# Automated Detection and Reporting of Build Cache Misses

## DPE Summit 2023

Etienne Studer, SVP of Engineering, Develocity

# Increased DevProd through Build Caching

**Build caching is a powerful technology to accelerate build performance on CI and locally.**

# How Build Caching works

**Build Caching avoids doing the same work already done before by reusing the work unit output stored in a local or remote build cache.**

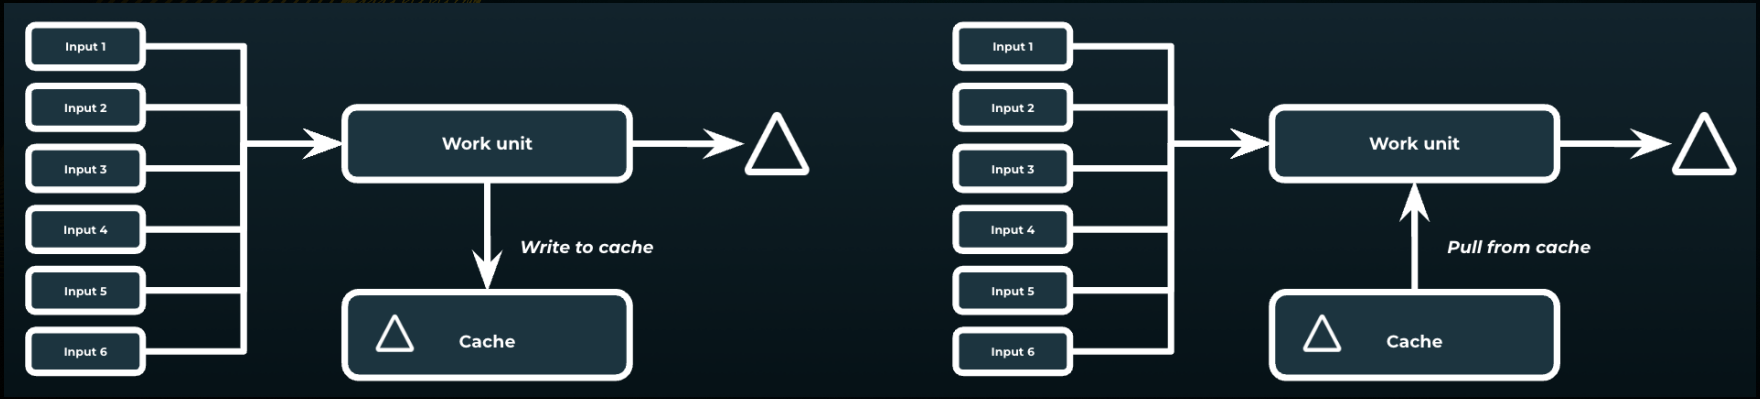

# Volatility in the inputs of work units

**Volatile inputs of work units cause unnecessary execution of work units.**

# Verification if a build is fully cacheable

When a build is executed twice under the same conditions, all of the second build's cache-compatible work units should take their output from the build cache.

# Automation of the verification process

**Use the Build Validation Scripts to make the process reliable, repeatable, automated, and productive.**

https://github.com/gradle/gradle-enterprise-build-validation-scripts

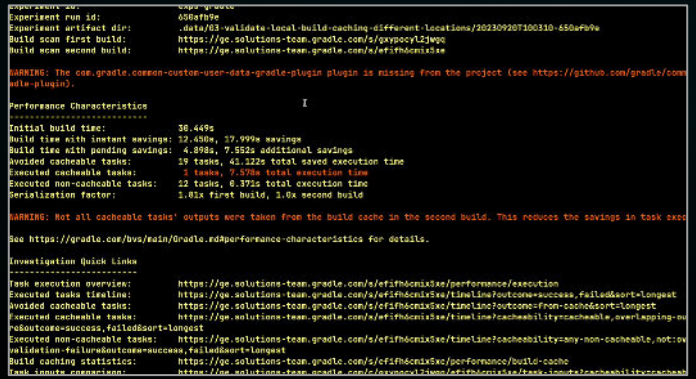

# Investigation of build cache misses

**Use Develocity Build Scan Comparison to investigate changes in task inputs that caused a build cache miss.**

https://ge.solutions-team.gradle.com/c/gxypocyl2jwgq/efifh6cmix5xe/task-inputs

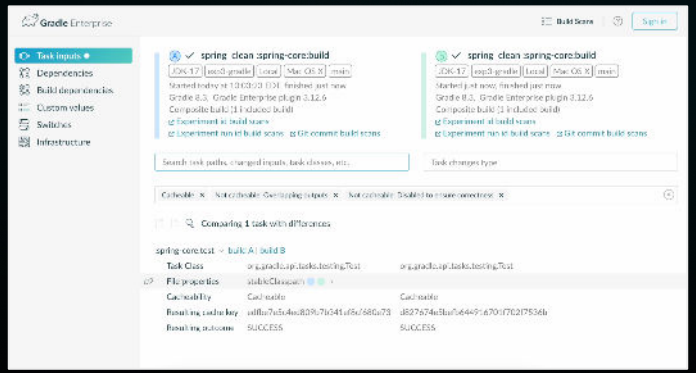

# Demo: BVS-driven build cacheability verification

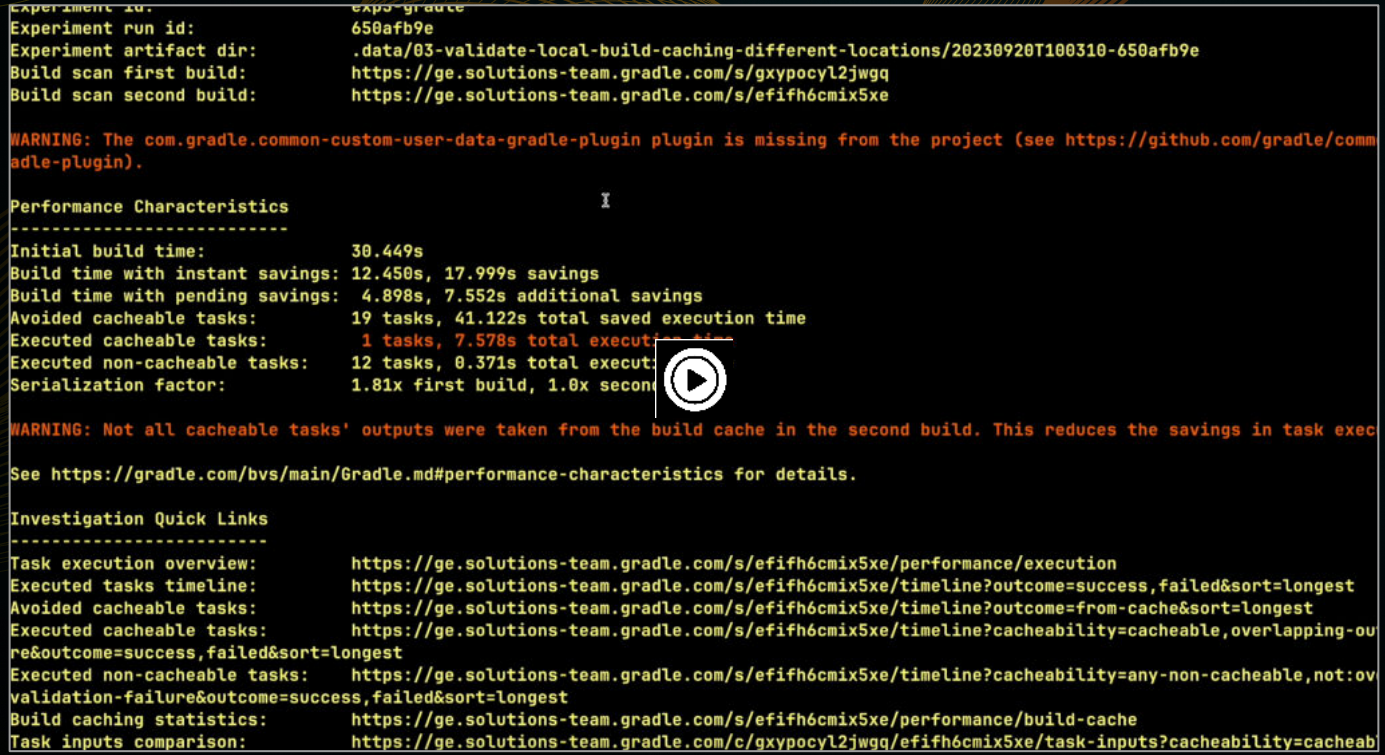

# Build caching regressions are costly

**Build changes that make a build no longer fully cacheable must be detected and notifications issued so the regression can be promptly investigated and fixed.**

# Automation of the regression detection process

**The Build Validation Scripts should be run on a continuous base in CI and cause the CI run to fail if a build is
no longer fully cacheable.**

./03-validate-local-build-caching-different-locations.sh -r https://github.com/spring-projects/spring-framework --fail-if-not-fully-cacheable

# Example: Continuously verified OSS projects

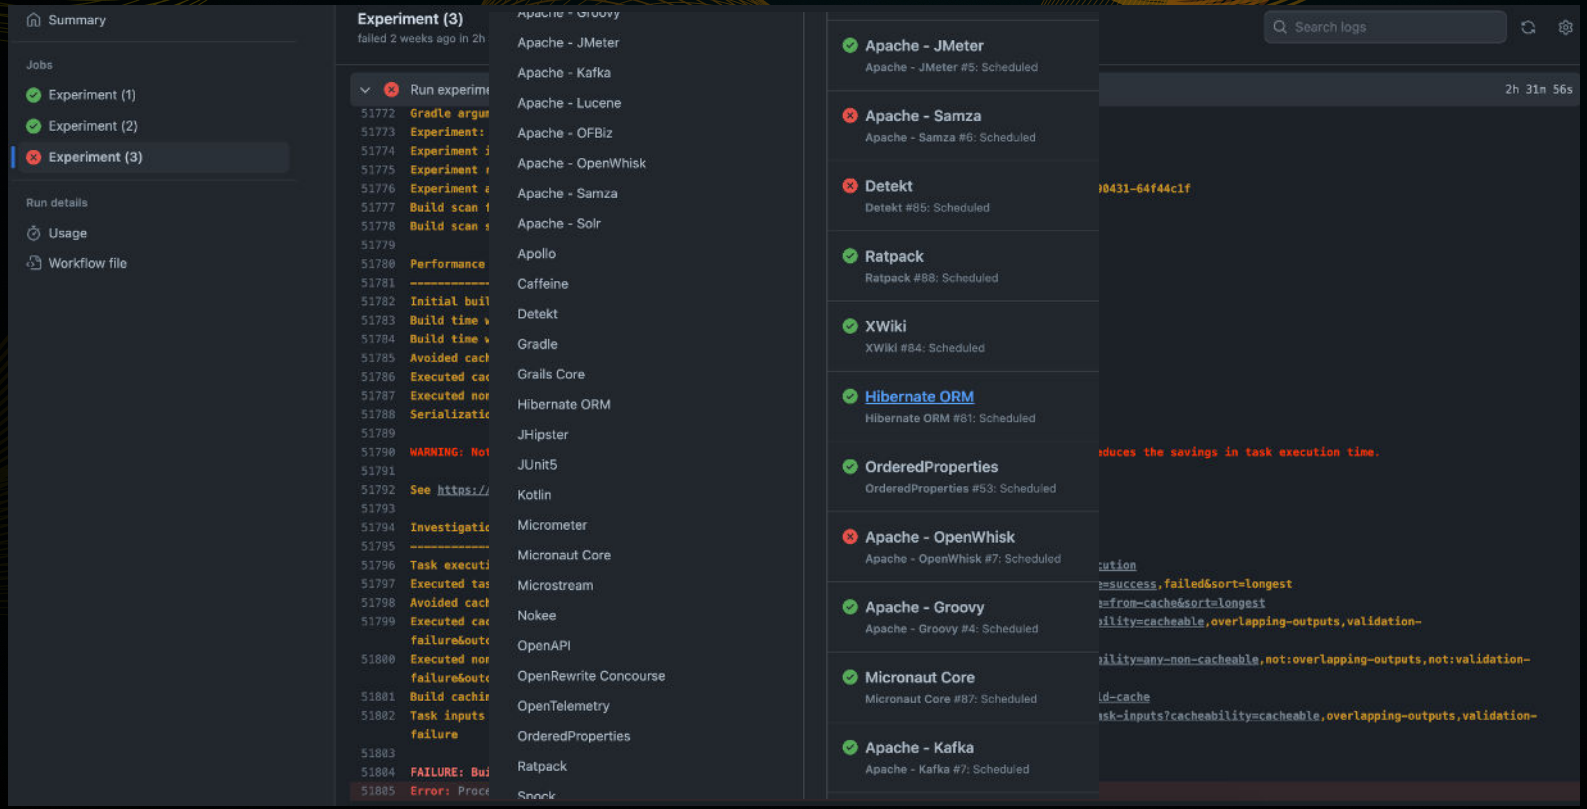


https://github.com/gradle/gradle-enterprise-oss-projects/actions

# Cross-machine regression detection

**The verification process can be extended by orchestrating the execution of two builds of the same version on different machines and submitting them to the Build Validation Scripts for verification of full cacheability.**

```
./04-validate-remote-build-caching-ci-ci.sh
-1 https://ge.solutions-team.gradle.com/s/
gxypocyl2jwgq
-2 https://ge.solutions-team.gradle.com/s/
efifh6cmix5xe --fail-if-not-fully-cacheable
```

# Regression detection generalization

**The verification process can be further generalized by automatically finding – amongst all builds – similar builds where full cacheability is expected and verifying them
for full cacheability.**

# Regression detection generalization

| | | | |
|---|---|---|---|
| **spring-authorization-server** | :spring-authorization-server-docs:asciidoctor | https://ge.spring.io/c/kgzhvgoftphtq/yf64zqfxta5xe | **Volatile input** |
| **spring-boot-build** | :spring-boot-project:spring-boot-actuator-autoconfigure:asciidoctor | https://ge.spring.io/c/6skln6qftnq2g/d6b2kc3yfgir4 | **Volatile input** |
| **spring-boot-build** | :spring-boot-project:spring-boot-actuator-autoconfigure:asciidoctorPdf | https://ge.spring.io/c/6skln6qftnq2g/d6b2kc3yfgir4 | **Volatile input** |
| **spring-boot-build** | :spring-boot-project:spring-boot-docs:asciidoctor | https://ge.spring.io/c/dbi36jrigy7im/ha3s5a3rlltj2 | **Volatile input** |
| **spring-boot-build** | :spring-boot-project:spring-boot-docs:asciidoctorMultipage | https://ge.spring.io/c/dbi36jrigy7im/ha3s5a3rlltj2 | **Volatile input** |
| **spring-boot-build** | :spring-boot-project:spring-boot-docs:asciidoctorPdf | https://ge.spring.io/c/dbi36jrigy7im/ha3s5a3rlltj2 | **Volatile input** |
| **javadoc-plugin** | :validatePlugins | https://ge.spring.io/c/vzkmbme7boyuy/6d7mcunkposza | **Eviction** |
| **spring-authorization-server** | :spring-authorization-server-docs:api | https://ge.spring.io/c/ejldvuvw7ushw/plvsiuinoiekw | **Eviction** |

Make your investment of achieving a fully cacheable build worthwhile by automating the continuous detections of caching regressions.

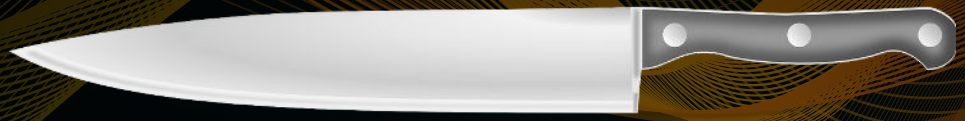

Keep the knife sharp!


THANKS
etienne@gradle.com
gradle.com