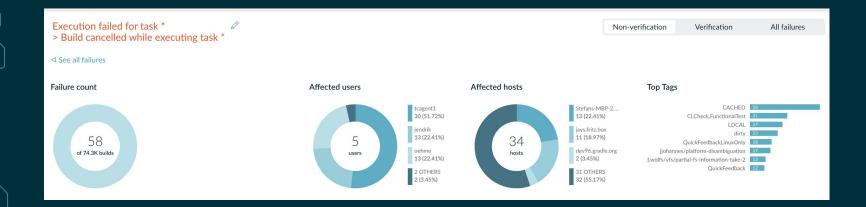
Failure Types

- Verification Failures
 - Syntax error detected by compilation
 - Code style violation detected by checkstyle
 - Misbehavior in the code detected by a JUnit test
- Non-Verification Failures
 - Flakey Test
 - Binary repository down
 - Out of memory exception while running the build
- Slow Builds





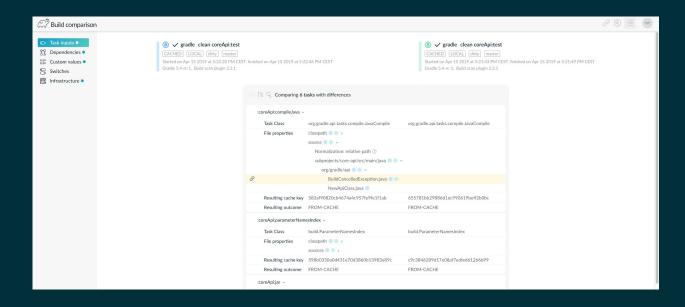
Impact analysis



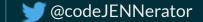




Root Cause Analysis



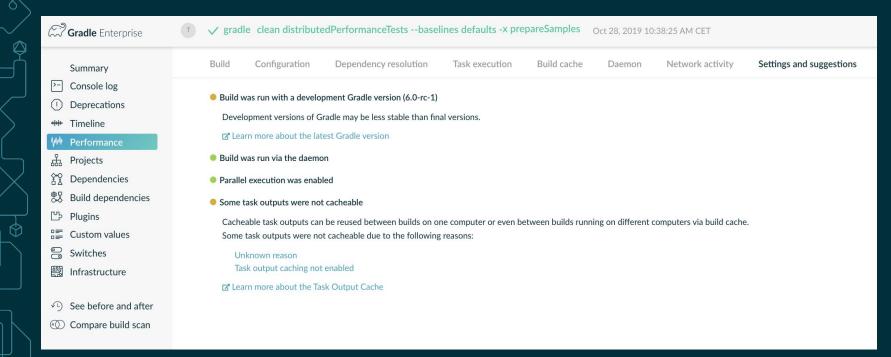






Optimize!





ttps://e.grdev.net/s/vpui4db7vx6xo/performance/suggestions



Build Caching

Inputs

- Gradle Tasks
- Maven Goal Executions

Outputs

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







Cacheable Task/Goal Executions

Gradle Compile/Maven Compile

- Source Files
- Compile Classpath
- Java version
- Compiler configuration
- etc...

Gradle Test/Maven Surefire

- Test Source Files
- Runtime Classpath
- Java version
- System properties
- etc...

Caching is a generic feature and applies to all tasks/goals. For IO-bound tasks/goals caching has no benefits (e.g. clean, copy).





Caching is effective for multi-module builds

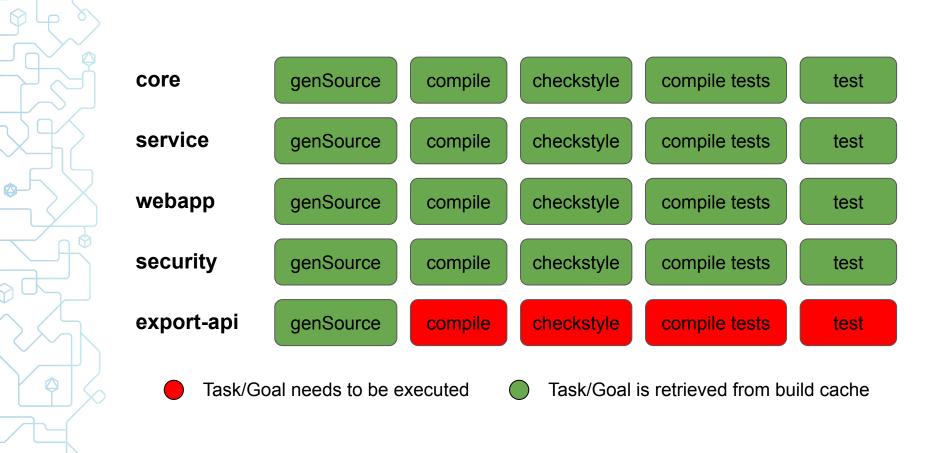
Builds with a single module will only moderately benefit from the cache



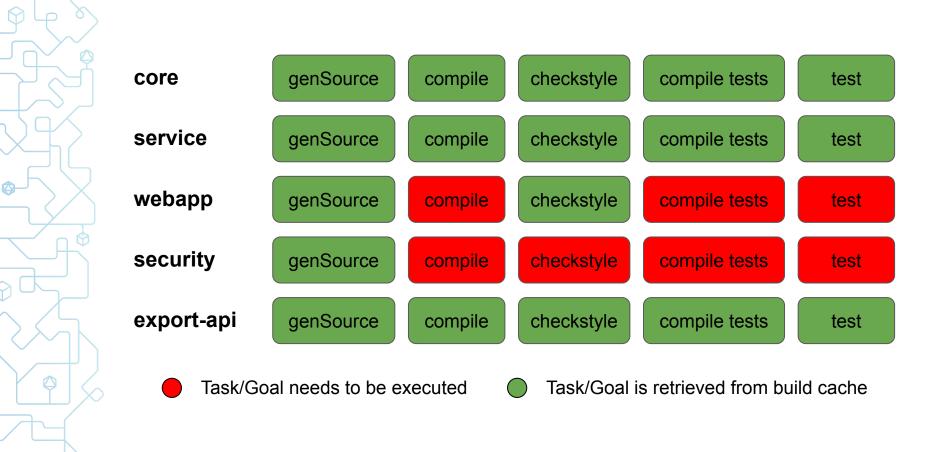












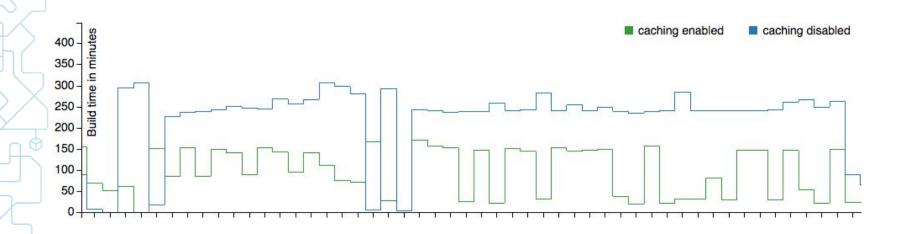


Cache effectiveness

- Even with only a few modules a cache significantly reduces build and test times
- For larger multi-module builds often 50% of modules are leaf modules
 - Build times is reduced by approximately 1/n with n being the number of modules
- Checking the inputs and downloading & unpacking items of the cache introduces overhead.
- Overhead is often very small compared to benefits
- Overhead should be measured and monitored too



Gradle CI Builds

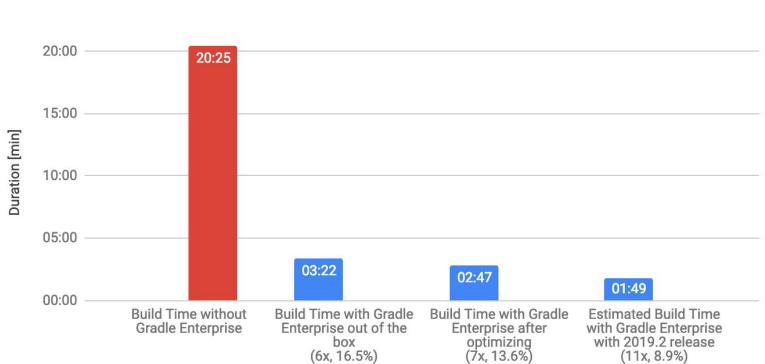


Build times are >80% faster for Gradle Core
Dramatically better caching results due to build scans.





Spring Boot build time for compile and unit tests (fully cached) https://spring.io/projects/spring-boot 25:00 20:00 20:25 15:00







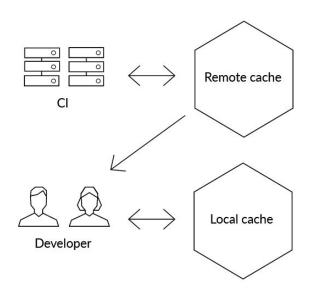
Local Build Cache

- Uses a cache directory on your local machine
- Speeds up development for single developer or build agent
- Reuses build results when switching branches locally





Remote Build Cache



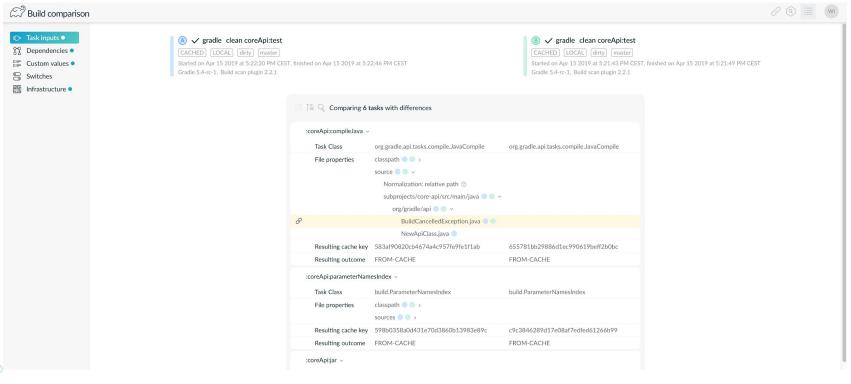
- Shared among different machines
- Speeds up development for the whole team
- Reuses build results among CI agents/jobs and individual developers





Build comparison ₹ Dependencies • Custom values Gradle 5.4-rc-1. Build scan plugin 2.2.1 Switches Infrastructure •





LinkedIn: Productivity at scale: How we improved build time with Gradle build cache SoundCloud: Solving Remote Build Cache Misses by Annoying Your Colleagues











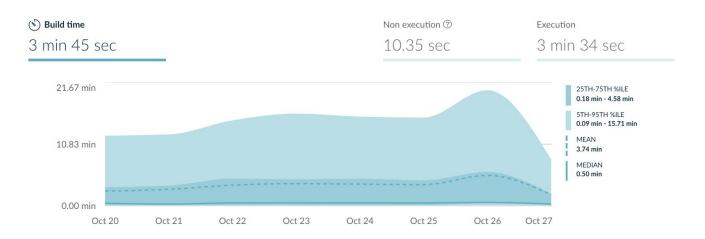


Iterate!





Performance regressions are easily introduced



- Infrastructure changes
- New annotation processors or versions of annotation processors
- Build logic configurations settings
- Code refactoring





What happens today with most regressions

- Unnoticed
- Noticed but unreported
- Reported but not addressed
 - Root cause is hard to detect (especially with flakey issues)
 - Overall impact and priority can not be determined
- Escalated after they have caused a lot of pain
 - Problem gets fixed after it has wasted a lot of time and caused a lot of frustration amongst developers.
- Result: The average build time is much higher than necessary and continuously increasing.



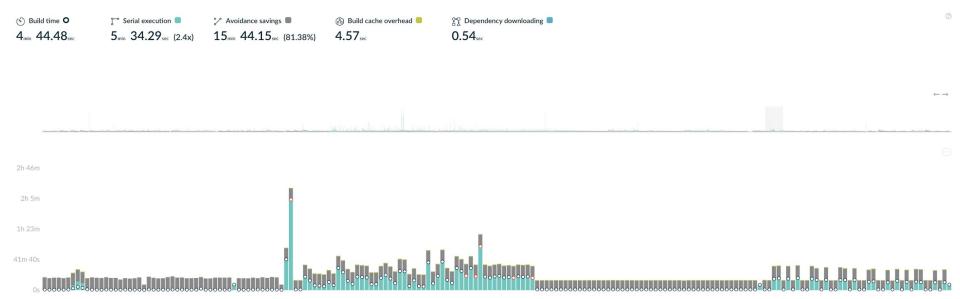




Performance Analytics







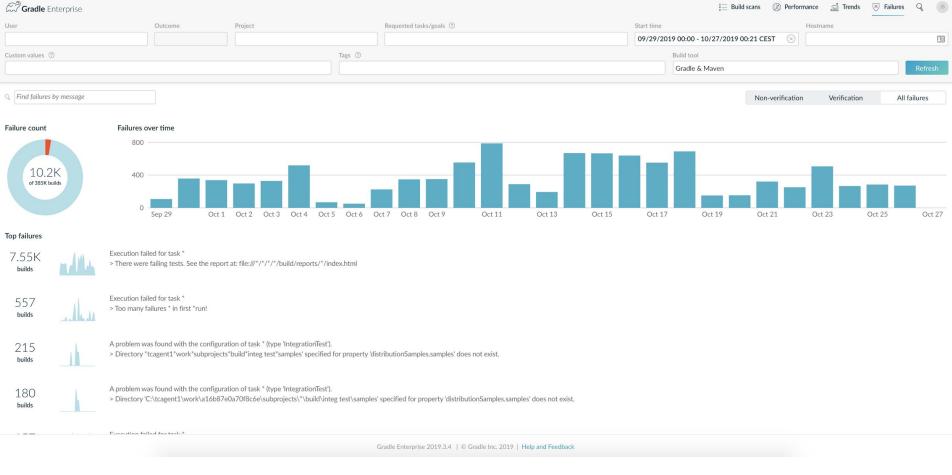




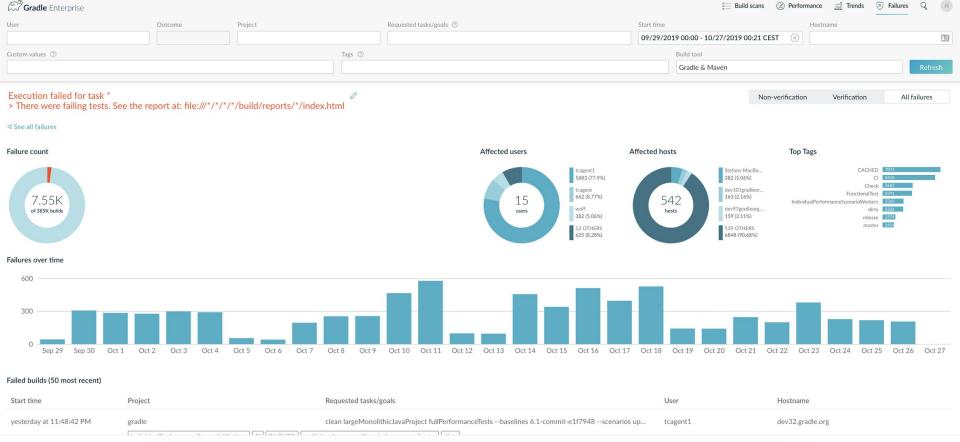
Failures Dashboard











Gradle Enterprise 2019.3.4 | © Gradle Inc. 2019 | Help and Feedback







Flaky tests are a problem for everyone

We're currently working on a solution

Plan to release this in Q4/2019

Interested in your ideas & thoughts



Conclusion

- Don't suffer in silence.
- The Path to Build Happiness is through Developer Productivity.
- Measure, Analyze, Optimize, and Iterate to achieve and maintain build happiness.





Resources



- Early Access Book: https://gradle.com/developer-productivity-engineering
- Try out build scans for Maven and Gradle for free: https://scans.gradle.com
- Gradle Enterprise docs and tutorials: https://docs.gradle.com







Thank you!



@codeJENNerator

