

DevProd for CI maintainers

DPE Summit 2023

Etienne Studer, SVP of Engineering, Develocity



Collect deep build data to understand the state of the developer toolchain on CI and make informed decisions what to improve.



Agenda

- **Capturing build data on CI**
- **Processing the captured build data**
- **Surfacing insights from the processed build data**





1

Capturing build data on CI



Collecting data about every CI build allows to understand the state of the developer toolchain on CI, including performance, reliability, and resource usage.



CI build data can be collected by configuring all projects' builds to capture and publish their build data to the server.

This requires the buy-in of the project owners and build modifications.

When there are many projects, data is required to prioritize which projects to reach out to first.



Alternatively, the CI build data can be collected by having the CI runner inject the configuration to capture and publish the build data into the invoked build.

This does not require touching the projects' builds.



CI plugins for Develocity

Jenkins Plugin

<https://github.com/jenkinsci/gradle-plugin>

Gitlab templates

<https://github.com/gradle/gradle-enterprise-gitlab-templates>

GitHub Gradle Build Action

<https://github.com/gradle/gradle-build-action>

TeamCity Plugin

<https://github.com/etiennestuder/teamcity-build-scan-plugin>

Bamboo Plugin

<https://github.com/gradle/gradle-enterprise-bamboo-plugin>



Example: Config-injection on Jenkins

Gradle Enterprise integration

Enable auto-injection

Gradle Enterprise connection settings

Gradle Enterprise server url [?](#)

Allow untrusted server [?](#)

Enforce Gradle Enterprise server url [?](#)

Gradle Enterprise access key [?](#)

The access key must be in the <server host name>=<access key> format. For more details please refer to the [documentation](#).

General settings

Auto-injection Git VCS repository filters [Beta] [?](#)

Check for the Gradle Enterprise build agent errors

Jenkins

Dashboard > gradle-freestyle > #3

- Back to Project
- Status
- Changes
- Console Output
- Edit Build Information
- Delete build '#3'
- Parameters
- Git Build Data
- Build Scans
- Previous Build
- Next Build

Build #3 (May 9, 2023, 10:56:14 PM)

No changes.

Started by anonymous user

Revision: 814b528652efbac9d54cdc2f33a33adab66cc93
Repository: <https://github.com/gradle/gradle.git>

- refs/remotes/origin/master

Build Scans

- <https://ge.mycompany.com/s/wjnl2qjzmu36>



Config-injection for Gradle and Maven

Gradle:

```
./gradlew build --init-script develocity-init.gradle
```

Maven:

```
mvn package -Dmaven.ext.class.path=develocity-extension.jar
```



All configuration related to capturing and publishing the build data to the server can be consolidated into a versioned convention plugin / extension.

This unifies the configuration of the data capturing, build caching, test acceleration, custom values, etc.

Example: [Gradle plugin](#) / [Maven extension](#)





2

Processing build data



Exporting build data into a big data store allows asking specific toolchain questions at scale.



Build models exposed by Develocity are described via OpenAPI specification. They currently cover build attributes, build cache performance, and project structure.

Build models are consumption-oriented.

Build model versions are backward-compatible.



Build models can be manually retrieved by writing a client that gets the data from the Develocity API.

This requires retries, resuming, paging, parallelizing, etc. and either ad-hoc analysis or saving in another store.

docs.gradle.com/enterprise/api-manual/ref/2023.3.html



Manually retrieve build models via API client

```
curl -s -H "$AUTH" "https://  
develocity.develocity.mycompany.com/api/builds  
?fromInstant=1694097000000&maxBuilds=3  
&query=tag:ci" | jq
```



Alternatively, build models can be automatically exported by Develocity and made available to a big data engine.

This allows to immediately ask specific toolchain questions.



Build model export, query & visualization



**Develocit
y**



AWS S3



AWS Athena



Grafana





Develocity – build model export

Develocity can be configured to export the build models to S3

Exported models are automatically updated on new versions

Newly available models are automatically exported.

```
{
  "id": "tzbyguougvos",
  "type": "gradle",
  "buildToolVersion": "7.3.3",
  "buildAgentVersion": "3.8.1",
  "modelVersion": {
    "string": "2023.1.6",
    "year": 2023,
    "release": 1,
    "patch": 6
  },
  "gradleAttributes": {
    "id": "tzbyguougvos",
    "buildStartTime": 1655007064459,
    "buildDuration": 13521,
    "gradleVersion": "7.3.3",
    "pluginVersion": "3.8.1",
    "rootProjectName": "object-storage-parent",
    "requestedTasks": [
      "tasks",
      "--all"
    ],
  },
  "hasFailed": false,
  "tags": [
    "CI",
    "HEAD",
    "github:action",
    "Linux"
  ],
}
```



AWS S3 – build model intermediate storage

Build models conform to the same OpenAPI-based schema as the Develocity API.

All build models of a build are in a single JSON stored in compressed format.


Stored build models are partitioned by build start time, one prefix per hour.

Designed for consumption by any Apache Hive-based engine.

Name	Type	Last modified
27d879e9e0c1...zip	obj	July 31, 2023, 19:17:10 UTC+02:00
21d8514a0d99...zip	obj	July 31, 2023, 19:16:10 UTC+02:00
23e9e520e9e0...zip	obj	July 31, 2023, 19:16:44 UTC+02:00
25e5e9d8f0e0...zip	obj	July 31, 2023, 19:05:15 UTC+02:00
24d9e9e0c1...zip	obj	July 31, 2023, 19:08:28 UTC+02:00
23e9e520e9e0...zip	obj	July 31, 2023, 19:16:44 UTC+02:00
25e5e9d8f0e0...zip	obj	July 31, 2023, 19:05:15 UTC+02:00
24d9e9e0c1...zip	obj	July 31, 2023, 19:08:28 UTC+02:00
23e9e520e9e0...zip	obj	July 31, 2023, 19:16:44 UTC+02:00
25e5e9d8f0e0...zip	obj	July 31, 2023, 19:05:15 UTC+02:00
24d9e9e0c1...zip	obj	July 31, 2023, 19:08:28 UTC+02:00

Name	Type	Last modified
starttime-2022-05-08	Folder	-
starttime-2022-06-01	Folder	-
starttime-2022-06-02	Folder	-
starttime-2022-06-03	Folder	-
starttime-2022-06-04	Folder	-
starttime-2022-06-05	Folder	-
starttime-2022-06-06	Folder	-
starttime-2022-06-07	Folder	-
starttime-2022-06-08	Folder	-
starttime-2022-06-09	Folder	-
starttime-2022-06-10	Folder	-
starttime-2022-06-11	Folder	-
starttime-2022-06-12	Folder	-





AWS Athena – build model query

Athena is a serverless big data query engine, with some caching to save query results.

The schema of the table is defined by the JSON schema of the build models, with each build model contained in a separate column.

SQL-like queries are run against tables and views and may include joins across build models.

```
Query 19 | X | Query 16 | X | Query 17 | X | Query 20 | X | Query 21 | X | Query 22 | X | + | -
1 CREATE EXTERNAL TABLE 'build' (
2 'id' STRING,
3 'buildStartTs' DATE,
4 'buildToolVersion' STRING,
5 'buildAppVersion' STRING,
6 'modelVersion' STRING, 'year' INT, 'string' STRING, 'patch' INT, 'release' INT,
7 'gradleAttributes' STRING, 'tags' ARRAY<STRING>, 'values' ARRAY<STRUCT<name: STRING, value: STRING, links: ARRAY<STRUCT<label: STRING
, url: STRING, buildStartTs: BIGINT, buildDuration: BIGINT, groupId: STRING, pluginVersion: STRING, nonProjectName: STRING,
'parentTasks': ARRAY<STRUCT<hasFailed: BOOLEAN, buildTool: STRING, buildOptions: STRUCT<buildToolPath: BOOLEAN,
'continueBuildEnabled': BOOLEAN, cleanEnabled: BOOLEAN, dryRunEnabled: BOOLEAN, excludeTasks: ARRAY<STRING>, offlineModeEnabled: BOOLEAN,
'runTasksEnabled: BOOLEAN, configurationCacheEnabled: BOOLEAN, configurationOnBandEnabled: BOOLEAN, continueOnErrorEnabled: BOOLEAN,
'filesystemWatchingEnabled: BOOLEAN, 'workspaceGradleUsers': INT, 'parallelProjectExecutionEnabled: BOOLEAN, 'refreshDependenciesEnabled: BOOLEAN,
, environment: STRUCT<username: STRING, operatingSystem: STRING, numberOfCores: INT, 'preScreen: STRING, 'javaVersion: STRING,
'jdkMemorySize: BIGINT, 'jvmCharset: STRING, 'jvmLocale: STRING, 'publicKeyName: STRING, 'localKeyName: STRING, 'localIpAddresses: ARRAY
<STRING>, 'hostnameFirstFailure: BOOLEAN, 'gradleEnterpriseSettings: STRUCT<noKubernetesEnabled: BOOLEAN, 'buildAgentCapturingEnabled
: BOOLEAN, 'taskInputsFileCapturingEnabled: BOOLEAN, 'buildAgentCapturingEnabled: BOOLEAN>,
8 'gradleBuildPerformance' STRUCT<id: STRING, 'buildTime: BIGINT, 'serialTaskExecutionTime: BIGINT, 'serializationFactor: DOUBLE, 'taskExecution:
ARRAY<STRUCT<duration: BIGINT, 'taskPath: STRING, 'taskType: STRING, 'avoidanceColors: STRING, 'fingerprintingDuration: BIGINT, 'avoidanceReasons:
BIGINT, 'nonCacheabilityCategory: STRING, 'nonCacheabilityReason: STRING, 'skipReasonMessage: STRING, 'cacheArtifacts: BIGINT,
'cacheReliabilityPercentage: BIGINT, 'avoidanceStrategy: STRUCT<total: BIGINT, 'ratio: DOUBLE, 'approach: BIGINT, 'localBuildCache: BIGINT,
'remoteBuildCache: BIGINT>, 'buildCaches: STRUCT<local: STRUCT<isEnabled: BOOLEAN, 'isEnabled: BOOLEAN, 'isSimplifiedForError: BOOLEAN,
'directory: STRING, 'remote: STRUCT<type: STRING, 'aliasName: STRING, 'url: STRING, 'isEnabled: BOOLEAN, 'isPushEnabled: BOOLEAN,
'isSubBuildCacheError: BOOLEAN, 'overhead: STRUCT<parsing: BIGINT, 'uploading: BIGINT, 'downloading: BIGINT, 'waiting: BIGINT>,
'effectiveTaskExecutionTime: BIGINT, 'effectiveWorkItemExecutionTime: BIGINT, 'serialBuildExecutionTime: BIGINT, 'taskFingerprintingStrategy: STRUCT
<count: INT, 'serialDuration: BIGINT>,
9 'gradleProjects' STRUCT<models: ARRAY<STRUCT<name: STRING, 'parent' INT, 'path: STRING>>
```

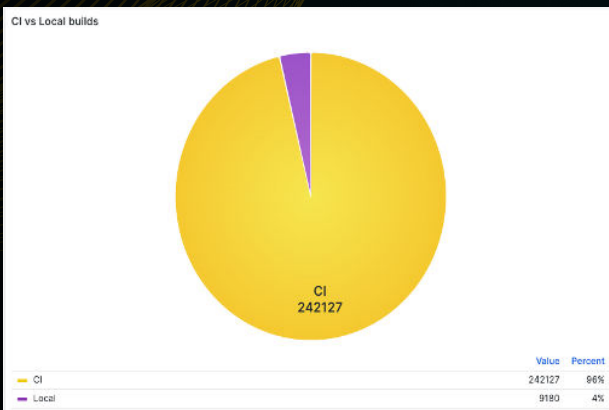




Grafana – build data visualization

Athena datasource is available for Grafana.

Charts are backed by queries into Athena tables and views.



Amazon Athena

Region: default (eu-central-1)

Data source: default (AwsDataCa...)

Database: default (etl)

Table: Choose

Column: Choose

Query result reuse

Enable:

TTL (mins): 60

```
1 SELECT CASE
2   WHEN isci = true THEN 'CI'
3   WHEN isci = false THEN 'Local'
4   ELSE 'Other'
5 END AS label, count(id)
6 FROM dashboarddata
7 WHERE $__dateFilter(startdate)
8 GROUP BY isci
```



IoC

Using Terraform to provision Athena resources.

Construct Grafana dashboard from JSON definitions.

Planned: Automatic creation and updates of Athena table definitions by Develocity





3

Surfacing insights from build data



Use the captured build data to identify usage and failure patterns, flakiness, build acceleration potential, resource waste, etc.



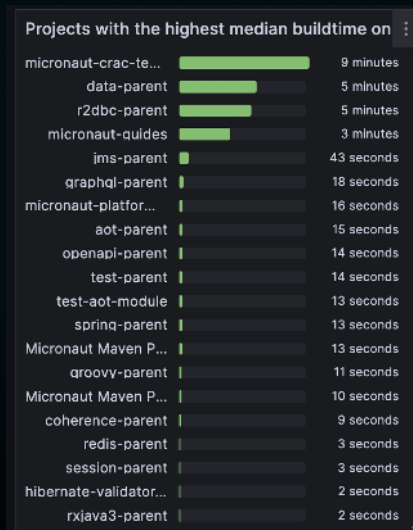
Overview

- How many projects, builds, build tools?



Contributors to build volume

- Projects with highest build count, highest total build duration, highest median build time?



Failure impact

- **Projects with highest build count of failed builds, total build duration, and median build time?**
- **Projects with highest total build duration of builds failing due to non-verification failures?**
- **Build agents with highest build count of failed builds, total build duration, and median build time?**
- ...



Build acceleration improvements

- **Projects that do not have build caching enabled with highest total execution and median execution time of generate & compile & test?**
- **Projects that have build caching enabled with highest total execution and median execution time of cacheable goals/tasks**
- **Projects that have build cache errors**
- **Projects that have long dependency download times**
- **Multi-module Maven or Gradle projects not built in parallel**
- **Multi-module Maven or Gradle projects built in parallel with only 1 worker**
- **Maven projects run with different Maven versions**
- **...**



**Prioritize actions for improvements
based on
quantitative impact analysis
and surfacing of the top-10 offenders.**





THANKS

etienne@gradle.com

gradle.com

