



**Big Data to Enable Global Disruption of the Grapevine-powered Industries**

# **D7.1 - Scalability and Robustness Experimental Methodology**

<b>DELIVERABLE NUMBER</b>	D7.1
<b>DELIVERABLE TITLE</b>	Scalability and Robustness Experimental Methodology
<b>RESPONSIBLE AUTHOR</b>	Ida Mele (CNR)



Co-funded by the Horizon 2020  
Framework Programme of the European Union

<b>GRANT AGREEMENT N.</b>	780751
<b>PROJECT ACRONYM</b>	BigDataGrapes
<b>PROJECT FULL NAME</b>	Big Data to Enable Global Disruption of the Grapevine-powered industries
<b>STARTING DATE (DUR.)</b>	01/01/2018 (36 months)
<b>ENDING DATE</b>	31/12/2020
<b>PROJECT WEBSITE</b>	<a href="http://www.bigdatagrapes.eu/">http://www.bigdatagrapes.eu/</a>
<b>COORDINATOR</b>	Panagiotis Zervas (Agroknow)
<b>ADDRESS</b>	110 Pentelis Str., Marousi, GR15126, Greece
<b>REPLY TO</b>	<a href="mailto:pythk@agroknow.com">pythk@agroknow.com</a>
<b>PHONE</b>	+30 210 6897 905
<b>EU PROJECT OFFICER</b>	Mr. Riku Leppanen
<b>WORKPACKAGE N.   TITLE</b>	WP7   Cross-sector Rigorous Experimental Testing
<b>WORKPACKAGE LEADER</b>	CNR
<b>DELIVERABLE N.   TITLE</b>	D7.1   Scalability and Robustness Experimental Methodology
<b>RESPONSIBLE AUTHOR</b>	Ida Mele (CNR)
<b>REPLY TO</b>	<a href="mailto:ida.mele@isti.cnr.it">ida.mele@isti.cnr.it</a>
<b>DOCUMENT URL</b>	<a href="http://www.bigdatagrapes.eu/">http://www.bigdatagrapes.eu/</a>
<b>DATE OF DELIVERY (CONTRACTUAL)</b>	30 September 2018 (M9), 30 June 2019 (M18, Updated Version)
<b>DATE OF DELIVERY (SUBMITTED)</b>	28 September 2018 (M9), 28 June 2019 (M18, Updated Version)
<b>VERSION   STATUS</b>	2.1   Final
<b>NATURE</b>	RE (REPORT)
<b>DISSEMINATION LEVEL</b>	PU (Public)
<b>AUTHORS (PARTNER)</b>	Ida Mele (CNR), Franco Maria Nardini (CNR), Raffaele Perego (CNR), Nicola Tonello (CNR), Vinicius Monteiro de Lira (CNR), Cristina Muntean (CNR), Salvatore Trani (CNR)
<b>CONTRIBUTORS</b>	Panagiotis Zervas (Agroknow), Todor Primov (ONTOTEXT), Aikaterini Kasimati (AUA), Giannis Stoitsis (Agroknow), Michalis Papakonstantinou (Agroknow), Panagiotis Rousis (Agroknow)
<b>REVIEWER</b>	Florian Schlenz (GEOCLEDIAN)

VERSION	MODIFICATION(S)	DATE	AUTHOR(S)
0.1	Initial draft	02/09/2018	Ida Mele (CNR), Franco Maria Nardini (CNR)
0.2	Added contributions	11/09/2018	Ida Mele (CNR), Franco Maria Nardini (CNR)
0.3	Added contributions	12/09/2018	Franco Maria Nardini (CNR), Raffaele Perego (CNR), Cristina Muntean (CNR)
0.4	First draft	14/09/2018	Franco Maria Nardini (CNR), Raffaele Perego (CNR), Ida Mele (CNR)
0.5	Added contributions	18/09/2018	Milena Yankova (CNR)
0.6	Input from partners	18/09/2018	Pythagoras Karampiperis (Agroknow), Antonis Koukourikos (Agroknow), Milena Yankova (ONTOTEXT), Vladimir Alexiev (ONTOTEXT)
0.9	Internal Review	21/09/2018	Stefan Scherer (GEOCLEDIAN)
1.0	Final edits after internal review	28/09/2018	Ida Mele (CNR), Nicola Tonello (CNR), Franco Maria Nardini (CNR), Raffaele Perego (CNR), Vinicius Monteiro de Lira (CNR), Cristina Muntean (CNR)
2.0	Update M18	26/06/2019	Ida Mele (CNR), Franco Maria Nardini (CNR), Raffaele Perego (CNR), Nicola Tonello (CNR), Vinicius Monteiro de Lira (CNR), Cristina Muntean (CNR), Mihalis Papakonstantinou (Agroknow), Panagiotis Zervas (Agroknow), Aikaterini Kasimati (AUA), Todor Primov (ONTOTEXT)
2.1	Consistency Update	15/01/2021	Salvatore Trani (CNR), Aikaterini Kasimati (AUA), Giannis Stoitsis (Agroknow), Michalis Papakonstantinou (Agroknow), Panagiotis Rousis (Agroknow)

PARTICIPANTS		CONTACT
<p>Agroknow IKE (Agroknow, Greece)</p>		<p>Nikos Manouselis Email: <a href="mailto:nikosm@agroknow.com">nikosm@agroknow.com</a></p>
<p>SIRMA AI (SAI, Bulgaria)</p>		<p>Todor Primov Email: <a href="mailto:todor.primov@ontotext.com">todor.primov@ontotext.com</a></p>
<p>Consiglio Nazionale Delle Ricerche (CNR, Italy)</p>		<p>Raffaele Perego Email: <a href="mailto:raffaele.perego@isti.cnr.it">raffaele.perego@isti.cnr.it</a></p>
<p>Katholieke Universiteit Leuven (KULeuven, Belgium)</p>		<p>Katrien Verbert Email: <a href="mailto:katrien.verbert@cs.kuleuven.be">katrien.verbert@cs.kuleuven.be</a></p>
<p>Geocledian GmbH (GEOCLEDIAN Germany)</p>		<p>Stefan Scherer Email: <a href="mailto:stefan.scherer@geocledian.com">stefan.scherer@geocledian.com</a></p>
<p>Institut National de la Recherche Agronomique (INRA, France)</p>		<p>Pascal Neveu Email: <a href="mailto:pascal.neveu@inra.fr">pascal.neveu@inra.fr</a></p>
<p>Agricultural University of Athens (AUA, Greece)</p>		<p>Katerina Biniari Email: <a href="mailto:kbiniari@aua.gr">kbiniari@aua.gr</a></p>
<p>Abaco SpA (ABACO, Italy)</p>		<p>Simone Parisi Email: <a href="mailto:s.parisi@abacogroup.eu">s.parisi@abacogroup.eu</a></p>
<p>SYMBEEOISIS EY ZHN S.A. (Symbeeosis, Greece)</p>	 Symbeeosis	<p>Konstantinos Rodopoulos Email: <a href="mailto:rodopoulos-k@symbeeosis.com">rodopoulos-k@symbeeosis.com</a></p>

## ACRONYMS LIST

BDG	BigDataGrapes
DSPS	Distributed Stream Processing Systems
IoT	Internet of Things
LDBC	Linked Data Benchmark Council
MIPS	Million instructions per second
MFLOPS	Million floating-point operations per second
MUDD	Multi-dimensional data generator
OWL	Ontology Web Language
RDF	Resource Description Framework
SPARQL	Symantec Protocol and RDF Query Language
SNB	Social Network Benchmark
SPB	Semantic Publishing Benchmark
SDPS	Streaming Data Processing Systems
TPC	Transaction Processing Performance Council

## EXECUTIVE SUMMARY

The deliverable D7.1, “Scalability and Robustness Experimental Methodology” consists in a report describing the methodology for assessing the performance of a big data system. In particular, the purpose of the task 7.1 is to develop and to implement a rigorous automated testing methodology for measuring and comparing the efficiency of the components in a big data system. The methodology considers the characteristics of the system as well as the heterogeneity and distributed nature of big data.

In this deliverable, we present the general concepts related to big data and its properties (i.e., volume, velocity, variety, and veracity). We analyze the state-of-the-art for big data benchmarking considering different challenges which range from preserving the 4V properties of big data to streaming and scalability issues. Moreover, we also provide a description of state-of-the-art tools that can be used for assessing and monitoring the big data system’s performance.

Besides reviewing the literature, we present the steps that can be followed for providing a rigorous testing of the BDG system. We believe that it would be better to follow a layered design where the user interfaces are at the top in order to provide easy access to the benchmarking for the user. Below the interfaces, there are the functional and execution layers. The former allows to capture the data and test generators as well as the metrics; the latter represents the basic operations for configuring the system, converting the data, and analyzing the results.

Another contribution of this deliverable is providing guidelines that can help the process of rigorous testing a big data system. We believe that a good approach would be following a standardized benchmarking methodology which is divided into different stages going from the selection of the application domain to the execution of the tests.

Since in the BDG project the semantic infrastructure is represented by graph databases, we also describe the main limitations of the current benchmarking in the context of relational databases and semantic repositories. Also, we provide some valid solutions for our project, for example, the benchmarks proposed by the Linked Data Benchmark Council (LDBC) which ensure linearity, reliability, repeatability, and easy to measure of the metrics. Additionally, LDBC is open for submissions of novel industry benchmarks which may represent specificity of data distribution in big data applications, and this makes it particularly suitable for the BDG project.

Another contribution of this deliverable is a proposal on the metrics to use for assessing the performance of the BDG system. Such metrics are chosen based on the datasets employed in the use cases and pilots of the BDG project. This deliverable has been updated at M18. In this new version we update the review of the literature with new contributions that have been published and we present an updated list of scalability metrics identified on the basis of the work done within the definition of the use cases and pilots.

This deliverable has been updated, despite not required, in order to be consistent with the other deliverables, by including the description of use case D (Food Protection) missing at the time of the second update at M18 and updating accordingly the list of the datasets used by the BDG project and the list of scalability metrics. Version 2.1 includes such an update and has been submitted on 15th January 2021.

## TABLE OF CONTENTS

1	INTRODUCTION.....	8
2	BIG DATA BENCHMARKING .....	10
2.1	LAYERED DESIGN OF BIG DATA BENCHMARKS.....	10
2.2	GUIDELINES.....	12
2.2.1	The benchmark test generation process. ....	13
2.2.2	The preparation of a big data system benchmark. ....	13
2.2.3	Benchmarking Requirements for BDG. ....	14
3	STATE OF THE ART .....	15
3.1	METHODOLOGIES .....	15
3.1.1	Scalability .....	15
3.1.2	Cloud Storage .....	16
3.1.3	Big data properties.....	16
3.1.4	Streaming .....	17
3.1.5	Security.....	18
3.2	TOOLS .....	18
3.2.1	Nagios.....	18
3.2.2	MetricBeat .....	19
4	EVALUATION OF THE BDG SEMANTIC COMPONENTS.....	20
5	BDG PERFORMANCE CRITERIA .....	23
5.1	SCALABILITY ISSUES AND METRICS.....	23
6	SUMMARY.....	37
7	REFERENCES.....	38

**LIST OF TABLES**

Table 1: Datasets of the BDG project with corresponding operations, format, and proposed metrics. ....24

Table 2: Wine Grapes Pilot.....29

Table 3: Wine Making Pilot.....30

Table 4: Farm Management Pilot..... 31

Table 5: Natural Cosmetics Pilot ..... 33

Table 6: Food Protection Pilot .....34

**LIST OF FIGURES**

Figure 1: Big Data Benchmarking Process ..... 10

Figure 2: Design Layers of a Big Data Benchmark..... 11

Figure 3: The big data generation process ..... 12

Figure 4: The benchmark test generation process ..... 13

Figure 5: The five-stage benchmarking methodology for a big data system..... 14

Figure 6: Example of a MetricBeat’s Module ..... 19



## 1 INTRODUCTION

Big data is a large umbrella under which we group any voluminous amount of structured, semi-structured, and unstructured data that has the potential to be mined along with the whole set of scalable approaches, tools, and methods used for processing and analyzing them. Big data represents an important research area due to the availability of large datasets that can be generated from various sources, e.g., Internet of Things (IoT), social media, and multimedia applications.

Big data systems allow to capture, store, manage, analyze, and visualize huge amount of heterogeneous data by exploiting highly scalable and robust tools and methods. Due to the complexity of the data managed and the scalability requirements, the standard techniques and methodologies traditionally used for measuring and assessing computational performance of software components suffer certain limitations while dealing with big data systems. To evaluate and compare the performance of big data systems, it is thus important to develop benchmarks and methodologies explicitly tailored for them.

In big data benchmarking, data is modeled with four dimensions: Volume, Velocity, Variety, and Veracity, which are also known as the “4V” properties of big data (Han et al.,2015). They are briefly described in the following of this section.

1. *Volume* represents the amount of data such as Terabyte (TB) or Petabyte (PB). Nowadays, data are generated faster than ever and the current speed is expected to increase exponentially over the next decade according to International Data Corporation (IDC). This means that big data generators must generate different volumes of data as input. The data volume has different meanings which depend on the workload. For example, in text processing (e.g., sort), the volume is represented by the amount of textual data. In graph analysis (e.g., social network), the volume is represented by the number of vertices in the graph. In image processing, the volume is represented by the number of pictures/second or the number of pixels/ images.
2. *Velocity* reflects the speed of generating, updating, or processing data which are also known as *data generation rate*, *data updating frequency*, and *processing speed*. The last one is particularly important in real-time applications.
3. *Variety* denotes the range of data types. Data can be structured (e.g., tables), unstructured (e.g., text, images, audio and video files), and semi-structured (e.g., web logs, .xls or .csv files). Big data benchmarking requires data generators that can support the whole spectrum of data types.
4. *Veracity* determines whether the data used in benchmarking reflects the characteristics of raw data. The best way to ensure veracity is to use real datasets. In case these datasets are not available, the alternative would be generating the synthetic datasets either randomly or from some statistic distributions. Since this could affect the veracity, another more recommendable approach would consist in synthetic data generation based on some examples from real datasets. Such synthetically generated datasets would be the input of workloads and guarantee the reality and credibility of the benchmarking results (Han et al.,2015).

In big data system benchmarking, it is of valuable importance to keep in mind these properties in order to make sure that they are preserved by the datasets used for the assessment of the system performance. If this is not respected the system may be tested over a too trivial dataset and the performance could be overestimated or conversely, the dataset may be too challenging, and the system performance could result poor. This is especially true when real datasets are not available or they need to be sampled, thus the data generation/sampling process has to ensure the 4V properties of big data. In Section 2, we will discuss the data generation requirements based on a layered design where the functional layer is devoted to preserve all the 4V properties of big data.

Various big data tools such as Hadoop, Storm, Spark are often employed in industry and research communities since they allow huge dataset to be distributed and processed in parallel. Big data applications use big data analytics techniques to efficiently analyze large datasets. However, choosing the suitable big data tools is difficult due to the challenges in processing and applying big data.

The document is structured as follows. We describe the high-level process for benchmarking big data systems in Section 2. This process consists of several steps, and it goes from planning the benchmarking to the final analysis and evaluation of the results. Then, we enlarge upon the process of benchmarking, providing a description of the layered design. Often the development of a benchmarking process is made of three layers: user interface, function, and execution layers. The function layer is the core of the designing, and we analyze it in more details. In particular, we explain the several steps that are needed to generate the datasets that would be used in the evaluation of the system performances. Such datasets must preserve the properties of big data, especially when they are created synthetically or sampled from the real data. Besides data, at this layer the tests must be decided, and they depend on the metrics chosen for the evaluation. Section 2 concludes with the process for generating tests, the five-stage benchmarking methodology, and the requirements for benchmarking the BDG project. The guidelines are general on purpose since the BDG system is not in the final stage and can change during the project.

In Section 3, we review the state-of-the-art in terms of methodologies and tools. Previous research works on benchmarking suites for big data analytics and databases are described. We organize our literature review in different topics based on scalability, cloud storage, preserving the properties of big data, streaming, and security.


In Section 4, we present the limitations of the most popular relational benchmarks and semantic repositories (e.g., UniProt, DBpedia) when applied to the BDG project. We propose to use the benchmarks by the Linked Data Benchmark Council (LDBC) which seem to be particularly suitable for big data projects as they ensure the linearity, reliability, repeatability, and easy to measure for the performance metrics. Moreover, we provide details on the Shapes Constraint Language (SHACL) which is a data-modeling language used for defining constraints in RDF graphs.

In Section 5, we describe the use cases of the BDG project, and we also propose some metrics that would be used in the benchmarking process for rigorous testing the efficiency of the BDG system. To highlight the metrics that we propose for the assessment of the system performance, we show in a table the use cases grouped by datasets and for each dataset we highlight the operations, formats and the corresponding metrics that can be employed during the assessments. Furthermore, we show for each pilot of the BDG project the datasets employed with the corresponding metrics. In this new updated version of the deliverable (M18) we update the review of the literature with new contributions that have been published and we present an updated list of scalability metrics identified on the basis of the work done within the definition of the use cases and pilots. This section has been updated on 15th January 2021, resulting in the final 2.1 version, in order to include the Food Protection use case missing in the previous update at M18 and by updating the list of datasets used within the BDG project and the list of scalability metrics.

Lastly, Section 6 concludes the deliverable.

## 2 BIG DATA BENCHMARKING

Big data benchmarks are developed to evaluate and compare the performance of big data systems and architectures. Figure 1 shows the benchmarking process for a big data system. It basically consists of five steps: (1) the *planning* step, at which the benchmarking object, the application domain, and the evaluation metrics are determined; (2) the *data generation* step when data is created; (3) the *test generation* step at which the possible tests are determined followed by (4) the *execution* step when the benchmark is conducted and finally (5) the *evaluation and analysis* step which takes care of analyzing and reporting the results.



**Figure 1: Big Data Benchmarking Process**

Successful and efficient benchmarking can provide realistic and accurate measuring of big data systems and thereby it promotes the development of a big data technology. Also, it helps the system owners to make decisions for planning big data systems. For example, benchmarking results can help the system owners to identify the performance bottlenecks in big data systems. The analysis of such bottlenecks can then be leveraged for optimizing the system configuration and the resource allocation.

### 2.1 LAYERED DESIGN OF BIG DATA BENCHMARKS

Figure 2 shows the design layers of a big data benchmark. There are three layers:

1. *User Interface Layer*. It provides interfaces which help the system owners to specify their benchmarking requirements (e.g., data, workloads, metrics and the preferred data volume and velocity).
2. *Function Layer*. It has three components: *data generators*, *test generators*, and *metrics*. Briefly, *data generators* are designed to produce datasets covering different data types and application domains, ensuring the 4V properties of big data in these datasets. The *test generators* automatically generate tests with comprehensive workloads for big data systems. The *metrics* can be tailored for the user needs or the architecture.
3. *Execution Layer*. It offers functions for supporting the execution of the benchmark tests over the different software stacks. Specifically, the system configuration tools can be used for running a generated test in a specified software stack. The data format conversion tools are used to transform a generated dataset into a format suitable for the test. Finally, the *result analyzer* and the *result reporter* analyze and display the evaluation results.




Figure 2: Design Layers of a Big Data Benchmark

**Functional Layer.** In this paragraph, we describe the different components of the *functional layer* which represents the core of the layered design of a benchmark. The *functional layer* has 3 important components:

1. *Data Generators.* They aim to generate datasets preserving the 4V properties of big data. The data can be represented by textual content, graphs, or tables. Figure 3 shows the process for generating datasets which can be summarized with 4 steps. The first step is for selecting the dataset so that the *variety* of big data is ensured. At this step, real datasets are selected to cover representative application domains and different data sources and types. Sometimes, the generator tools are employed to directly make synthetic datasets, which means that the datasets are independent of real data. At the second step, there is the data processing. Each data generator employs a data model to capture and preserve the important characteristics in one or multiple datasets of a specific data type. At this step, different models have to be developed to capture the characteristics of data and preserving the data *veracity*. Moreover, different data types (e.g., tables, text, streams, and graphs) have to be considered. At the data processing step, the data sampling is also important as it allows to scale down the dataset sizes. At the third step, there is the data generation so that *volume and velocity* can be controlled according to the user requirements for supporting different data generation rates (e.g., parallel and distributed execution over multiple machines). At the fourth step, after the dataset is generated, the format conversion tools transform it into a format that can be used as the input data of a specific workload.
2. *Test Generators.* They generate automatic benchmarking tests. The aim of these generators is abstracting the workload of the current big data systems to create a set of *operations and workload patterns* that can be used in big data processing. As reported in Figure 2, test generators consist of three parts:
  - a. *Abstract operations.* They represent the abstracted processing actions performed on the datasets (i.e., operators). In the phase of generating tests, the operations can be divided into three categories: (i) element operation, (ii) single-set operation, and (iii) double-set operation, according to the number of datasets processed by them.

- b. *Abstract workload.* These are abstracted patterns which are designed to combine operations to form complex processing tasks. In the test generator, three workload patterns can be abstracted: (i) a single-operation pattern which contains one single operation; (ii) a multi-operation pattern which is for multiple operations; and (iii) an iterative-operation pattern which is for operations that are executed iteratively. The difference between a multi-operation and an iterative-operation pattern is that the former contains a finite number of operations, while the latter only provides stopping conditions, therefore the exact number of operations can be known only at run time.
  - c. *Prescriptions.* They include the information needed to produce a benchmarking test, such as datasets, a set of operations and the workload patterns, a method to generate workload, and the evaluation metrics.
3. The *metrics* (single or multiple) can be divided into two types: user-perceived and architecture metrics<sup>1</sup>. User-perceived metrics represent observable metrics that are easy to understand for the users (e.g., duration of a test, request latency, and total throughput). While the user-perceived metrics are used to compare performances of workloads of the same category, the architecture metrics are designed to compare workloads from different categories, and they should also take energy consumption and cost efficiency into account. Some examples of architecture metrics are million instructions per second (MIPS) and million floating-point operations per second (MFLOPS).

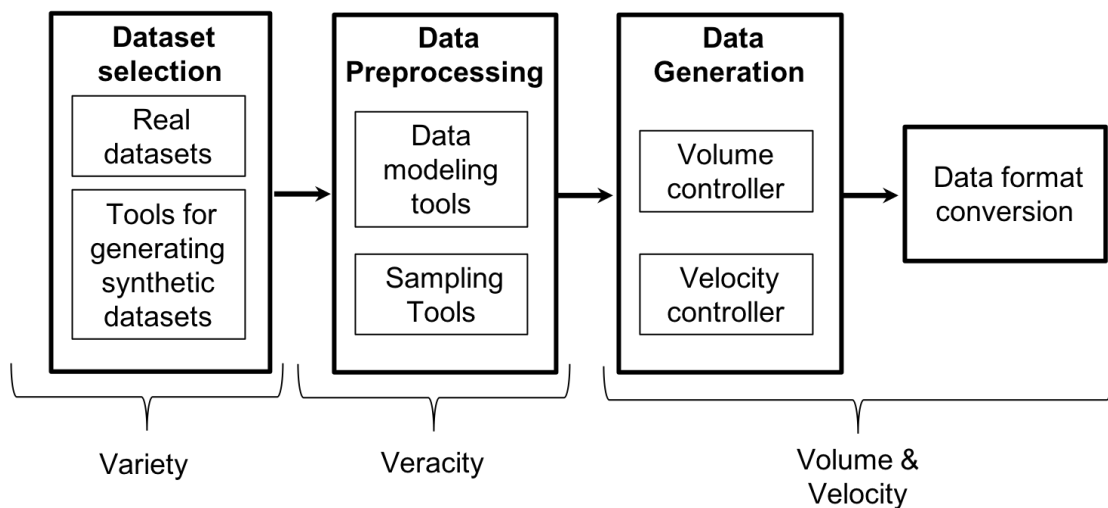


Figure 3: The big data generation process

## 2.2 GUIDELINES

As we have seen in the previous section, big data benchmarks are developed to evaluate and compare the performance of big data systems and architectures. Hence, benchmarking must provide realistic and accurate measuring of big data systems in order to address two objectives:

1. Promoting the development of big data technology, i.e., developing new architectures (processors, memory systems, and network systems), proposing innovative theories, algorithms and techniques to manage big data and to extract their value and hidden knowledge.
2. Assisting the system owners to make decisions for deciding the system features, for tuning system configurations, validating deployment strategies, and conducting other efforts to improve systems

<sup>1</sup> <http://prof.ict.ac.cn/BigDataBench/>

performance. For example, benchmarking results can identify the performance bottlenecks in big data systems, thus guiding the optimization of the system configuration and the resource allocation.

We now provide some general rules for testing a big data system for the BDG project.

### 2.2.1 The benchmark test generation process.

In Figure 4, we report the steps for generating a benchmark test. As we can see, at step 1, the dataset is selected. It can be a real dataset or a dataset generated by an automatic tool. At step 2, a set of abstracted operations that can be performed over the data must be decided. This is followed by step 3 at which the set of operations are combined based on the workload patterns. After that, at step 4, a prescription is generated. Lastly, at step 5, a test for a specific system and software stack is created based on the prescription and the system configuration tools. Following these steps, we can create automatically the workloads for the different use cases.




Figure 4: The benchmark test generation process

### 2.2.2 The preparation of a big data system benchmark.

Figure 5 shows a typical benchmarking methodology for big data systems which consists of five stages. At the stage 1, the application domain is selected followed by stage 2 for choosing the applications that are more suitable for the domain, e.g., the data models, data operations, workload patterns, and metrics. Stage 3 consists in implementing the data generation tools which will produce datasets with the 4V properties. Plus, at this stage the workloads are implemented to support application-specific benchmarking tests. At stage 4, the target system is determined, and the input data is prepared for testing the system. Such operations depend on a *prescription* which includes all the information needed to produce a benchmarking test, including input data, workloads, a method to generate test, and the evaluation metrics. Finally, the benchmark test is carried out at stage 5 and the results can be analyzed and evaluated.

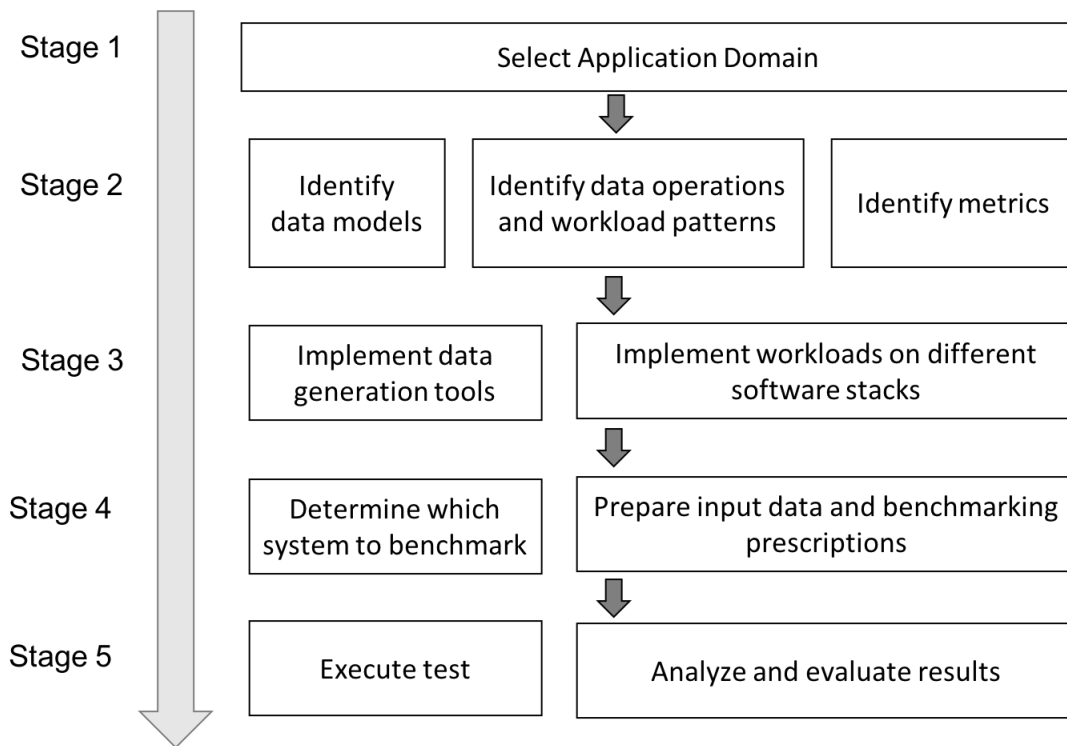


Figure 5: The five-stage benchmarking methodology for a big data system

### 2.2.3 Benchmarking Requirements for BDG.

When using real data as input of a big data system the data veracity is automatically guaranteed. However, very often it is difficult to use real data since the volume and velocity of this type of data cannot be used due to the different requirements. Hence, in some scenarios we may want to apply synthetic data which preserves the 4V properties of big data. As we have seen, the data should cover different data sources and types to ensure *variety*. This synthetic data can be used as workload of some tasks, such as *Sort* and *WordCount*, in micro benchmarks, and as input for database basic operations, such as *Read*, *Write*, and *Scan*.

To preserve the data veracity there must be different models which can capture the characteristic of real data of different types such as table, text, stream, and graph data. Then, volume and velocity are controlled to fulfill user requirements (as also shown in Figure 3). For example, the data generation can be parallel and distributed over multiple machines for supporting different data generation rates. At the final step, after a dataset is generated, the format conversion tools transform this dataset into an appropriate format which can be used as the input of a workload running on a specific system.

## 3 STATE OF THE ART

### 3.1 METHODOLOGIES

Several commercial and open source products have been released for big data storage and processing. Hence, the users can choose which system best suits their needs, and big data system developers have to decide how to evaluate their systems with regard to the big data processing needs. System benchmarking is the classic way of meeting the above requests.

Big data benchmark effort consists in implementing different workloads to evaluate some specific types of systems or architectures.

Regarding big data analysis, most of the benchmarks target the popular Hadoop platform (REN et al., 2012) Pavlo et al. (2009) presented a micro benchmark for big data analytics for comparing Hadoop-based analytics to a row-based Relational DataBase Management System (RDBMS) and a column-based RDBMS one. Also, Gridmix (2013) and HiBench (2010) are benchmarking suites for Hadoop MapReduce and Hive. While GridMix includes micro-benchmarks for only text data, Hibench covers more data types and software stacks. It has four categories of workloads and the inputs of such workloads can be datasets of fixed size or scalable and synthetic datasets. Finally, LinkBench (Armstrong et al., 2013) is based on traces from “social graph” databases (e.g., Facebook data). The authors characterized the data and query workload in many dimensions providing a realistic and challenging test for persistent storage of social and web service data.

Lastly, novel data-analytics approaches are requested for handling large datasets. Two common issues in big-data analytics are: (1) how to build high performance platforms to efficiently analyze big data? and (2) how to design a mining algorithm to find the useful things from big data? Solutions have been presented over the last years and they are reported in surveys on big-data analytics (Tsai et al., 2015) and (Zhang et al., 2016). The former summarizes recent work on knowledge discovery in large databases, while the latter categorizes the existing parallel data processing systems for big data.

#### 3.1.1 Scalability

In order to deal with scalability issues, Ming et al. (2014) proposed a Big Data Generator Suite (BDGS) for generating synthetic data representing big-data workloads. Using real datasets as seed, BDGS generates synthetic data by scaling the seed data and making sure that the characteristics of raw data is preserved. This approach allows to have fully scalability since each data generator can produce synthetic datasets, and its tools for the data format conversion can transform the format of the datasets. The users can specify their preferred data size, and the limit can only be bounded by the storage size and by the BDGS parallelism in terms of the nodes and its running time.

BigDataBench<sup>2</sup> (Wang et al, 2014) applied a similar idea to the scenario of internet services. It is based on a wide class of application domains (search engine, e-commerce, and social network). The effectiveness of this framework was tested over six real life datasets that cover three representative data types (structured, semi-structured, and unstructured data), and three data sources (text, graph, and table). Overall, the framework covers several micro benchmarks, three dominant internet services (i.e., Search Engine, Social Network, and E-commerce), and two fast emerging big data domains (multimedia and bioinformatics).

---

<sup>2</sup> <http://prof.ict.ac.cn/BigDataBench/>



### 3.1.2 Cloud Storage

Much work focused on benchmarking the cloud and datacenter infrastructure (Ferdman et al., 2010; Cooper et al., 2010). The Yahoo! Cloud Serving Benchmark (YCSB) (Cooper et al., 2010) was designed to provide tools for comparing different cloud data serving systems. This framework is extensible as it supports easy definition of new workloads and makes easy to benchmark new systems. Furthermore, it was used for comparing two non-relational databases (Cassandra and HBase) against one geographically distributed database (Yahoo's PNUTS) and a traditional relational database (MySQL). The CloudSuite benchmark (Ferdman et al., 2011) was implemented to test cloud service architectures using the performance counters of modern servers. The authors discovered that the modern processor micro-architecture is inefficient for running study scale-out workloads. Such inefficiency is due to the mismatch between the workload needs and modern processors, particularly in the organization of instruction and data memory systems and the processor core micro-architecture.

### 3.1.3 Big data properties

All the existing benchmarking techniques may have some limitations with respect to the 4V properties of big data. Concerning the *volume*, benchmarks like HiBench (2010) and LinkBench (Armstrong et al., 2013) are only partially scalable. Instead, BDGS (Ming et al. 2014) represents a fully scalable big data generator suite. Also, *velocity* is oftentimes underestimated, for example, LinkBench (Armstrong et al., 2013), CloudSuite (Ferdman et al., 2011) and BigDataBench (Wang et al., 2014) try to capture dynamic adjustment of data generation speed using parallel strategies. For the *variety*, most of the benchmarks support only one type of data (e.g., Hibench is suitable for unstructured data) and cannot handle multiple data types. Same limitation is encountered for the *veracity* since GridMix (2013), HiBench, and YCSB (Cooper et al., 2010) have a generation process of synthetic data which is independent of real raw data. In particular, HiBench randomly generates synthetic data using the programs in the Hadoop distribution or other statistic distributions. On the other hand, TPC-DS (Ghazal et al., 2013) and BigDataBench (Wang et al., 2014) try to preserve the characteristics of real data. In particular, the Transaction Processing Performance Council (TPC) proposes a series of benchmarks to test the performance of DBMSs in the decision support systems. One of them is the TPC-DS (Ghazal et al., 2013) which implements a multi-dimensional data generator (MUDD).

Han et al. (2015) reviewed all the benchmarking techniques and also proposed a comprehensive workload for big data systems by taking into consideration the big data issues, namely, generating data which preserves the 4V properties (i.e., volume, velocity, variety and veracity) and generating tests with comprehensive workloads for big data systems.

Similarly, Zhu et al. (2014) tried to generate comprehensive big data workloads called *BigOP*. It is an end-to-end system benchmarking framework featuring the abstraction of representative operations, workload patterns, and prescribed tests. Its abstraction model guides the development of the benchmark and also enables automatic generation of tests with comprehensive workloads. The framework was tested on three big data processing systems, i.e., Hadoop, Spark and MySQL Cluster involving relational data, text and graph data, as well as all operations and workload patterns.

Another example of end-to-end big data benchmark which models different types of data but for the specific business model of product retailers is BigBench (Ghazal et al., 2013). It covers a data model and synthetic data generator that addresses the variety, velocity, and volume aspects of big data systems containing structured, semi-structured, and unstructured data. The structured part of the BigBench data model is adopted from the TPC-DS benchmark, which is enriched with semi-structured and unstructured data components. The semi-structured part captures user clicks on the retailer's website. The unstructured data captures product reviews submitted online. The data generator designed for BigBench provides scalable volumes of raw data based on a scale factor. The BigBench workload is designed around a set of queries against a data model. From a business

prospective, the queries cover the different categories of big data analytics. From a technical prospective, the queries are designed to span three different dimensions based on data sources, query processing types, and analytic techniques. The feasibility of BigBench was shown by implementing it on the Teradata Aster Database.

Dimitrov et al. (2013) studied how the conventional optimizations, such as caching, prediction, and prefetching, can be applied to big data workload. They analyzed spatial and temporal reference patterns to bring out several insights related to memory and platform usages (e.g., memory footprints, read-write ratios, bandwidths and latencies). Indeed, trends may intersect by characterizing the memory access patterns of various Hadoop and noSQL big data workloads. Using memory DIMM traces collected using special hardware, the authors analyzed the spatial and temporal reference patterns to bring out several insights related to memory and platform usages, such as memory footprints, read-write ratios, bandwidths, latencies.

### 3.1.4 Streaming

Chintapalli et al. (2016) focused on real-time streaming benchmarks. Instead of testing speed-of-light event processing, they created a full data pipeline to closely mimic the real-world production scenarios and compared the performance of three data engines: Flink<sup>3</sup>, Storm<sup>4</sup> and Spark Streaming<sup>5,6</sup>

Karimov et al. (2018) presented a framework for benchmarking Streaming Data Processing Systems (SDPS). Differently from other works, their benchmark was specifically tailored for small structured messages. The evaluation consisted on measuring the performance for three open-source SDPS (i.e., Apache Storm, Spark, and Flink) with regard to the throughput and latency of windowed operations (e.g., windowed aggregations and joins). Such operations are important for applications like video games, which require fast processing of large-scale online data feeds from different sources and the windowed aggregations/joins are used for monitoring user feeds. The typical use cases are tracking the in-application-purchases and monitoring the advertisements. The workload adopted for this type of benchmark is derived from an online video game application called Rovio and monitors the user's actions in a given game to ensure that the services work as expected.

In “Benchmarking Distributed Stream Processing Platforms for IoT Applications”, Shukla and Simmhan (2016) developed a benchmark suite and performance metrics to evaluate Distributed Stream Processing Systems (DSPS) for streaming IoT applications. Their benchmark included 13 common IoT tasks which were classified across different functional categories in order to form micro-benchmarks. They also proposed two applications: statistical summarization and predictive analytics that leverage the dataflow compositional features of DSPS. The benchmark was validated using Apache Storm DSPS over datasets consisting of observations from smart cities and transportation. The contribution of their work is threefold: (1) they proposed a rigorous classification of the different characteristics of streaming applications and of their data sources; (2) they described performance metrics of DSPS that are necessary to meet the latency and scalability needs of streaming IoT applications; (3) they proposed an IoT Benchmark for DSPS based on representative micro-benchmark tasks. Overall, they aimed at evaluating DSPS in terms of performance and scalability and offer a baseline for big data researchers and developers to uniformly compare DSPS platforms for different IoT domains.

---

<sup>3</sup> <http://storm.apache.org/>

<sup>4</sup> <http://flink.apache.org/>

<sup>5</sup> <http://spark.apache.org/>

<sup>6</sup> <https://tsicilian.wordpress.com/2015/02/16/streaming-big-data-storm-spark-and-samza/>

In a follow-up work (Shukla et al., 2017), the authors classified common characteristics of streaming applications and their composition semantics and data sources. They also proposed categories of tasks that are frequently used by IoT applications and the key features of input data streams from sensors. The proposed real-time IoT benchmark is called RIoT Bench and it is based on 27 representative microbenchmark tasks.

### 3.1.5 Security

Cermak et al. (2016) in “A performance benchmark for NetFlow data analysis on distributed stream processing systems” presented a novel performance benchmark which is based on common security analysis algorithms of NetFlow data to determine the suitability of distributed stream processing systems. More in details, the authors contributed to give specifications of the challenges and requirements for distributed stream processing systems for effective flow analysis. Also, they proposed a benchmark for measuring the performance of a distributed stream processing system based on flow data processing. Lastly, they compared different distributed stream processing systems to quantify their suitability for the flow analysis. Although the benchmark can be used for any distributed stream processing systems, the authors took into account three popular systems: Samza, Storm, and Spark.

## 3.2 TOOLS

### 3.2.1 Nagios

Nagios is an open source tool for network monitoring<sup>7</sup>. It can be used for ensuring the proper functioning of systems, applications, services, and business processes of any IT infrastructures. More in detail, Nagios checks the network for a timely detection of problems that could be caused by data overload, crashed servers, or network connections. Also, it is able to monitor availability, uptime, and response time of every node on the network and it can deliver the results in a variety of visual representations and reports.

In case of a failure, Nagios alerts the technical staff promptly, allowing them to fix the problem or to take proper countermeasures before the outages may start affecting business processes and the services for end-users or customers.

Nagios consists of 3 main components:

1. *Nagios XI*. It allows to monitor the health of the entire network (e.g., reducing downtime, detecting network incidents). It can be installed on different platforms such as Microsoft, VMware, and Linux. Moreover, there are plenty of third-party add-ons for monitoring all in-house and external applications, services, and systems.
2. *Nagios Log Server*. It can be used for making sense of the data, identifying trends as well as finding security threats. It makes simpler the process of searching log data. It has no data limits and allows to get all the log data in one location ensuring high availability.
3. *Nagios Fusion*. It offers high visibility and scalability to the network. It helps to solve problems that come with multiple networks and geographical separation. It also allows to visualize multiple Nagios XI and Core servers in one location with the purpose of simplifying the network management.

---

<sup>7</sup> <https://www.nagios.com/>

### 3.2.2 MetricBeat

In the BDG project we need to analyze data in an effective and efficient way, so we plan to use MetricBeat<sup>8</sup>, a lightweight tool for data analytics. It can collect metrics for monitoring the servers and their services, and it is often employed for statistics on systems and services, such as system-level CPU usage, memory, file system, disk IO, and network IO statistics for every process running on the system.

MetricBeat is compatible with the most popular operating systems (i.e., Linux, Windows, and Mac OS). Also, it is part of the Elastic Stack<sup>9</sup> and can work with Logstash, Elasticsearch, and Kibana. Elasticsearch has become popular among Web companies interested in monitoring the provided services for improving their effectiveness, availability, and incident analysis or prediction. These companies often use Elasticsearch to collect data about their services, but very often they cannot easily extract information from such datasets due to the huge amount of data and its complexity. For this reason, machine learning techniques are employed to leverage huge data and extract useful information. As an example, Elasticsearch plus X-Pack can be used for anomaly detection in the infrastructure and applications.

MetricBeat makes the data shipping to Elasticsearch and Kibana (for visualization) easy. Due to its versatility and ease of use, we plan to use MetricBeat for monitoring the metrics of the different components of the BDG system. The tool uses modules and sets of metrics (metricsets). A module defines the basic logic for collecting data from a service (e.g., how to connect, which metrics to collect, how often to collect metrics). Each module has one or more metricsets. This allows to retrieve with a single request to the remote system a list of related metrics rather than each metric as a separate event. Figure 6 reports an example of a MetricBeat module. The metric retrieval is performed by periodically interrogating the host system based on the specified period. Because multiple metricsets can send requests to the same service, MetricBeat reuses connections whenever it is possible. If MetricBeat cannot connect to the host system within the time specified by the timeout value, it returns an error (e.g., it sends an event error to the specified output, allowing the monitoring and diagnosis for errors). More details about MetricBeat and the metrics used for monitoring the BDG system are in the deliverable 7.2.

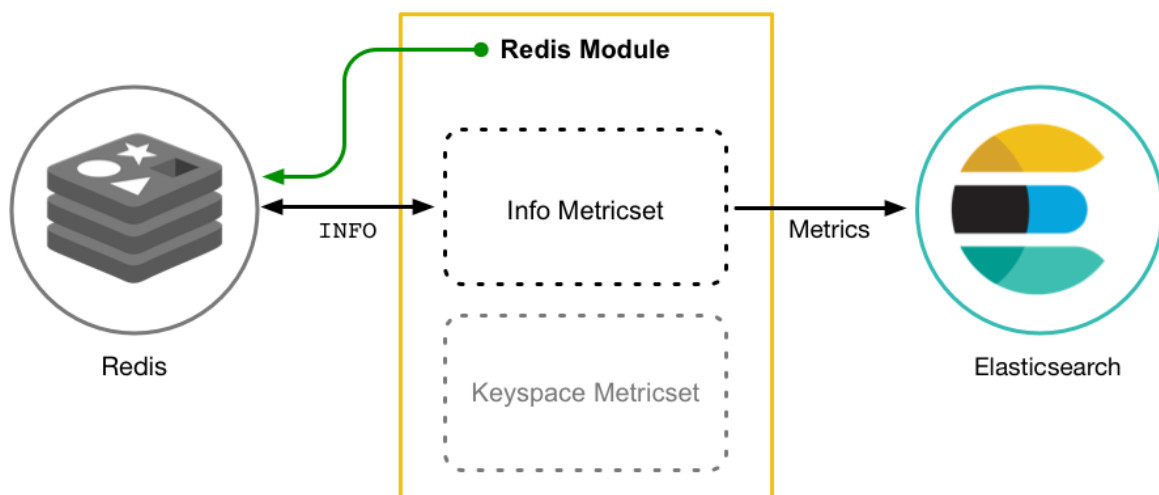


Figure 6: Example of a MetricBeat’s Module

<sup>8</sup> <https://www.elastic.co/products/beats/metricbeat>

<sup>9</sup> <https://www.elastic.co/>

## 4 EVALUATION OF THE BDG SEMANTIC COMPONENTS

In the context of relational databases there exist some well-established benchmarks, namely the well-known **TPC benchmarks**<sup>10</sup>. Although relational databases are not considered in this project, relational benchmarks introduced several best practice techniques that could serve as inspiration for benchmarking.

One particular limitation of the most influential and commonly used TPC benchmarks is that it uniformly distributed and uncorrelated data produced by generators of data and tests. While this simplifies the setup, it also ignores very challenging big data problems. Another influential benchmark is Star Schema Benchmark<sup>11</sup> which uses a schema that is closer to the canonical form. A star schema is a normalized schema that has a main table and several dimension tables (fact tables) that represent the points of the star.

Since graph databases that form the backbone of the Semantic infrastructure of BigDataGrapes aim at different types of queries (e.g. subgraph, supergraph, pattern match, reachability, shortest path query, etc.), these widespread benchmarks are not adequate for evaluating their performance.

Benchmarking RDF systems presents different challenges than the ones posed in relational database engines. As a result, existing relational benchmarks are not really suitable for RDF benchmarking. This is due to the intricacies of RDF data, which a) are expressed using a simple data model based on the concept of triple and b) is not necessarily accompanied by a schema.

The most popular benchmark for semantic repositories developed to facilitate the evaluation of semantic repositories in a standard and systematic way with support for RDF and OWL is **Lehigh University Benchmark (LUBM)**, defined in “An Evaluation of Knowledge Base Systems for Large OWL Datasets” (Guo et al., 2004). The purpose of the benchmark is to measure the performance of storing and querying of large amounts of data that are created for realistic Semantic Web systems. It employs synthetically generated datasets using a fixed OWL ontology of university organizations, lecturers, teachers, students and courses. It consists of a university domain ontology, customizable and repeatable synthetic data, a set of test queries, and several performance metrics. The complexity of the language constructs used is between OWL Horst and OWL DLP (Fischer et al., 2008; Ontotext Lab., 2009). 14 queries are defined that are used to check the query evaluation correctness and speed of repositories that have loaded a given dataset. The biggest standard dataset is LUBM(50) (i.e. it contains synthetic data for 50 universities). Its size is 6.8 million explicit statements, distributed in 1,000 RDF/XML files with total size of 600 megabytes. For the purposes of scalability measurements many groups have used the LUBM generator to create bigger datasets. This is the dataset which is adapted practically by all major semantic repository vendors. Through the years it played a considerable role in the development of the field. LUMBL benchmark evaluates the performance with respect to extensional queries over a large data set against single ontology in OWL which makes it restricted and does not replicated the big data needs of BigDataGrapes project.

**UniProt (Universal Protein Resource)**, <http://www.uniprot.org>) is the world's most comprehensive and most popular dataset of information on proteins, created by joining several other resources (Swiss-Prot, TrEMBL, and PIR). UniProt RDF<sup>12</sup> is an RDF representation of the dataset; it is based on an OWL ontology, expressed in a sub-language of OWL Lite, that is more expressive than OWL Horst, but still tractable. It represents one of the largest datasets distributed in RDF and OWL. Processing UniProt is often used as benchmark for scalability and reasoning capabilities of semantic repositories. Still, very few of the repositories are capable to load UniProt and perform materialization on top of it or to interpret its semantics otherwise.

---

<sup>10</sup> <http://www.tpc.org>.

<sup>11</sup> [http://simile.mit.edu/wiki/Dataset:\\_Barton](http://simile.mit.edu/wiki/Dataset:_Barton).

<sup>12</sup> <http://dev.isb-sib.ch/projects/uniprot-rdf>

**DBPedia**<sup>13</sup> is a dataset represented in RDF the infobox of Wikipedia together with other information related to or derived from Wikipedia. It serves as a connectivity hub of the Linked Open Data (LOD) initiative<sup>14</sup> The diversity of the information represented in DBPedia and the fact that it represents encyclopedic knowledge, makes it an excellent resource for benchmarking semantic repositories. DBPedia version 3.3 includes 362 million unique statements without the **owl:sameAs** links. When used for benchmarking of semantic repositories it has several advantages, compared to the straight-forward usage of DBPedia, and it is more diverse as it represents data and data modelling patterns from several other datasets as well. One should consider for instance that DBPedia and Geonames are datasets of a very different nature. A repository which performs well on Geonames could show poor performance when dealing with DBPedia for various reasons; one of them being that Geonames uses few tens of different predicates, whereas DBPedia uses around a hundred thousand predicates, which makes it difficult to apply to BigDataGrapes datasets.

**Berlin SPARQL Benchmark**, (Bizer and Sch, 2008) is another benchmark that focuses on evaluating performance of query engines in an e-commerce use case: searching products and navigating through related information. Randomized “query mixes” (of 25 queries each) are evaluated continuously through a client application that communicates with the repository through SPARQL endpoint. The benchmark enables evaluation with respect to the changing sizes of the dataset and differing numbers of simultaneous clients.

Although created for the benchmarking of SPARQL engines, the design of BSBM favors relational databases and other raw-store-based implementations, as long as they deal with a single fixed data schema and uniform dense data representation. Generally, RDF databases are designed to deal efficiently with diverse data, integrated from multiple data sources, encoded against different data schema, resulting in sparse data tables in relational databases.

A particularly suited set of benchmarks for desired graph database functionalities is **LDBC**<sup>15</sup> (**Linked Data Benchmark Council**). As a starting point, LDBC considers the following proposal of desirable characteristics for performance metrics (Lilja, 2005):

- **Linearity:**
  - It means that the value of the performance metric should be linearly proportional to the actual system performance. That is, if the value of the metric changes by a ratio, the actual performance of the system should change by the same ratio.
- **Reliability:**
  - A performance metric is reliable if a system A always outperforms system B when the corresponding values of the metric for both systems indicate that system A should outperform system B (i.e. larger value better performance)
- **Repeatability:**
  - A performance metric is repeatable if the same value of the metric is measured each time the experiment is performed. It implies that a good metric is deterministic.
- **Easy to measure:** If a metric is not easy to measure, it is unlikely that anyone will actually use it

LDBC benchmarks includes performance metrics as well as cost-based metrics (price/performance). Additionally, LDBC considers new approaches according to the requirements in current technologies and application domains. For example, a metric for robustness can be used to measure the ability of a database to

---

<sup>13</sup> More information about DBPedia is provided by its developers in Chapter 6 "Semantic Annotation and Retrieval: Web of Data".

<sup>14</sup> <http://esw.w3.org/topic/SweoIG/TaskForces/CommunityProjects/LinkingOpenData/>

<sup>15</sup> <http://ldbncouncil.org/>

perform well under a variety of conditions, including adverse runtime conditions such as unexpected data skew or resource contention (Wierner et al., 2009).

The LDBC currently has developed two new benchmarks: one for the [Social Network Benchmark \(SNB\)](#) and one for the [Semantic Publishing Benchmark \(SPB\)](#). The SNB aims at testing graph data management technologies for three scenarios: interactive (transaction query workload), business intelligence (analytical query workload) and graph analytics (graph analysis algorithms, such as PageRank). A wide variety of systems could potentially execute one or more workloads of the SNB. Currently, the Interactive Workload is in draft release stage. The Semantic Publishing Benchmark is based on the BBC news website and models a mixed query and update workload with a limited amount of semantic inferencing.

LDBC is open for submissions of new industry specific benchmarks (datasets) to reflect the specificity of data distribution in big data applications which makes it particularly suitable for BigDataGrapes project.

**Shapes Constraint Language (SHACL)** is a common data modeling language developed and maintained by a W3C Working Group in 2014 with the goal to “produce a language for defining structural constraints on RDF graphs”. SHACL is used for validating RDF graphs according to a set of predefined conditions. These conditions are provided as the so called “shapes” but can be also other constructs expressed in the form of an RDF graph. Most of the key objectives for BDG project depend on the integration of heterogeneous data across various connected knowledge domains. SHACL is frequently used in formalization of the data description (commonly referred as “shaping the data”). The basic data modeling capabilities of SHACL allow the describing of the valid properties for each of the classes in the data set. Speaking about data validation capabilities, SHACL can be used for definition of constraints on data, which improves the quality of the data as it enforces compliance with predefined models instead of using procedural code.

SHACL standard is divided in two parts:

- SHACL Core - describes a core RDF vocabulary to define common shapes and constraints;
- SHACL-SPARQL - describes an extension mechanism in terms of SPARQL.

SHACL is implemented in RDF4J (RDF for Java) as a stackable sail that is notification-enabled. It stores SHACL Shapes in a special named graph, the SHACL\_SHAPE\_GRAPH (SHACLShapeGraph). The base sail needs to be a NotifyingSail as well. The implementation enables automated validation of constraints on RDF databases in RDF4J, which until now was only possible using ad-hoc approaches.

RDF4J comprises a large collection of libraries, utilities and APIs. The important components for this section are:

- RDF4J classes and interfaces (API) - provide a uniform access to the SAIL components from multiple vendors/publishers;
- RDF4J server application.

Programmatically, GraphDB can be used via the RDF4J Java framework of classes and interfaces.

## 5 BDG PERFORMANCE CRITERIA

### 5.1 SCALABILITY ISSUES AND METRICS

We now propose some metrics that can be used for assessing the performance of the big data system tailored on the BDG project. The project has three use cases which share datasets and hence bottlenecks that would need to be considered in the benchmarking.

**Use Case A. Earth Observation Data Anomaly Detection and Classification.** The purpose of this use case is to develop models that differentiate between Earth Observation data issues and anomalies. This is for triggering warnings to farmers concerning farm management practices or damage events. The datasets involved in this use case are: Sentinel-2 and Landsat-8, provided by GEOCLEDIAN partner. Plus, data from AUA PA Lab consisting of: soil and elevation maps, IoT stationary data, canopy sensing and laser scanner as well as images captured by drones. The first type of data is in GEOTIFF and PNG format plus metadata in JSON format. The remaining type of data is mostly semi-structured such as .csv and .xls files.

**Use Case B. Yield and Quality Prediction.** The purpose of this use case is to leverage historical earth observation data combined with additional relevant information to make educated guesses about yield and wine quality. Regarding the yield prediction, GEOCLEDIAN would provide Earth Observation data (e.g., Sentinel-2, Landsat-8). Other datasets for weather, soil, elevation, exposition, variety, anomalies (e.g., pests/diseases), management practices, and historical yield/quality are provided by other partners. Earth Observation datasets are in GEOTIFF and PNG format plus metadata in JSON format. For predicting biological efficacy, the data would be provided by AUA PA Lab and APIGEA. Most of this data is semi-structured (.csv and .xls files). Finally, for the crop quality prediction, data used for predicting biological efficacy would be enriched with data collected by proximal sensors, multispectral and thermal cameras in the JSON, GEOTIFF, PNG format and by the land-based weather data which is binary.

**Use Case C. Farm management.** This use case would take care of optimizing the farm practices and of management. This would mean modelling climate, sunlight exposure, soil quality, slope and topography to predict the vine specific needs. Indeed, an inadequate management, such as over or under cropping, irrigating, spraying, inadequate pruning and poor canopy management, can affect the grape quality. Each plant has its own specific needs and can be monitored by sensors. AUA PA Lab would provide sensor data which would be integrated with the satellite imagery geo-spatial data from GEOCLEDIAN in order to initiate the management procedures required to produce high-quality grapes.

**Use Case D. Food Protection.** The purpose of this use case would be to enhance the current digital solution with new modules that will address further needs of the grape and wine supply chain. This includes 1) a software module that will be able to predict emerging and increasing risks for chemical hazards in the grapes and wines supply chain; 2) a price prediction dashboard that will include algorithms able to predict the prices of agricultural products, including grapes; 3) a food fraud dashboard that will help experts working in the food industry to perform an effective vulnerability assessment for products. The enhancement will mainly focus on the further development of Agroknow's Big Data platform with new software modules that will enable advanced data analysis and risk prediction using machine learning and deep learning methods, using food recalls and border rejections, pricing data, surveillance studies data about hazards and fraud in food products, laboratory testing data, country risk and corruption index data.

The use cases rely on satellite imagery geo-spatial data. Due to the heavy task of image processing, the main bottleneck would be represented by the num. of images/sec and the num. of pixels/sec in images. Moreover, other datasets with measurements and sensor data would be leveraged to predict the grape quality and improving the management and production in vineyards. Due to this variety of data (i.e., GEOTIFF, JSON, and



PNG plus .csv, .xls, and binary files), the benchmarking should be able to model heterogeneous data and to handle their potentially huge size. Finally, one would need fast processing of streaming data in order to take decisions on-the-fly.

We summarized in a table the datasets with the purpose of highlighting the metrics that would be used to evaluate the performance of the system during the benchmarking process. In Table 1, for each dataset we report its operations, the format and the proposed metrics. We would like to highlight that the proposed metrics are high level and can be used in combination with many other standard low-level metrics, such as CPU, memory, and network usage to assess the overall performance of the BDG platform. Lastly, for each dataset, we show the use cases that make use of it. As we can see, several datasets are shared by different use cases.

**Table 1: Datasets of the BDG project with corresponding operations, format, and proposed metrics.**

Name	Operations	Format	Metrics	Use Case Where the Dataset is used
<b>Sentinel-2</b>	Preprocessing, Atmospheric Corrections, Parcel Extraction, Normalized Difference Vegetation Index (NDVI), other Vegetation Indices calculation	JSON, GEOTIFF, PNG	Num. of images/sec, Num. of pixel/image, Processing time, Request latency, Total Throughput	Use case A (Data Anomaly Detection and Classification) Use case B1 (Yield Prediction) Use Case B2 (Predicting Biological Efficacy) Use Case B3-1 (Crop quality prediction, table grapes) Use Case B3-2 (Crop quality prediction, winemaking) Use Case C1 (Opt. of Farm Practices) Use Case C2 (Mng. Zone Delineation)
<b>Landsat-8</b>	Preprocessing, Atmospheric Corrections, Normalized Difference Vegetation Index (NDVI), other Vegetation Indices calculation	JSON, GEOTIFF, PNG	Num. of images/sec, Num. of pixel/image, Processing time, Request latency, Total Throughput	Use case A (Data Anomaly Detection and Classification) Use case B1 (Yield Prediction) Use Case B2 (Predicting Biological Efficacy) Use Case B3-1 (Crop quality prediction, table grapes) Use Case B3-2 (Crop quality prediction, winemaking) Use Case C1 (Opt. of Farm Practices) Use Case C2 (Mng. Zone Delineation)
<b>IoT Stationary Data</b>	Data filtering for outliers	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Yield Mapping</b>	N/A	XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes)

				Use Case C2 (Mng. of Zone Delineation)
<b>Canopy Sensing and Vegetation Indexes</b>	Data filtering for outliers	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Drone Imagery</b>	Data filtering for outliers	GEOTIFF	Num. of images/sec, Num. of pixel/image, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Eca Sensing</b>	Data filtering for outliers	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Crop Calendar</b>	N/A	DOC, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Topographic Data and Elevation Maps</b>	N/A	CSV, XLS, XML	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Grape Berry Mechanical Properties</b>	N/A	XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Phenolic Composition</b>	N/A	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use case B1 (Yield Prediction) Use Case B3-1 (Crop quality prediction, table grapes) Use Case C2 (Mng. of Zone Delineation)
<b>Vine Leaf Extract Var. 1</b>	Ultrasounds, maceration	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B2 (Predicting Biological Efficacy)
<b>Vine Leaf Extract Var. 2</b>	Ultrasounds, maceration	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B2 (Predicting Biological Efficacy)
<b>Genetic Data</b>	Query on the databases	CSV	Num. of queries/sec Num. of rows/sec, Num. of columns/sec,	Use Case B3-2 (Crop quality prediction, winemaking)

			Num. of bytes/sec, Processing time	
<b>Soil/Plot Characteristics [Plot Management]</b>	Query on the information system	CSV, SHP	Num. of queries/sec Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Num. of pixels/sec, Num. of images/sec, Num. of pixel/image, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Harvest Information [Plot Management]</b>	Digitized had-written forms	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Climatic data</b>	Query on the databases	CSV	Num. of queries/sec Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Laboratory Analysis on Must, Grape and Wine [Grape and berry Mechanical and Chemical Properties] [Qualitative and Quantitative Characteristics of Must]</b>	N/A	CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Alcoholic Fermentation Monitoring - Online Data [Winemaking Activities]</b>	Query on the information system	CSV	Num. of queries/sec Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Winemaking Activities</b>	Digitized had-written forms	CSV	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>Sensory Analysis</b>	N/A	CSV	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case B3-2 (Crop quality prediction, winemaking)
<b>SENTEK DRILL &amp; DROP TRISCAN</b>	Comparison with water balance output	JSON, CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case C1 (Opt. of Farm Practices)
<b>Raingauge Metos Weather Station</b>	Water balance modelling	JSON, CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case C1 (Opt. of Farm Practices)
<b>Air Temperature/ Humidity Metos Weather Station</b>	Water balance modelling	JSON, CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case C1 (Opt. of Farm Practices)

<b>IR Temperature Meteos Weather Station</b>	Comparison with water balance output	JSON, CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time	Use Case C1 (Opt. of Farm Practices)
<b>Leaf Moisture Sensor</b>	Pest monitoring	JSON, CSV, XLS	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec. Processing time	Use Case C1 (Opt. of Farm Practices)
<b>Soil Profile</b>	Water balance modelling	SHP	Num. of pixels/sec, Num. of images/sec, Num. of pixel/image, Processing time	Use Case C1 (Opt. of Farm Practices)
<b>Lab testing data</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	xsl, csv	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec Num. of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Pricing data</b>	Preprocessing , enrichment with products and hazards terms, indices	xsl, csv, API	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec Num. of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>UK Food Standards Agency</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>RASFF – Rapid Alert System for Food and Feed</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	XML, Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Food Standards Australia New Zealand</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>FDA Recalls, Market Withdrawals, &amp; Safety Alerts, warning letters, import refusals and inspection citations</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	OpenFDA Food API, html, xsl, json, xml	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>EFET - Hellenic Food Safety Authority</b>	Preprocessing, parse pdf, transformation to internal data model enrichment with products and hazards terms, indices	Pdf, Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E

<b>Japanese Imported Foods Inspection Services</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html, xsl	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Czech Agriculture and Food Inspection Authority</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Healthy Canadians food alert information website</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Food Safety Authority of Ireland</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>German Food Safety: warnings and information to the public</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Hong-Kong-Center for Food Safety</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Open Food Facts</b>	Preprocessing , transformation to internal data model enrichment with products terms, store in SQL, indices	Html, API	Num. of recalls per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>ProMED-mail</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of outbreaks per week Num. of hazards covered Number of ingredients covered	Use Case D (Grape and Wine Quality Risk Assessment (safety)) Use Case E
<b>Food fraud news by EU</b>	Preprocessing , parse PDF, transformation to internal data model enrichment with products and hazards terms, indices	Html, pdf	Num. of fraud news per month Num. of fraud issues covered Number of ingredients covered	Use Case D

<b>Country risk indicators</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html, xls	Num. of indicators per month Number of countries covered	Use Case D
<b>News items from MEDISYS and professional websites</b>	Preprocessing , transformation to internal data model enrichment with products and hazards terms, indices	Html	Num. of outbreaks per week Num. of hazards covered Number of ingredients covered	Use Case D

The BigDataGrapes project consists of different pilots, and each of them is characterized by a set of datasets. We identify five pilots:

- *Wine Grapes*
- *Wine Making*
- *Farm Management*
- *Natural Cosmetics*
- *Food Protection*

In the following, we report one table for each pilot summarizing the name of the dataset together with its description, priority, provenance, data type format and approximative size.

**Table 2: Wine Grapes Pilot**

<b>Name</b>	<b>Description</b>	<b>Priority</b>	<b>Provenance</b>	<b>Data Type Format</b>	<b>Data Size</b>	<b>Metrics</b>
<b>Yield Mapping</b>	Yield data	Essential	Laboratory equipment	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Grape and Berry Mechanical Properties</b>	Measurements	Essential	Laboratory equipment	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Classical Analytical Techniques (HPLC)</b>	Phenolic composition data	Essential	Laboratory equipment	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Topographic Data and Elevation Maps</b>	Spatial data (boundaries and elevation data)	Essential	Remote sensing	CSV, XLS, XML	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Canopy Sensing and Vegetation Indices</b>	Canopy sensing data	Essential	Proximal sensors	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time

<b>IoT Stationary Data</b>	Soil moisture data, meteorological parameters	Essential	IoT data	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Drone Imagery</b>	Drone images	Essential	Multispectral and thermal cameras	GEOTIFF	GB	Num. of images/sec, Num. of pixel/image, Processing time
<b>Crop Calendar</b>	Records of crop growth stages and agricultural operations	Essential	Log files	DOC, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Eca Sensing</b>	Geo-referenced soil electrical conductivity data	Essential	Proximal sensors	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time

Table 3: Wine Making Pilot

Name	Description	Priority	Provenance	Data Type Format	Data Size	Metrics
<b>Genetic Data</b>	Genetic profile, morphological description, origin, etc.	Essential	French Network of Grapevine Repositories (Database of the collections)	CSV, API	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Soil Characteristics</b>	Texture, pH, etc.	Essential	Field measurement	XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Plot Management</b>	Treatments/fertilizing (when, what, how much), ground handling, or tasks related to the culture management, pest control, water status, yield, etc.	Essential	Field measurement	PDF, DOC, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Climatic Data</b>	Rainfall, temperature, radiation, etc.	Essential	Field measurement	XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Grape and Berry Mechanical and Chemical Properties</b>	Anthocyanin content, weight, length, width, density, etc.	Essential	Field measurement	XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time

<b>Qualitative and Quantitative Characteristics of Must</b>	Sugar content, alcohol, pH, etc.	Essential	Laboratory equipment	N/A	MB	N/A
<b>Winemaking Activities</b>	Bioconversion of sugar into ethanol and CO <sub>2</sub> , monitoring of alcoholic fermentation and sugar content, yeast characteristics, etc.	Essential	Laboratory equipment	XLS, PDF	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Sensory Analysis</b>	Expert panel of tasters' sensory analysis (wine bitterness, astringency, phenol content, aroma, etc.)	Essential	Expert analysis	XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Wine Commercial Information</b>	Number of bottles produced, number of bottles sold	Additional	Selling point	XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time

Table 4: Farm Management Pilot

Name	Description	Priority	Provenance	Data Type Format	Data Size	Metrics
<b>Sentinel-2</b>	Sentinel-2A/B MSI spectral bands, Vegetation index time series & advanced products	Essential	Copernicus EO Programme, ESA	JSON, GEOTIFF, PNG	150 GB/year* site	Num. of images/sec, Num. of pixel/image, Processing time
<b>Landsat-8</b>	Landsat-8 OLI spectral bands, Vegetation index time series & advanced products	Essential	USGS, NASA	JSON, GEOTIFF, PNG	6 GB/year* site	Num. of images/sec, Num. of pixel/image, Processing time
<b>Chemical and Physical Info on Grapes</b>	Anthocyanins, pH, brix values during maturation	Additional	Excel table file	XLS	TBD	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time



<b>Day by Day Activities in term of Treatments, Fertilization, Field Operation</b>	Diary where farmer or operators can record and /or plan all the activities on their fields	Essential	SITI4farmer	TXT	TBD	Num. of rows/sec, Num. of bytes/sec, Processing time
<b>Plot and Fields Information Georeferenced</b>	Information of plots position, shaping, cultures, type of seed, dates, and everything related on the culture and the farm itself (form official and not official point of view)	Essential	Form the field through SITI4farmer	JSON	TBD	Data insert duration, Processing time
<b>Relative Humidity</b>	Relative humidity (RH) is the ratio of the partial pressure of water vapor to the equilibrium vapor pressure of water at a given temperature	Essential	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time
<b>Air Temperature</b>	N/A	Essential	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time
<b>Global Solar Radiation</b>	Power per unit area received from the Sun in the form of electromagnetic radiation in the wavelength range of the measuring instrument	Additional	Field Sensors	JSON	200 Kb/year	Num. of bytes/sec, Processing time
<b>Wind Speed and Direction</b>	N/A	Essential	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time
	N/A	Additional	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time

<b>Soil Temperature</b>						
<b>Soil Moisture</b>	Measurement of the water in the large and intermediate size pores that can move about in the soil and be easily used by plants	Essential	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time
<b>Precipitation</b>	Rainfall measurements	Essential	Field Sensors	JSON	200 Kb/year	Num. of bytes/sec, Processing time
<b>Infrared Surface Temperature</b>	Temperature Surface calculated with infrared measurements	Essential	Field Sensors	JSON	400 Kb/year	Num. of bytes/sec, Processing time

Table 5: Natural Cosmetics Pilot

Name	Description	Priority	Provenance	Data Type Format	Data Size	Metrics
<b>Agiorgitiko Samples UAE (11 samples)</b>	Data on biological efficacy of samples of Agiorgitiko dried vine leaves, developed with Ultrasound Assisted Extraction	Essential	Laboratory testing	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Agiorgitiko Samples MAC (11 samples)</b>	Data on biological efficacy of samples of Agiorgitiko dried vine leaves, developed with Maceration	Essential	Laboratory testing	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Mandilaria Samples UAE (5 samples)</b>	Data on biological efficacy of samples of Mandilaria dried vine leaves, developed with Ultrasound Assisted Extraction	Essential	Laboratory testing	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time
<b>Mandilaria Samples MAC (5 samples)</b>	Data on biological efficacy of samples of Mandilaria dried vine	Essential	Laboratory testing	CSV, XLS	MB	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Processing time

	leaves, developed with Maceration					
<b>Weather Data</b>	Weather data on the regions selected for sample gathering	Essential	Open source data	N/A	N/A	N/A

**Table 6: Food Protection Pilot**

Name	Description	Priority	Provenance	Data Type Format	Data Size	Metrics
<b>Lab testing data</b>	Laboratory testing results for fruits, vegetables, meat, dairy and cereals collected from 34 countries	Essential	European Food Safety Agency, National Authorities of 33 EU countries and the United States	csv, xls, html, pdf	100G B per year	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Num. of ingredients covered
<b>Pricing data</b>	Pricing data for agricultural commodities and products including grapes from the European Union	Essential	Essential	csv, xls	500 MB per year	Num. of rows/sec, Num. of columns/sec, Num. of bytes/sec, Num. of ingredients covered
<b>UK Food Standards Agency</b>	Food recalls data announced by the National Authority	Essential	FSA - National Authority of the United Kingdom	Html	10 MB per year	Num. of recalls per week
<b>RASFF – Rapid Alert System for Food and Feed</b>	Food safety issues (recalls and border rejections) within the European Union that are collected by Rapid Alert System for Food and Feed (RASFF)	Essential	European Food Safety Agency	XML, Html	800 MB per year	Num. of recalls per week, Num. of hazards covered, Number of ingredients covered
<b>Food Standards Australia New Zealand</b>	Food recalls data announced by the National Authority	Essential	National Authorities of Australia and New Zealand	Html	10 MB per year	Num. of recalls per week, Num. of hazards covered, Number of ingredients covered
<b>FDA Recalls, Market Withdrawals, &amp; Safety Alerts,</b>	Food recalls and import refusals data announced by the National Authority	Essential	The website of Food and Drug Administratio	OpenFDA Food API, html, xls, json, xml	300 MB per year	Num. of recalls per week, Num. of hazards covered

warning letters, import refusals and inspection citations			n of the United States, Food Safety and Inspection Service			Number of ingredients covered
EFET - Hellenic Food Safety Authority	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Greece	Pdf, Html	300 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Japanese Imported Foods Inspection Services	Border rejections data announced by the National Authority	Essential	The website of Ministry of Health of Japan	Html, xls	700 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Czech Agriculture and Food Inspection Authority	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Czech Republic	Html	100 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Healthy Canadians food alert information website	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Canada	Html	10 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Food Safety Authority of Ireland	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Ireland	Html	10 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
German Food Safety: warnings and information to the public	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Germany	Html	100 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Hong-Kong-Center for Food Safety	Food recalls data announced by the National Authority	Essential	The website of National Food Safety Authority of Hong Kong	Html	10 MB per year	Num. of recalls per week Num. of hazards covered Number of ingredients covered
Open Food Facts	Dataset with information about the food product brands that is collected by consumers	Essential	Website of Open Food Facts	Html, API	200 MB	Num. of recalls per week Num. of hazards covered Number of ingredients covered

<b>ProMED-mail</b>	Information about the health outbreaks collected by ProMED	Essential	Web site of the ProMED-mail organization	Html	100 MB per year	Num. of outbreaks per week Num. of hazards covered Number of ingredients covered
<b>Food fraud news by EU</b>	Food fraud cases that are collected and curated by the Knowledge Centre for Food Fraud and Quality	Essential	Web site of the Knowledge Centre for Food Fraud and Quality	Html, pdf	10 MB per year	Num. of fraud news per month Num. of fraud issues covered Number of ingredients covered
<b>Country risk indicators</b>	Dataset with risk indicators that are announced by the Organisation for Economic Co-operation and Development	Essential	The website of the Organisation for Economic Co-operation and Development	Html, xls	10 MB per year	Num. of indicators per month Number of countries covered
<b>News items from MEDISYS and professional websites</b>	News items about food sector that are collected by MEDISYS EU service	Essential	The MEDISYS EU service	Html	10 MB per year	Num. of outbreaks per week Num. of hazards covered Number of ingredients covered

## 6 SUMMARY

The deliverable 7.1, “Scalability and Robustness Experimental Methodology”, belongs to WP7 which is about “Cross-sector Rigorous Experimental Testing”. The purpose of this deliverable is to report on the methodology that is used for testing the scalability and robustness of a big data system and to propose the metrics for the assessment.

We have presented the process for implementing an automated and rigorous experiment which would allow to measure the efficiency of the big-data-system components. The benchmarking has to consider not only the characteristics of the system but also the properties of big data. As further contribution, we have provided the guidelines that are followed in the assessment of the BDG system.

Part of this deliverable is a review of the state-of-the-art about benchmarking methodologies. We also describe tools and technologies (e.g., Nagios and MetricBeat) that can be used for online monitoring of Big Data systems. Finally, we have proposed some metrics that are used in the assessment of the performance of the BDG platform.

This deliverable has been updated at M18. In this new version we update the review of the literature with new contributions that have been published and we present an updated list of scalability metrics identified on the basis of the work done within the definition of the use cases and pilots.

Finally, a consistency update of this deliverable has been submitted at 15th January 2021 by including the Food Protection use case, missing at the time of the second update at M18, and updating accordingly the list of datasets used within the BDG project and the list of scalability metrics to be used for each pilot.

## 7 REFERENCES

- Armstrong, T., G., Ponnekanti, V., Borthakur, D., Callaghan, M. (2013). Linkbench: a database benchmark based on the Facebook social graph. In Proceedings of the SIGMOD, ACM, 1185-1196.
- Bizer, C. and Schultz, A.(2008): Benchmarking the Performance of Storage Systems that expose SPARQL Endpoints. In: Proceedings of the 4th International Workshop on Scalable Semantic Web knowledge Base Systems, SSWS2008.
- Čermak, M., Tovarnak, D., Lastovicka, M. and Čeleda, P.(2016). A performance benchmark for NetFlow data analysis on distributed stream processing systems. NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium, Istanbul, pp. 919-924. doi: 10.1109/NOMS.2016.7502926.
- Chintapalli, S., Dagit, D., Evans, B., Farivar, R., Graves, T., Holderbaugh, M., Liu, Z., Nusbaum, K., Patil, K., Peng, B.J. and Poulosky, P.(2016). Benchmarking Streaming Computation Engines: Storm, Flink and Spark Streaming. IEEE International Parallel and Distributed Processing Symposium Workshop.
- Cooper, B.F., Silberstein, A., Tam, E., Ramakrishnan, R. and Sears, R.(2010). Benchmarking cloud serving systems with YCSB. In Proceedings of the 1st ACM symposium on Cloud computing ACM, pages 143-154.
- Dimitrov, M., Kumar, K., Lu, P., Viswiswanathan, V., Willhalm, T.(2013). Memory System Characterization of Big Data Workloads. IEEE International Conference on Big Data.
- Ferdman, M., Adileh, A., Kocberber, O., Volos, S., Alisafae, M., Jevdjic, D., Kaynak, C., Popescu, C., A.D., Ailamaki, A. and Falsafi, B.(2011). Clearing the Clouds: A Study of Emerging Workloads on Modern Hardware. In Proceedings of the 17th Conference on Architectural Support for Programming Languages and Operating Systems-ASPLOS 2012, pages 1-11.
- Fischer, F., Keller, U., Kiryakov, A., Huang, Z., Momtchev, V. and Simperl, E. (2008): Initial Knowledge Representation Formalism. LarKC project deliverable D1.1.3, [http://www.larkc.eu/wp-content/uploads/2008/01/larkc\\_d113-initial-knowledge-representation-formalism\\_m7.pdf](http://www.larkc.eu/wp-content/uploads/2008/01/larkc_d113-initial-knowledge-representation-formalism_m7.pdf).
- Ghazal, A., Rabl, T., Hu, M., Raab, F., Poess, M., Crolotte, A. and JACOBSEN, H.A. (2013). BigBench: Towards an Industry Standard Benchmark for Big Data Analytics. In SIGMOD ACM, New York, New York, 2013, pages 197-1208.
- GridMix. (2013). Available online: <https://hadoop.apache.org/docs/r1.2.1/gridmix.html>.
- Guo, Y., Pan, Z. and Heflin, J. (2004): An Evaluation of Knowledge Base Systems for Large OWL Datasets. Journal of Web Semantics, 3(2), 2005, pp. 158-182.
- Han, R., Zhen, J., Gao, W., Xinhui, T. and Wang, L. (2015). Benchmarking Big Data Systems: State-of-the-Art and Future Directions. TECHNICAL REPORT. ICT, ACS.
- Huang, S., Huang, J., Dai, J., Xie, T. and Huang B. (2010). The HiBench benchmark suite: Characterization of the MapReduce-based data analysis. In Data Engineering Workshops (ICDEW), 2010 IEEE 26th International Conference on IEEE, pages 41-51.
- Karimov, J., RABL, T., Katsifodimos, A., Samarev, R., Heiskanen, H., Markl, V. (2018). Benchmarking Distributed Stream Data Processing Systems. ICDE.
- Lilja, D. J.(2005). Measuring Computer Performance - A Practitioner's Guide. Cambridge University Press.
- Ming, Z., Luo, C., Gao, W., Han, R., Yang, Q., Wang, L. and Zhan, J. (2014). BDGS: A Scalable Big Data Generator Suite in Big Data Benchmarking. In CoRR.
- Ontotext Lab. (2009). RDF(S), Rules, and OWL Dialects. [http://www.ontotext.com/inference/rdfs\\_rules\\_owl.html](http://www.ontotext.com/inference/rdfs_rules_owl.html)

- Pavlo, A., Paulson, E., Rasin, A., Abadi, D. J., Dewitt, D. J., Madden, S. and Stonebraker, M. (2009). A comparison of approaches to large-scale data analysis. In Proceedings of the 2009 ACM SIGMOD International Conference on Management of data, SIGMOD '09, pages 165–178.
- Ren, Z., Xu, X., Wan, J., Shi, W., and Zhou, M. (2012). Workload characterization on a production Hadoop cluster: A case study on Taobao. In IEEE International Symposium on Workload Characterization, pages 3–13.
- Shukla, A., Chaturvedi, S., Simmhan, Y.(2017). RIoT Bench: An IoT benchmark for distributed stream processing systems. Concurrency Computat: Pract Exper.
- Shukla A. and Simmhan, Y. (2016). Benchmarking Distributed Stream Processing Platforms for IoT Applications. CoRR abs/1606.07621. arXiv:1606.07621 <http://arxiv.org/abs/1606.07621>.
- Tsai, C.-W., Lai, C.-F., Chao, H.-C., and Vasilakos, A. V. (2015). Big data analytics: a survey. In Journal of Big Data. 2:21. Springer.
- Wang, L., Zhan, J., Luo, C., Zhu, Y., Yang, Q., He, Y., Gao, W., Jia, Z., Shi, Y., Zhang, S., Zhen, C., Lu, G., Zhan, K. and Qiu., B. (2014). BigDataBench: a Big Data Benchmark Suite from Internet Services. In Proceedings of the The 20th IEEE International Symposium On High Performance Computer Architecture (HPCA 2014).
- Wierner, J. L., Kuno, H. and Graefe, G.(2009). Benchmarking query execution robustness. In TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC).
- Zhu, Y., Zhan, J., Weng, C., Nambiar, R., Zhang, J., Chen, X. and Wang, L.(2014). Generating comprehensive big data workloads as a benchmarking framework. The 19th International Conference on Database Systems for Advanced Applications (DASFAA 2014).
- Zhang, Y., Cao, T., Li, S., Tian, X., Yuan, L., Jia, H., and Vasilakos, A. V. (2016). Parallel Processing Systems for Big Data: A Survey. In Proceedings of the IEEE. Vol. 104, No.11, November 2016.