



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

D7.2 - Experimental Report on Current Datasets

DELIVERABLE NUMBER	D7.2
DELIVERABLE TITLE	Experimental Report on Current Datasets
RESPONSIBLE AUTHOR	Panagiotis Zervas (Agroknow)



GRANT AGREEMENT N.	780751
PROJECT ACRONYM	BigDataGrapes
PROJECT FULL NAME	Big Data to Enable Global Disruption of the Grapevine-powered industries
STARTING DATE (DUR.)	01/01/2018 (36 months)
ENDING DATE	31/12/2020
PROJECT WEBSITE	http://www.bigdatagrapes.eu/
COORDINATOR	Panagiotis Zervas
ADDRESS	110 Pentelis Str., Marousi, GR15126, Greece
REPLY TO	pzervas@agroknow.com
PHONE	+30 210 6897 905
EU PROJECT OFFICER	Ms. Annamária Nagy
WORKPACKAGE N. TITLE	WP7 Cross-sector Rigorous Experimental Testing
WORKPACKAGE LEADER	Agroknow
DELIVERABLE N. TITLE	D7.2 Experimental Report on Current Datasets
RESPONSIBLE AUTHOR	Panagiotis Zervas (Agroknow)
REPLY TO	pzervas@agroknow.com
DOCUMENT URL	http://www.bigdatagrapes.eu/
DATE OF DELIVERY (CONTRACTUAL)	30 June 2019 (M18)
DATE OF DELIVERY (SUBMITTED)	28 June 2019 (M18)
VERSION STATUS	1.0 Final
NATURE	R (Report)
DISSEMINATION LEVEL	PU (public)
AUTHORS (PARTNER)	Panagiotis Zervas (Agroknow), Mihalis Papakonstadinou (Agroknow)
CONTRIBUTORS	Franco Maria Nardini (CNR), Todor Primov (CNR), Nikola Rusinov (ONTOTEXT)
REVIEWER	Franco Maria Nardini (CNR)

VERSION	MODIFICATION(S)	DATE	AUTHOR(S)
0.1	Table of Contents	02/6/2019	Panagiotis Zervas (Agroknow), Mihalis Papakonstadinou (Agroknow)
0.4	Initial version	04/6/2019	Panagiotis Zervas (Agroknow), Mihalis Papakonstadinou (Agroknow)
0.6	Input from partners	06/6/2019	Franco Maria Nardini (CNR), Todor Primov (CNR), Nikola Rusinov (ONTOTEXT)
0.7	Internal Review	17/06/2019	Franco Maria Nardini (CNR)
0.9	Partners review, final comments and edits	24/6/2019	Franco Maria Nardini (CNR), Todor Primov (CNR), Nikola Rusinov (ONTOTEXT)
1.0	Final version	28/6/2019	Panagiotis Zervas (Agroknow), Mihalis Papakonstadinou (Agroknow)

PARTICIPANTS

CONTACT

Agroknow IKE
(Agroknow, Greece)



Panagiotis Zervas
Email: pzervas@agroknow.com

Ontotext AD
(ONTOTEXT, Bulgaria)



Todor Primov
Email: todor.primov@ontotext.com

Consiglio Nazionale Delle
Ricerche
(CNR, Italy)



Raffaele Perego
Email: raffaele.perego@isti.cnr.it

Katholieke Universiteit Leuven
(KULeuven, Belgium)



KatrienVerbert
Email: katrien.verbert@cs.kuleuven.be

Geocledian GmbH
(GEOCLEDIAN Germany)



Stefan Scherer
Email: stefan.scherer@geocledian.com

Institut National de la Recherché
Agronomique
(INRA, France)



Pascal Neveu
Email: pascal.neveu@inra.fr

Agricultural University of Athens
(AUA, Greece)



Katerina Biniari
Email: kbiniari@aua.gr

Abaco SpA
(ABACO, Italy)



Simone Parisi
Email: s.parisi@abacogroup.eu

SYMBEEOSIS LONG LIVE LIFE S.A.
(Symbeosis, Greece)



Symbeosis

Kostas Gardikis
Email: c.gardikis@gmail.com

ACRONYMS LIST

BDG	BigDataGrapes
API	Application Programming Interface
CPU	Central Processing Unit
AUA	Agricultural University of Athens
RDF	Resource Description Framework
UI	User Interface
JVM	Java Virtual Machine

EXECUTIVE SUMMARY

The deliverable D7.2 “Experimental Report on Current Datasets” consists of a report describing the outcomes of the experimentation performed on the datasets provided by each pilot, utilizing the Big Data Grapes stack.

This report starts by outlining the methodology and metrics we employ for our experimentation. We focus our experimentation on a pilot level, initially analysing the provided datasets and evaluating them against the Big Data Vs, as suggested and described in detail in D7.1.

We move on to identify and document the data flows each pilot adheres to throughout the Big Data Grapes stack, highlighting the frameworks and components involved in the process.

We experiment on each of the identified data flow steps by describing 3 usage scenario for each. During the execution of each scenario for each step, we monitor the chosen performance metrics, showing the respective diagrams and analysing the outcomes.

This deliverable also presents an end to end report for each of the pilot’s data flows, identifying bottlenecks and making suggestions to improve the performance where needed.

Table of contents

1 Introduction	10
2 Methodology	11
2.1 Steps	11
2.2 Metrics.....	11
3 Rigorous testing experiments	12
3.1 Farm management Pilot	12
3.1.1 Dataset Analysis & Evaluation.....	12
3.1.2 Data Flows	14
3.1.3 Experimentation.....	15
3.1.4 End to End Report.....	20
3.2 Natural Cosmetics Pilot	20
3.2.1 Dataset Analysis & Evaluation	21
3.2.2 Data Flows.....	22
3.2.3 Experimentation	22
3.2.4 End to End Report	29
3.3 Table and Wine Grapes Pilot	30
3.3.1 Dataset Analysis & Evaluation	30
3.3.2 Data Flows.....	36
3.3.3 Experimentation	36
3.3.4 End to End Report	42
3.4 Wine Making Pilot	42
3.4.1 Dataset Analysis & Evaluation	42
3.4.2 Data Flows.....	46
3.4.3 Experimentation.....	47
3.4.4 End to End Report.....	50
4 Conclusions.....	51

LIST OF FIGURES

Figure 1: Abstract overview of sensor data ingestion	15
Figure 2: Completion time of sensor data ingestion	16
Figure 3: CPU usage during sensor data ingestion	16
Figure 4: Network usage during sensor data ingestion	16
Figure 5: Memory usage during sensor data ingestion	16
Figure 6: Completion time of satellite image processing data ingestion.....	17
Figure 7: CPU usage during satellite image processing data ingestion.....	18
Figure 8: Network usage during satellite image processing data ingestion.....	18
Figure 9: Memory usage during satellite image processing data ingestion	18
Figure 10: Completion time of semantically enriched data ingestion	19
Figure 11: CPU usage during ingestion of semantically enriched data.....	19
Figure 12: Network usage during ingestion of semantically enriched data	19
Figure 13: Memory usage during ingestion of semantically enriched data.....	20
Figure 14: Completion time of dataset upload step	23
Figure 15: CPU usage during dataset upload	23
Figure 16: Network usage during dataset upload	24
Figure 17: Memory usage during dataset upload	24
Figure 18: Completion time of enriched data ingestion	25
Figure 19: CPU usage during ingestion of enriched data	25
Figure 20: Network usage during enriched data ingestion.....	25
Figure 21: Memory usage during enriched data ingestion	25
Figure 22: Completion time of satellite image processing data ingestion	26
Figure 23: CPU usage during satellite image processing data ingestion.....	27
Figure 24: Network usage during satellite image processing data ingestion	27
Figure 25: Memory usage during ingestion of satellite image processing data	27
Figure 26: Completion time of data correlation step	28
Figure 27: CPU usage during correlation of data.....	28
Figure 28: Network usage during correlation of data	28
Figure 29: Memory usage during correlation of data.....	29
Figure 30: Completion time for dataset upload	37
Figure 31: CPU usage during dataset upload	37
Figure 32: Network usage during dataset upload	37

Figure 33: Memory usage during dataset upload.....	38
Figure 34: Completion time for data streams pipeline.....	38
Figure 35: CPU usage during data streams pipeline	39
Figure 36: Network usage during data streams pipeline	39
Figure 37: Memory usage during data streams pipeline	39
Figure 38: Completion time for ingestion of enriched data.....	40
Figure 39: CPU usage during ingestion of enriched data	41
Figure 40: Network usage during ingestion of enriched data	41
Figure 41: Memory usage during ingestion of enriched data	41
Figure 42: Completion time for dataset upload.....	47
Figure 43: CPU usage during dataset upload.....	47
Figure 44: Network usage during dataset upload.....	48
Figure 45: Memory usage during dataset upload.....	48
Figure 46: Completion time for the ingestion of semantically enriched data	49
Figure 47: CPU usage during the ingestion of semantically enriched data.....	49
Figure 48: Network usage during the ingestion of semantically enriched data.....	49
Figure 49: Memory usage during ingestion of semantically enriched data.....	49

LIST OF TABLES

Table 1: Sentinel-2 satellite images dataset.....	12
Table 2: Landsat-8 satellite images dataset	12
Table 3: Satellite image processing dataset.....	13
Table 4: Environmental & Field indicators dataset.....	14
Table 5: Laboratory tests dataset.....	21
Table 6: Field metadata & Geographical information dataset.....	21
Table 7: Geographical & Field metadata dataset.....	30
Table 8: Soil quality dataset	31
Table 9: Vegetation indices estimation dataset	31
Table 10: Vegetation indices dataset.....	32
Table 11: Radioactive transfer & biophysical characteristics dataset	32
Table 12:Laboratory tests dataset	33
Table 13: IoT data.....	34
Table 14: Yield data dataset	34
Table 15: Qualitative data dataset	35

Table 16: Photosynthesis data dataset.....	35
Table 17: Experimentation results of rdfization & semantic enrichment.....	40
Table 18: Genetic data dataset.....	42
Table 19: Soil characteristics dataset.....	43
Table 20: Plot management dataset	43
Table 21: Climatic data dataset	44
Table 22: Laboratory analysis dataset	45
Table 23: Winemaking activities dataset.....	45
Table 24: Sensory analysis dataset	46
Table 25: Experimentation outcomes	51

1 INTRODUCTION

The BigDataGrapes (BDG) platform aims at targeting the needs of the grapevine industry using Big Data technologies, frameworks and components. In this report we perform an experimentation over the currently provided datasets by focusing on the needs of each use case as described in D2.1.

This report is structured as follows: first, we describe the methodology we use and the metrics we employ for the experimentation. Both of them are focused around the 4Vs of Big Data, as described in D7.1. We then present the description of each pilot's datasets and their respective data flows through the BDG stack. We perform this analysis by performing step by step and end to end tests for each dataset while monitoring the chosen metrics. Finally, we conclude this report by reporting the outcomes of this process, the identification of bottlenecks and heavy-duty tasks. We also provide suggestions on how to improve the performance of the whole BDG stack.

It should be noted all the experiments are performed on the working version of the BDG stack by using the currently deployed and usable components and tools. All of the Application Programming Interface (API) calls done throughout the experimentation support user authentication by API keys. The user authentication information is stored in a MySQL instance. All of the data flows are performed using Apache Nifi and the data uploads are performed either through API endpoints or using Filebeat from the Elastic stack, in case of internal data transfers.

2 METHODOLOGY

In this section we describe the methodology we use for the experimentation over the provided datasets, along with the steps we perform and the respective metrics we use for each of them.

As described in D7.1, we focus our experimentation around the 4Vs of Big Data: volume, velocity, variety and veracity. Based on the characteristics of each dataset, we choose to employ techniques and methods described in the bibliography (Deliverable D7.1, see BDGS and BigDataBench) to ensure the presence of every V for each of the datasets provided by the pilots. We focus our experimentation on the performance of the BDG stack by tracking several system indicators as described in section 2.2. Finally, we employ MetricBeat, i.e., the proposed tool for monitoring the scalability of BD platform reviewed in D7.1, over our Elastic stack to monitor and report our chosen metrics.

Since all of the BDG stack has been deployed in a microservice architecture, our step by step and end to end experimentation is done over the API endpoints provided by the platform. To make the experimentation easier we developed a python script that simulate these endpoint calls.

To showcase the potential of the deployed stack in terms of scalability we perform the experimentation using three usage scenarios for each step of the process by gradually increasing the requests made towards the platform.

2.1 STEPS

We choose to perform the experimentation on a pilot level. We identified a set of steps we follow for each of them. First, we identify and abstractly describe the provided datasets for each pilot. Then we move on with evaluating them against the Vs of Big Data, generating where applicable synthetic data series to cover any of the 4 Vs that are not covered by the provided datasets. We then identify the data flow each dataset will follow in the BDG stack, denoting the steps of this flow. Finally, using a python script that will simulate bursts of this data flows throughout the BDG stack, we will monitor and report our chosen metrics for this benchmark in real-time.

It should be noted that all of the experimentation is done using the actual provided data by each pilot and involve the already integrated and working components of the BDG stack.

2.2 METRICS

As mentioned in the previous sections, we will perform step by step and end to end experimentation for each pilot. In order for us to track and report the performance of the stack we chose to monitor the following indicators:

- Completion time, both on a step by step level, as well as on the whole end to end data flow,
- CPU (central processing unit) usage, we will track the CPU usage by each of the components as they are triggered by the flow,
- Memory usage of each of the components and technologies,
- Network usage in terms of bytes, we employ this metric, since the whole stack is based on a microservice architecture.

3 RIGOROUS TESTING EXPERIMENTS

3.1 FARM MANAGEMENT PILOT

The farm management pilot aims at developing a system that can support the farmer in his/her data collection, as well as in his/her day to day work. To that end, this specific pilot has granted access to datasets varying from satellite images and their respective processing to environmental and field indicators tracking by sensors deployed in the fields. In the following sections we analyze the provided datasets, following the methodologies described in D7.1. We identify the data flows of this specific use case inside the stack and experiment on each of them, using the metrics and the methodology described in the previous section.

3.1.1 Dataset Analysis & Evaluation

This pilot is characterized by datasets concerning satellite images from the external services of Sentinel-2 and Landsat-8 and the respective processing. All of this data is accessible through API endpoints provided by GEOCLIDEAN. It is also characterized by datasets concerning the environmental indicators and sensor data coming from deployed installations on the respective fields. In the tables below, we further describe these datasets and evaluate them against the 4Vs of Big Data.

Table 1: Sentinel-2 satellite images dataset

Dataset	Sentinel-2 Satellite Images	
Metadata/Description	This dataset contains the satellite images collected from Sentinel-2 for the fields provided by AUA, INRA, ABACO and Symbeeosis	
Provider	GEOCLIDEAN	
Sample	Sample of this data can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~2K scenes with a total size of ~1.5TB of data which is an acceptable value
	Velocity	The provided dataset is accessed through a REST service and shows the desired velocity since it is generated and ingested in real-time
	Variety	In terms of variety the provided dataset shows one type of data
	Veracity	There is no need for synthetic data generation

Table 2: Landsat-8 satellite images dataset

Dataset	Landsat-8 Satellite Images
Metadata/Description	This dataset contains the satellite images collected from Landsat-8 for the fields

	provided by AUA, INRA, oO and Symbeeosis	
Provider	GEOCLIDEAN	
Sample	Sample of this data can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~1.2K scenes with a total size of ~1.2TB of data which is an acceptable value
	Velocity	The provided dataset is accessed through a REST service and shows the desired velocity since it is generated and ingested in real-time
	Variety	In terms of variety the provided dataset shows one type of data
	Veracity	There is no need for synthetic data generation

Table 3: Satellite image processing dataset

Dataset	Satellite Image Processing Dataset	
Metadata/Description	This dataset contains the satellite image processing outcomes for the provided fields	
Provider	GEOCLIDEAN	
Sample	Sample of this data can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~100K of data points
	Velocity	The provided dataset is accessed through a REST service and shows the desired velocity since it is generated and ingested in real-time
	Variety	In terms of variety the provided dataset has 19 different data types along with metadata information & images, that are spread across 83 fields
	Veracity	Given the combination of the previous Vs, we have a moderate volume of data that shows high velocity, being produced and

	available in real-time and with an acceptable variety. Thus, we believe there is no need for synthetic data generation
--	--

Table 4: Environmental & Field indicators dataset

Dataset	Environmental & Field Indicators	
Metadata/Description	This dataset contains environmental & field indicators by sensors deployed in ABACO's field	
Provider	ABACO	
Sample	Sample of the data can be accessed here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~230K data points
	Velocity	The provided dataset is accessed through a REST service and shows the desired velocity since it is generated and ingested in real-time
	Variety	In terms of variety the provided dataset has 35 different data types
	Veracity	There is no need for synthetic data generation

Concluding the dataset analysis for this specific pilot, we observe that we have high volume datasets, having ~300K records that require a total amount of storage of ~2.7TB and are constantly growing. The provided datasets are all available through REST services generated and ingested in real-time and cover 56 different data types. Based on these observations we do not believe that there is a need for synthetic data generation and will move on with the experimentation with the provided datasets.

3.1.2 Data Flows

In this section we describe the steps that the provided datasets follow inside the BDG stack. As with all the use cases we will experiment on, we will focus on the ones that are already completed and integrated into the stack that serve this specific use case.

Initially, the data collected from the sensors deployed in the field are ingested into the platform. The next step involves the ingestion of the image processing data for each of the fields to be ingested. After this step, the data are converted into Resource Description Framework (RDF) and stored into the knowledge graph of the platform for their semantic enrichment. Finally, the data are extracted from the knowledge graph, and sent to the predictive analytics service for the analysis to happen. In its current state, this final step, utilizes several BDG

components. More specifically it uses Apache Hadoop, Hue and Spark. However, the datasets it currently uses are not the ones provided by this pilot, leading us to exclude this step from our experimentation.

3.1.3 Experimentation

In this section we perform the experimentation for each of the identified steps, using the datasets provided for this use case. To that end we use the ~230K sensor data coming from the field and the ~100K of satellite image processing outcomes coming from the GEOCLIDEAN service. As described in the methodology section we perform our experimentation using 3 scenarios for each step, to better showcase the scalability potential of the deployed stack.

3.1.3.1 Ingestion of Sensor Data

This step involves the ingestion of the sensor data, through the data stream component, Apache Kafka, of the BDG stack through a web service and stored into BDG's MongoDB for the metadata on the field, or initially into Apache Cassandra to take advantage of the high write throughput and then into Elasticsearch for the actual numerical values.

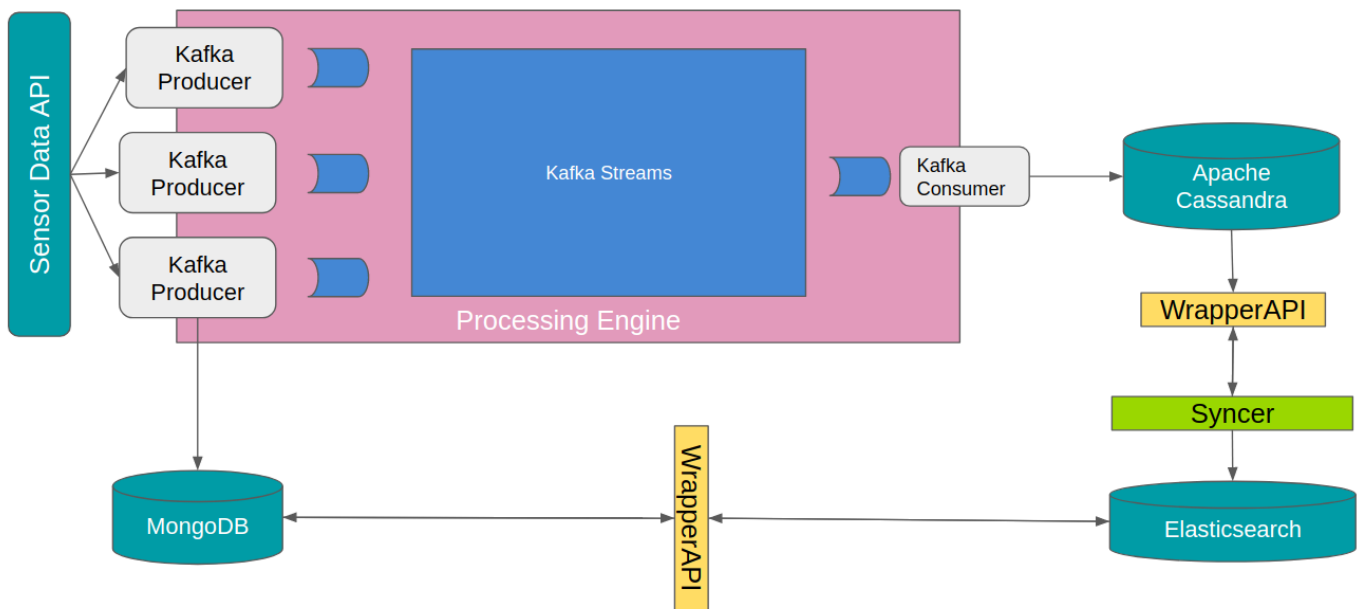


Figure 1: Abstract overview of sensor data ingestion

Figure 1 shows an overview of this specific process. As shown in the figure, the data is initially collected by Kafka producers and split between the metadata information, which is sent to MongoDB and actual data points that are inserted into Kafka topics to be processed by Kafka streams and collected by a Kafka consumer that sends the processed data to the Apache Cassandra. The last step of this step is the syncing between the data that reside in the Apache Cassandra instance and Elasticsearch.

To experiment on this step we identify the following usage scenarios, one with a Kafka setting of 1 producer/broker/consumer for all the data that reside into this dataset, one in which we will use a setting of 3 producers/brokers/consumers and a final in which we have a setting of 35, creating an 1-1 scenario for each data type to test the concurrency capabilities of the respective component.

Time vs. Kafka Brokers

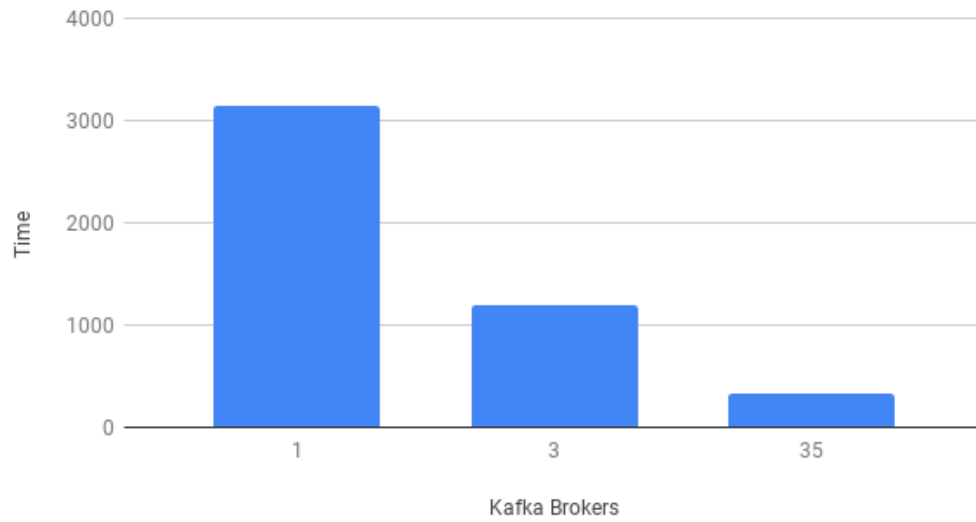


Figure 2: Completion time of sensor data ingestion




Figure 3: CPU usage during sensor data ingestion




Figure 4: Network usage during sensor data ingestion

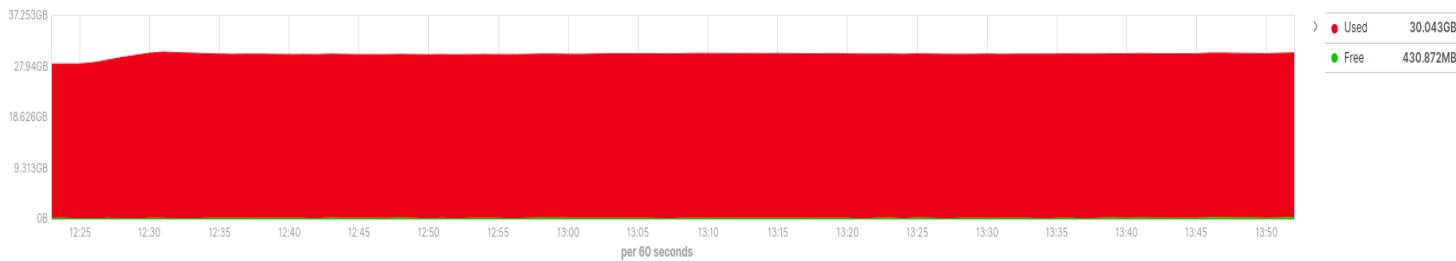


Figure 5: Memory usage during sensor data ingestion

As shown in the above figures, we observe the highest completion time for the ingestion scenario involving the setting of 1 producer/broker/consumer for Apache Kafka, taking 3,138 seconds to complete. During this scenario we have an average of 50% CPU usage for the whole stack, with a peak of 150%. The network traffic during this scenario shows an average of 500KB/s of inbound traffic. In the next setting of 2 producer/broker/consumers we observe on average the same CPU usage, without a spike as with the previous one. Figure 4 shows a much higher network traffic, on average 1.9MB/s and the whole scenario takes 1,192 seconds to complete. In our final experiment, in which we have one dedicated producer, broker and consumer for each of the observed indicators, we have a much lower completion time, 328 seconds, and on average the same CPU usage as in the previous scenario, 50%. In terms of network traffic, we have a small increase in comparison to the previous one, leading to an inbound traffic of 2MB/s. Throughout all scenarios we observe the same memory usage of the whole system with very little increases, easily explained by the fact that all of the involved components are Java based, set to reserve the required memory upon startup.

Following our experimentation, we conclude that the setting of 1 producer/broker/consumer for all of the data, is ineffective for this amount of data, since it takes a lot of time to be completed. We consider the final usage scenario as the best one with dedicated Apache Kafka pipelines for each indicator, since it takes little time to be completed without showing higher values in network traffic and CPU, as compared with the second one.

3.1.3.2 Ingestion of Satellite Image Processing Data

This step concerns the harvesting and storage of the satellite image processing dataset for each field through GEOCLIDEAN's service. In the following use cases we experiment on this specifically for each pilot's field, so we chose for this step to experiment on the overall ingestion process of this service, involving all of the 100K data points. As with all the steps, we perform our experimentation with 3 different usage scenarios, in the first we will ingest 50% of the data sequentially, all of the data are ingested in a sequential way, and finally we employ a scenario in which all of the data is ingested concurrently using 20 workers, each one responsible for the ingestion of the data for the two fields.




Figure 6: Completion time of satellite image processing data ingestion




Figure 7: CPU usage during satellite image processing data ingestion




Figure 8: Network usage during satellite image processing data ingestion




Figure 9: Memory usage during satellite image processing data ingestion

The figures illustrated above present the completion time, CPU usage, network traffic and memory usage of the whole stack, throughout our experimentation for this specific step. As expected, we observe the highest completion time for the second scenario where all of the data is ingested sequentially and a very high decrease in terms of completion time for the final one where 20 concurrent workers are employed. More specifically, 780 seconds are required for the first scenario to be completed, 1380 for the second and only 3.2 for the final one. In terms of CPU usage, each scenario shows on average the same usage, ~40-50%, with a spike of 400% for the first one and a spike of 100% for the second. The network traffic shows a steady increase for each step, starting at 300-350KB/s for the first, 450-500KB/s for the second and ~700KB/s for the final one. As in the previous step, the memory usage remains on the same levels.

Concluding our experimentation for this step, we observed that the best usage scenario is the third one, where 20 concurrent workers are employed. This leads us to believe that the best way to involve the ingestion of satellite images for the BDG stack is through frequent offline tasks with a high number of workers deployed in parallel.

3.1.3.3 Rdfization & Semantic Enrichment

In section 3.3.3.3 we thoroughly experiment on this step with higher volume datasets so do not conduct the same experiment here with the provided datasets.

3.1.3.4 Ingestion of Semantically Enriched Data

In this step we extract the semantically enriched data from GraphDB and store it in BDG’s Elasticsearch instance. The 3 scenarios we use for our experimentation are the following: in one we consider the ingestion of 1 day’s data, ~1,000 data points, in the next one month’s data, which is approximately 35K data points and in the final one we will attempt to extract and store all of the 330K of data points provided. All of the experiments will be performed using 3 concurrently working processes and in batches of 1,000 objects per request.

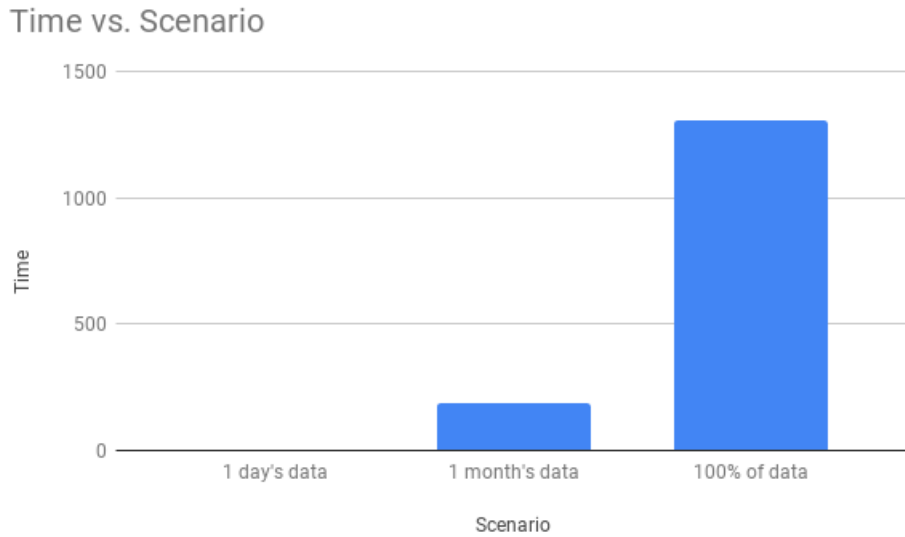


Figure 10: Completion time of semantically enriched data ingestion




Figure 11: CPU usage during ingestion of semantically enriched data




Figure 12: Network usage during ingestion of semantically enriched data




Figure 13: Memory usage during ingestion of semantically enriched data

Figure 10 illustrates the completion time for all of the chosen scenarios for this step. As we can see the ingestion of 1 day's data takes just under 2 seconds to complete, while 184 seconds are required in order to ingest 1 month's data. The ingestion of all of the data takes the longest as it takes 1307 seconds to complete. In terms of CPU the first 2 scenarios show on average the same usage (roughly 60-70%), while the last one consumes on average 100% of CPU with a peak of 260%. The network traffic shows the same behaviour for the last 2 scenarios with an average of 2-2.5MB/s while the first one has under 1MB/s, more specifically it consumes 800KB/s. Finally, the memory usage remains around the same levels throughout our experimentation.

As presented by our analysis for this step, we consider the best usage scenario to be the real-life one involving the ingestion of 1 day's data. To that end and taking into account the high network usage of the other 2 scenarios, we believe that the highest performance for this step can be achieved with the frequent extraction of smaller batches of data, happening on the background of the whole platform.

3.1.4 End to End Report

In conclusion to our experimentation for the data flows for this specific pilot we can see that overall the performance of the stack shows the necessary scalability. In particular the data pipeline step, involving Apache Kafka, MongoDB and Elasticsearch shows very high scalability when performed with high parallelization and low performance in terms of completion time if done otherwise. For the satellite image processing dataset ingestion, we consider that the best performance is also achieved when increasing its concurrency. However high network usage is observed in this case, which leads us to believe that further experimentation is needed to ensure that this observation does not create a problem when the stack is used by many concurrent users and pilots. The rdfization process shows a nice performance when performed using the command line tool. Finally, for the extraction of semantically enriched data, an excellent performance is observed in the real-life scenario. However, as the data volumes increase, high completion times are observed which leads us to believe that further experimentation is needed, along with possible changes in terms of employed technologies/frameworks for this specific step.

It should be noted at this point, as mentioned in the introductory section, that all of the described experimentation has been performed for the completed steps that are integrated into the BDG stack. Upon completion and integration of the rest, the experimentation will be updated to include the new results.

3.2 NATURAL COSMETICS PILOT

The natural cosmetics pilot focuses on the prediction of the biological efficacy of pharmaceutical plants. Currently this pilot has provided a dataset covering the lab experiments performed on the plants which is linked with satellite image processing of the respective fields in order for the correlation of lab tests and satellite images to be made possible.

3.2.1 Dataset Analysis & Evaluation

As mentioned in the introductory section, this pilot has provided 2 datasets. One is the lab tests performed on pharmaceutical plants and the other is the geographical information and related metadata of the fields the tests were performed against. Using the latter another dataset is also employed, that of the satellite image processing for these specific fields. In the tables below we describe these datasets and evaluate them against the 4Vs of Big Data.

Table 5: Laboratory tests dataset

Dataset	Laboratory Tests	
Metadata/Description	This dataset contains the laboratory tests performed on the pharmaceutical plants	
Provider	Symbeeosis	
Sample	A sample of this dataset can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset does not show a high value, leading to the need of the generation of synthetic data
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 10 different types of tests, which can be considered valid for our experimentation
	Veracity	In the context of veracity, we will follow the methodologies suggested in D7.1, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided. More specifically the methodology described here will be followed for the generation of synthetic data.

Table 6: Field metadata & Geographical information dataset

Dataset	Field Metadata & Geographical Information
Metadata/Description	This dataset contains the metadata of the fields the laboratory tests were performed in. It contains geographical information, grape variety cultivated and date the samples were collected

Provider	Symbeeosis	
Sample	A sample of this dataset can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset does not show a high value, leading to the need of the generation of synthetic data
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 6 different types of data values, which although is not a high value cannot be increased
	Veracity	In the context of veracity, due to the need of synthetic data generation, we will follow the statistical distribution of the raw data provided

This pilot is also using the dataset of satellite image processing which is described in section 3.1.1.

As described in this section the need for synthetic data generation has been identified. We perform this generation following the methodologies described in D7.1.

3.2.2 Data Flows

In this section we describe the workflows the Natural Cosmetics pilot follows inside the BDG stack, during which the experimentation will be performed. It should be noted that the identified steps are the currently completed and integrated ones into the BDG stack that support this specific use case.

First, the dataset is uploaded into the system, either by using the provided UI (User Interface) per pilot, or by using the specific API endpoint as described in D6.2. Second, the preprocessed dataset is sent over to the Rdfization service for the RDF generation and then it is imported into the GraphDB storage engine for the semantic enrichment of the provided data. The next step of this pilot is to export the enriched data from GraphDB. The exported semantically enriched data is then inserted into the Elasticsearch instance along with the satellite image processing data that cover this pilot's fields. Finally, all of the data that resides into Elasticsearch is exported and used for the correlation of the provided values by querying the specific API endpoint.

3.2.3 Experimentation

In this section we perform the rigorous testing experimentation for each of the identified steps, using the synthetic dataset generated. The dataset that was generated contains 3M laboratory tests that cover 200K fields and all of the generated records follow the statistical distribution of the initial one provided by the pilot. As

described in the methodology section we will perform our experimentation using 3 scenarios for each step, to better showcase the scalability potential of the deployed stack.

3.2.3.1 Dataset Upload

To experiment on this step, we choose a scenario that involves high amounts of data, along with high concurrency in the requests made towards the respective API endpoint. For the needs of this specific experiment we assume as usage scenario that under which the dataset upload requests involve 10,000 entities and the concurrency of them is set to 1,000, 5,000 and 10,000 respectively for each testing scenario.

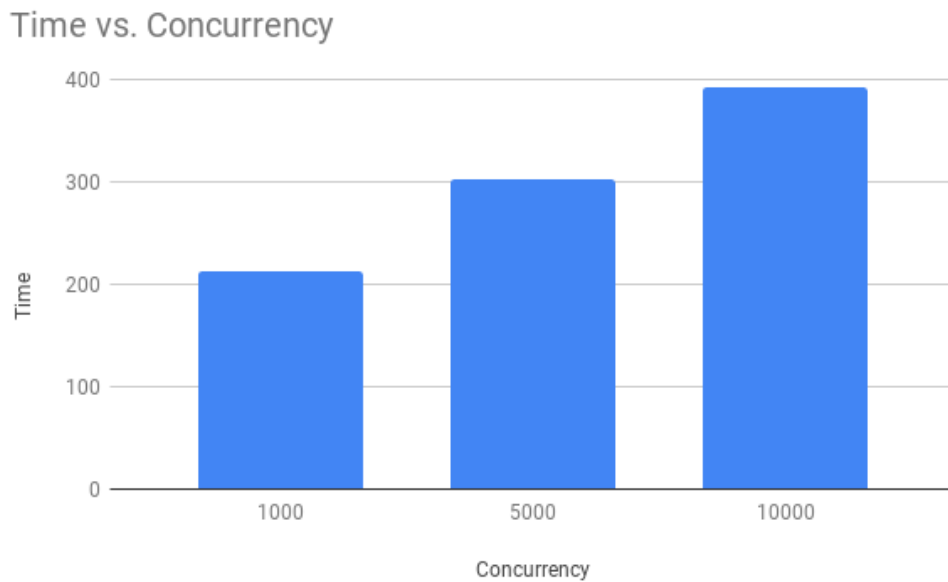


Figure 14: Completion time of dataset upload step

Figure 14 shows the completion time against the concurrency of the requests to the stack. As it is demonstrated in this figure, completion time shows a high increase when moving from 1000 to 5000 concurrent requests, whereas when moving to 10000 concurrent requests the completion time although increased has a lower increase rate.




Figure 15: CPU usage during dataset upload




Figure 16: Network usage during dataset upload




Figure 17: Memory usage during dataset upload

The figures shown above depict the CPU usage, network traffic in bytes and memory used by the stack, throughout the execution of the three scenarios. We observe a slight increase in terms of CPU usage between the scenarios starting from an average of 70% for the one with 1000 concurrent requests, 75% for the next one and an average of 80-85% for the last scenario with the 10000 concurrent requests, with a peak of 120%. In terms of network traffic, we observe a steady rate of inbound traffic for the first 2 scenarios and a high increase for the last one leading to 800KB/s, the outbound traffic shows a linear increase for the three scenarios with an average of 300KB/s, 700KB/s and 1.2MB/s respectively for each. Finally, a stable memory usage is observed which is due to the fact that most of the components in the stack are based on the JVM (Java Virtual Machine) which can be set to reserve the required memory upon startup.

We conclude our experimentation on this step by observing that the best usage of the platform taking into account all of the monitored metrics is the first one, where small amounts of data are ingested into the stack. To that end we change the initial ingestion of data for this pilot to have its datasets uploaded in small batches of concurrent requests.

3.2.3.2 Rdfization & Semantic Enrichment

In section 3.3.3.3 we thoroughly experiment on this step with higher volume datasets so do not conduct the same experiment here with the provided datasets.

3.2.3.3 Ingestion of Enriched Data

In this step we experiment with the ingestion of the semantically enriched data for the Natural Cosmetics pilot. To fully investigate this step, we identify 3 usage scenarios, based on the data we will extract from GraphDB and ingest into Elasticsearch. To this end, the first scenario involves the ingestion of 100K of the laboratory test, the second one the ingestion of 1M laboratory test and the final one the ingestion of 3M lab tests. The ingestion of the laboratory tests is done in batches of 1000.

Time vs. #ofRecords




Figure 18: Completion time of enriched data ingestion




Figure 19: CPU usage during ingestion of enriched data




Figure 20: Network usage during enriched data ingestion




Figure 21: Memory usage during enriched data ingestion

As demonstrated in the figures above, we observe a steady increase in terms of completion time for each scenario, easily explained by the increase of the total number of records required for each scenario. Interestingly though the completion time does not show a linear increase with respect to the number of records, since in order to ingest all of the synthetically generated data for this pilot (3M records), we observe a roughly 2 times increase in terms of completion time compared to the second scenario where 1M records are ingested. In terms of CPU usage, as shown in the respective figure, we observe a steady usage of an average of 50%, spiking to 150% and 520% for the second and third scenario respectively. The network traffic also shows a steady increase throughout the employed scenarios, starting from 800KB/s for the first scenario, 1.5-2MB/s for the second and 2.5MB/s for the last one. As shown in Figure 21, the memory usage remains steady, as is the case with the previous steps of all the pilots.

In conclusion of the experimentation with the current step, we observe a close to linear increase in terms of completion time as the volume of the ingested data increases. However, due to the high network traffic generated in the last 2 scenarios, we believe that this ingestion should be triggered once a record is inserted into GraphDB and the enrichment is done, or in smaller batches of record, so that it does not affect other components running at the same time.

3.2.3.4 Ingestion of Satellite Image Processing Data

We move on with the experimentation over the ingestion of the Natural Cosmetics related dataset that comes through GEOCLIDEAN'S API, containing the satellite image processing data. For this step we will follow an approach similar to the previous one. To that end, since the actual supplied fields by Symbeeosis are 13, we experiment on this step in the following 3 scenarios: one under which the ingestion of 1 field's processed image data occurs, one in which 7 fields are ingested concurrently and finally one where all of the supplied fields are ingested at the same time.

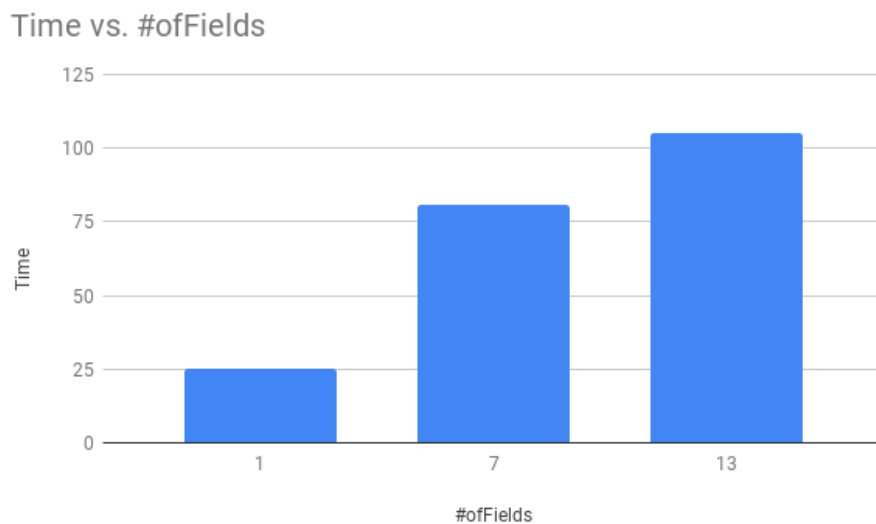


Figure 22: Completion time of satellite image processing data ingestion




Figure 23: CPU usage during satellite image processing data ingestion




Figure 24: Network usage during satellite image processing data ingestion

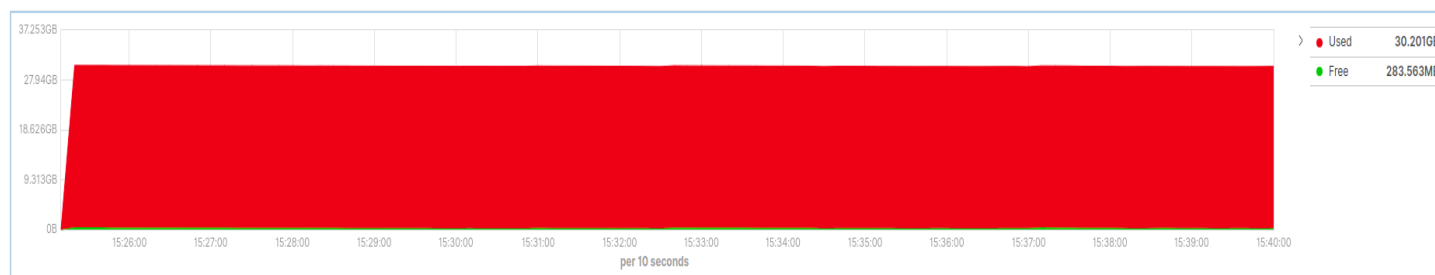


Figure 25: Memory usage during ingestion of satellite image processing data

As shown in the figures above, we observe a low completion time for each step, varying from 25 seconds to 110 seconds. In terms of CPU usage, a low value is observed, on average 40% for the first two scenarios and a 80-90% usage for the last one where all of the image processing data for every field is ingested. Finally, in terms of network usage we observe a steady value of average 500KB/s for each of the scenarios employed and a steady memory usage for the whole stack.

As further described in section 3.1.3.2, where we experimented with the ingestion of all of the data that reside in GEOCLIDEAN’s service, we consider this step to have a very high efficiency value considering the metrics we monitor.

3.2.3.5 Correlation of Data

The final step for this pilot’s datasets is the correlation one, in which the semantically enriched laboratory tests will be correlated with the results of the satellite image processing. We identify the following usage scenarios for this step: in the first one we will call the correlation API with 230K samples from the lab tests, which is on average the data from one field, in the second one we will move on with the usage of 50% of the fields data, a total of 1.5M records. In the final usage scenario, we will employ all of the 3M synthetically generated data series. It should be noted that in all scenarios covering more than one field, we evenly distribute the lab tests across all of the 13 supplied fields.

Time vs.#ofRecords

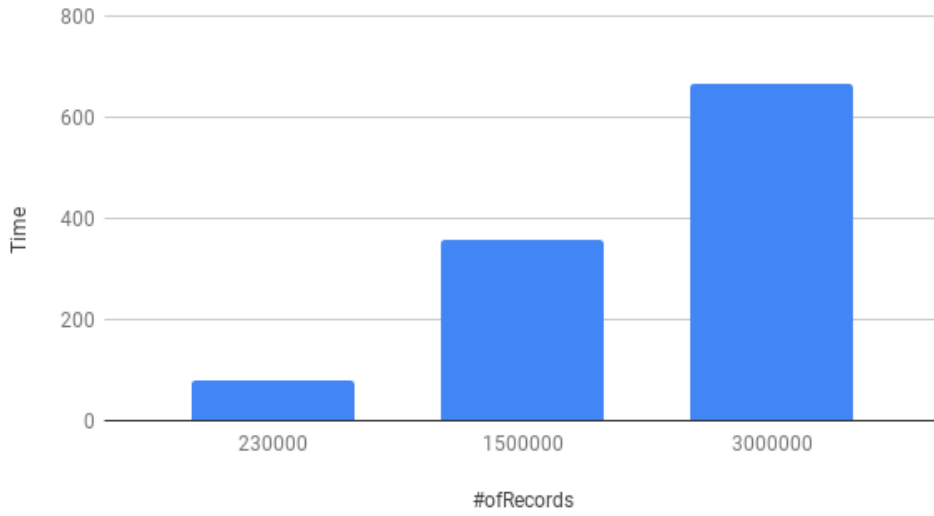


Figure 26: Completion time of data correlation step




Figure 27: CPU usage during correlation of data




Figure 28: Network usage during correlation of data



Figure 29: Memory usage during correlation of data

As shown in figure 26 the completion time increases as we supply more data to the correlation script. More specifically, 80 seconds are required for the completion of the first scenario, 356 for the second and 667 for the scenario under which all of the synthetically generated data are provided to the script. Interestingly, regardless of the size of the provided datasets, the CPU usage remains at the same levels across all experimentations, roughly 140%. This can be explained by the fact that the correlation script is configured to run in a single thread, regardless of the dataset's size. In terms of network traffic, we observe a steady increase as we progress through the scenarios starting from 300KB/s and leading to 650KB/s. This fluctuation can be explained by the fact that each correlation script execution results in a zip file generation and download by the respective API's consumer. Finally, in terms of memory we see that a steady value is observed, with the exception of an initial cache clear performed by the platform, resulting in an increase in the memory usage.

In conclusion to our experimentation for this step, we consider the last usage scenario as the real-life one. Correlation workflows are heavily dependent on the volume of the input datasets so the last usage scenario best simulates the desired production state of this step. drawback of this step is the execution time required for the actual correlation to happen. To ameliorate this, we consider splitting this step into 2 distinct substeps, one in which the correlation of the provided data happens using offline cron jobs and another that returns the produced histograms in real-time.

3.2.4 End to End Report

In conclusion to our experimentation for the Natural Cosmetics pilot and the data flows its dataset have inside the BDG stack, we consider the overall performance of the distinct steps as very successful. In the context of the data uploading step we observed that this step is an easily scalable one, since even with 10K concurrent requests the respective components did not show any downtime. However due to the high network usage as the concurrency increases, we note as an upper bound for the performance of the stack that of 5000-7000 concurrent requests at most. The rdfization step also shows very good performance when executed using the command line tool, a tool triggered by cron jobs installed in the platform. The extraction of the semantically enriched data is the one presenting a bottleneck for this pilot (as was the case for the previous). This leads us to believe that to achieve the best performance for this step, the exported data should either be split into smaller batches or further experimentation employing different components of the stack should be investigated. In the context of the satellite image processing data, we consider this step as a highly performant one, since its completion time and the monitored metrics show very good values for the provided fields of this pilot. Finally, the correlation of data, also shows very good performance, taking into consideration the nature of this specific step. In the future we plan on further experimenting on this step by increasing the concurrency of the requests sent to it to better showcase its scalability.

It should be noted at this point, as mentioned in the introductory section, that all of the described experimentation has been performed for the completed steps that are integrated into the BDG stack. Furthermore, due to the lack of Big Data coverage for the provided datasets, the experimentation was performed with high amounts of synthetic data, which although follow the statistical distribution of the provided may decrease the accuracy of the results.

3.3 TABLE AND WINE GRAPES PILOT

The table and wine grapes pilot focuses on the prediction of the quality of the final product of table and wine grapes. This is achieved through the provided datasets of this pilot; datasets that involve soil and weather data, as well as crop qualitative and quantitative data. In the sections that follow, we analyze the provided datasets against Big Data 4Vs, we identify the flows these datasets follow inside the BDG stack and we perform our experimentation on each using the methodology described in the introductory section.

3.3.1 Dataset Analysis & Evaluation

As mentioned in the introductory section, this pilot has provided 7 different datasets that cover 3 different fields, resulting in a total of 22 datasets, since one of the fields is lacking an IoT installation. In the tables below we describe these datasets and evaluate them against the 4Vs of Big Data.

Table 7: Geographical & Field metadata dataset

Dataset	Geographical & Field Metadata	
Metadata/Description	This dataset contains geographical information for each field along with the grape variety cultivated	
Provider	AUA	
Sample	Samples for this dataset can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided datasets cover 3 fields which can be further expanded into 12 considering the different varieties cultivated at each field, technically making it a different field.
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 5 different data types which is not a high variety value
	Veracity	In the context of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical

	distribution of the raw data provided
--	---------------------------------------

Table 8: Soil quality dataset

Dataset	EM38	
Metadata/Description	This dataset contains soil quality data	
Provider	AUA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided datasets have a total of 67K measurements, which if associated with the different data types lead to a dataset that has satisfying volume
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 10 different data types
	Veracity	In terms of veracity, there is no need for synthetic data generation

Table 9: Vegetation indices estimation dataset

Dataset	RapidScan	
Metadata/Description	This dataset contains data used to estimate vegetation indices such as NDVI and NDRE indices	
Provider	AUA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided datasets on RapidScan have a total of 400 rows, resulting to a need of synthetic data generation

	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 22 different data types, which is considered valid for our experimentation
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 10: Vegetation indices dataset

Dataset	SpectroSense	
Metadata/Description	This dataset contains data used to estimate LAI (Leaf Area Index) and NDVI vegetation indices	
Provider	AUA	
Sample	Samples for these datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the 17 provided datasets on RapidScan have a total of ~10.5K rows
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 12 different data types, which is considered valid for our experimentation
	Veracity	In terms of veracity, there is no need for synthetic data generation

Table 11: Radioactive transfer & biophysical characteristics dataset

Dataset	CropCircle
Metadata/Description	This dataset contains data used to estimate the radioactive transfer and the biophysical characteristics of plant canopies.

Provider	AUA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~100K rows
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset has 6 data types, including geographical information
	Veracity	In terms of veracity, there is no need for synthetic data generation

Table 12: Laboratory tests dataset

Dataset	Laboratory Tests	
Metadata/Description	This dataset contains laboratory tests performed on samples from the fields	
Provider	AUA	
Sample	Samples for these datasets can be found here , and here	
Evaluation	Big Data Vs	
	Volume	In terms of volume this dataset has ~1200 rows, which is not considered a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However for the needs of our experimentation we can mock the desired velocity
	Variety	The provided dataset has 9 different data types
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 13: IoT data

Dataset	IoT data	
Metadata/Description	This dataset contains data collected from IoT installations on 2 different fields	
Provider	AUA	
Sample	Samples for these datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided API endpoints show a total of 260K of data points that is constantly growing
	Velocity	The provided datasets are API endpoints that have bursts of new data every 5 minutes
	Variety	In terms of variety the provided dataset tracks 5 different data types, which is not a high value
	Veracity	In terms of veracity, there is no need for synthetic data generation due to the combination of volume and variety

Table 14: Yield data dataset

Dataset	Yield data	
Metadata/Description	This dataset contains data on the yield of a specific field	
Provider	AUA	
Sample	Samples for this datasets can be found here , and here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has 100 rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 4 different data types, along with metadata information on the field, which is not a high value

	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided
--	----------	--

Table 15: Qualitative data dataset

Dataset	Qualitative data	
Metadata/Description	This dataset contains data on the quality characters of the grapes of a specific field	
Provider	AUA	
Sample	Samples for these datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~50 rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 46 different observations
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 16: Photosynthesis data dataset

Dataset	Photosynthesis data	
Metadata/Description	This dataset contains data on the photosynthesis measurements taken from selected plants of a specific field.	
Provider	AUA	
Sample	Samples for these datasets can be found here	
Evaluation		

Big Data Vs	
Volume	In terms of volume the provided dataset has 250 rows, which is not a high value
Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
Variety	In terms of variety the provided dataset tracks 65 different observations
Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

This pilot is also using the dataset of satellite image processing which is described in section 3.1.1.

Due to the high variety observed across all of the AUA provided datasets and taking into account the high volume of most of them, we chose not to generate synthetic data for this use case and focus our experimentation on the provided ones.

3.3.2 Data Flows

In this section we describe the workflows the AUA pilot follows inside the BDG stack, during which the experimentation will be performed. It should be noted that the identified steps are the currently completed and integrated ones into the BDG stack that support this specific use case.

Initially, the datasets are uploaded into the system, either be the provided UI per pilot, or by the respective API endpoint as described in D6.2. Depending on whether it is metadata information of actual data, we will insert them into MongoDB and Elasticsearch respectively. The data provided by the IoT installations follow the data stream pipeline of the stack that involves their ingestion through Apache Kafka. All of the ingested data is initially stored on BDG's Elasticsearch instance. The cleaned and preprocessed data is then send over to the Rdfization service for the RDF generation and import into the GraphDB storage engine for the semantic enrichment. The next step for this pilot is to export the enriched data from GraphDB and the ingestion of it along with the dataset of the satellite image processing into the Elasticsearch instance. Finally, for this specific pilot's needs there are 2 different correlations that are expected: one that correlates the lab tests with the results from the satellite image processing and another that attempts to correlate the yield of a field with the respective lab tests and environmental indicators of the provided datasets, the last 2 however are not currently implemented and thus not taken into account in this experimentation.

3.3.3 Experimentation

In this section we perform the rigorous testing experimentation for each of the identified steps, using the data provided by AUA. As mentioned in the dataset evaluation section we do not generate synthetic data for this specific use case and we experiment with the provided datasets that show a total of ~440K records spread

across files and API endpoints. As described in the methodology section we perform our experimentation using 3 scenarios for each step, to better showcase the scalability potential of the deployed stack.

3.3.3.1 Dataset Upload

Similarly, to the approach followed for the other use cases, we experiment on the dataset upload step using the provided datasets. These show a volume of 180,000 records with a variety of 68 data types, spread across 7 datasets. We initially experiment on this step, uploading 50% of the provided data sequentially, we move on by attempting to upload all of the data sequentially and conclude this subsection by uploading all of the provided data with 22 concurrent requests, as the number of the datasets.

Time vs. Scenario

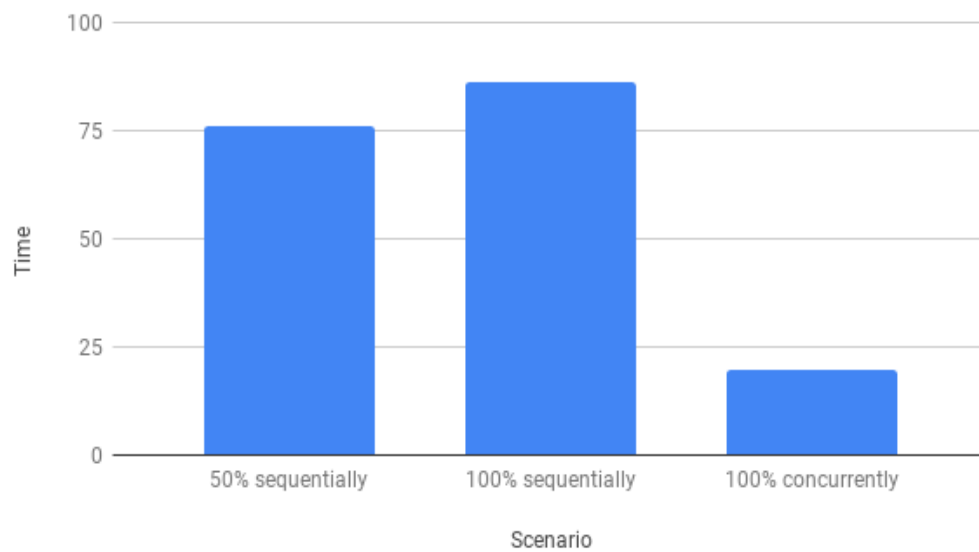


Figure 30: Completion time for dataset upload




Figure 31: CPU usage during dataset upload

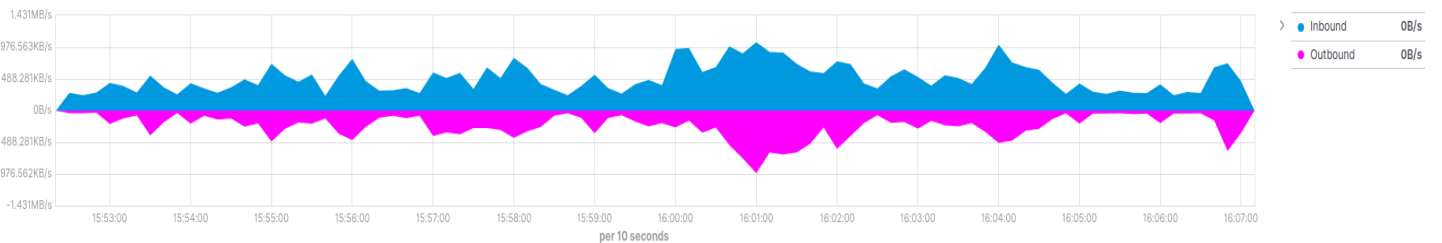


Figure 32: Network usage during dataset upload




Figure 33: Memory usage during dataset upload

As shown in the figures above we observe a low increase ingesting 50% or all of the data sequentially and a much lower value when we attempted ingesting all of the data concurrently. We also observe a steady CPU usage for both the sequential scenarios of 80% and a peak of 180% for the last fully parallelized one. In terms of network usage, we have a steady usage for the first 2 scenarios, an average of 400-500KB/s is observed and a slight increase for the last one of just above 500KB/s. Finally, the memory usage shows a very steady value throughout our experimentation.

In conclusion to this step's experimentation, we have observed that the fully parallelized experiment shows the best performance in consideration with the metrics we monitor, leading to us fully parallelizing this specific step.

3.3.3.2 Data streams Pipeline

To experiment on this step of the data flows for the AUA (Agricultural University of Athens) pilot, we will employ the following 3 usage scenarios: one under which we will ingest the latest data as they come, which is the real-life scenario, one under which we will ingest one's day data from the 2 available sensor APIs and a final one where we will harvest all of the data existent into the APIs and store them into the BDG stack. The metrics documented for this step involves the whole process of calling the sensor API to get the data, streaming them and performing further calculations on the data streams using Kafka and Kafka streams, modeling them into the internal BDG format and inserting them initially into Apache Cassandra to take advantage of the high write throughput and then into our Elasticsearch instance.

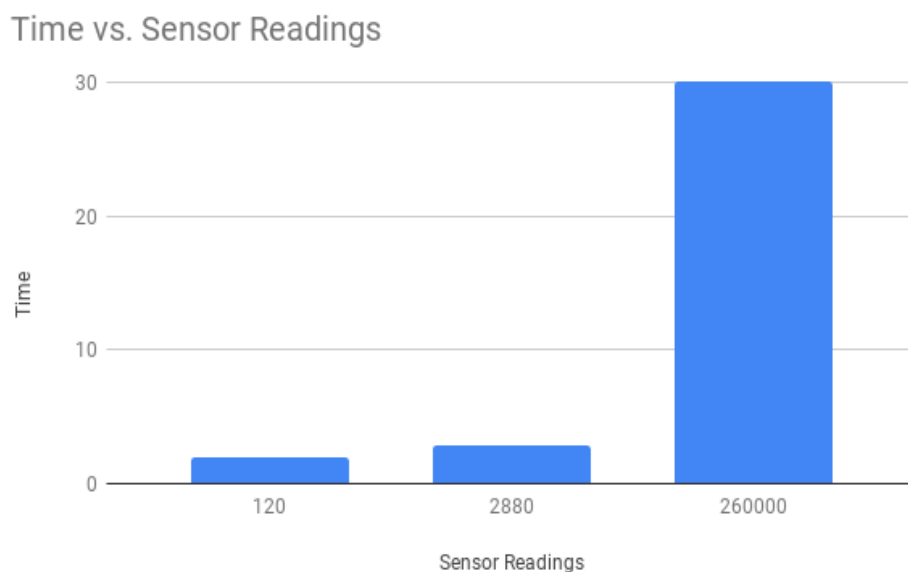


Figure 34: Completion time for data streams pipeline




Figure 35: CPU usage during data streams pipeline




Figure 36: Network usage during data streams pipeline




Figure 37: Memory usage during data streams pipeline

During our experimentation with the employed 3 scenarios for this step, we observe very low completion time for the first scenario where the latest data is ingested, 2 seconds, a minor increase in the case of the ingestion of 1 day's data, just below 3 seconds and an increase when attempting to ingest all of the currently generated data. The CPU usage of the whole stack shows a steady value for the all of the scenarios with an average of 40-45% and a spike of 300% for the last one. In terms of network traffic, we observe high fluctuations varying between 600KB/s for the first scenario, 2.3MB/s for the second and 9MB/s in the last scenario. Finally, the memory usage remains around the same value throughout our experimentation.

In conclusion of the experimentation for this step, we consider the usage of BDG's components involved as very successful since the real-life scenario, which is the currently deployed one, shows a very high performance in respect to the monitored metrics.

3.3.3.3 Rdfization & Semantic Enrichment

RDFi-zation is the process of converting tabular data to RDF. As described in D6.2, Section 2.2.4.1, we use the TARQL tool to perform this task. TARQL is built from the ground up to handle heavy loads and work on streaming data. In order to illustrate this, we ran the following test on the most voluminous dataset in the Table and Wines pilot - the IoT-stationary data (available [here](#)). This data consists of continuous meteorological observations on a given plot. The data is collected every 5 minutes, thus, one month of data corresponds to roughly 9000 observations, and 1 year to 110000 observations. The sensor reports 11 different values.

In order to illustrate the performance of the RDF-ization pipeline we ran an experiment using 3 datasets:

- 1 month of data, corresponding to data for August 2018 - 8839 observations
- the full Palivou IoT data, corresponding to data for months 08-12 2018 - 32095 observations
- A dataset corresponding to 6 years of (simulated) data - 609787 observations

Table 18 summarizes the results in terms of runtime¹ and number of resulting RDF triples.

Table 17: Experimentation results of rdfization & semantic enrichment

dataset	period	observations	runtime	resulting triples
palivou-08	1 month	8839	3,70s	141494
palivou-iot	5 months	32095	5,40s	513579
palivou-6y	6 years	609787	34,65s	9756453

As we can see, in the simulated 6y dataset the tool took 34 seconds to produce close to 10million triples from 6 years' worth of dense meteorological data thus illustrating that its performance is capable of handling large amounts of data in manageable time.

3.3.3.4 Ingestion of Enriched Data

In this step we experiment with the ingestion of the semantically enriched data for the AUA pilot. To fully investigate this step, we identify 3 usage scenarios, based on the data we will extract from GraphDB and ingest into Elasticsearch. To this end, the first scenario involves the ingestion of 10K rows of data, the second one the ingestion of 200K and the final one the ingestion of all 440K data points provided by this pilot. The ingestion of is done in batches of 1000. In the first 2 scenarios we evenly distribute the samples which we will ingest across all of the provided datasets by the pilot.




Figure 38: Completion time for ingestion of enriched data

¹ On a modern workstation with an intel i7 processor




Figure 39: CPU usage during ingestion of enriched data




Figure 40: Network usage during ingestion of enriched data




Figure 41: Memory usage during ingestion of enriched data

Following our experimentation for this step, we observe that the completion time for each scenario shows a steady increase, starting from 16 seconds for the ingestion of 10K records, 420 seconds are required to extract and store 200K of enriched data and for the final scenario, the ingestion of all of the provided data takes 613 seconds to complete. In terms of usage, we observe the same average usage throughout the scenarios, 50%, with different maximum values for each, 60%, 110% and 200% respectively. In terms of network traffic, we are seeing 1MB/s for the first scenario, increasing to 3.5MB/s for the second, leading to a 6.5MB/s for the last. Finally, the memory usage remains on the same values throughout our experimentation.

Concluding our experimentation for this step, we consider the first scenario as the most promising, since it has the lowest values on all metrics monitored, while having an acceptable network traffic. To that end, as was the case with other use cases for this specific step, we consider the best practise to be having offline tasks frequently querying for new data into the GraphDB engine, pushing changes into the Elastic stack.

3.3.3.5 Ingestion of Satellite Image Processing Data

This pilot covers 3 different fields on which the satellite image processing data are required. However, since this step is identical to the one required by the one described in 3.2.3.4 for the Natural Cosmetics pilot but requires less fields, we believe that the testing reported in that section clearly overlaps with this one, with the previous experimentation being on a higher scale.

3.3.4 End to End Report

In this section we conducted a thorough experimentation on the steps the data provided by the Table and Wine Grapes pilot follow inside the BDG stack. The initial step of the dataset upload shows excellent performance and scalability regardless of the increase in concurrency. Since the CPU and network usage show minor increases as we increase the concurrency, we consider this step as a highly performant one. The data pipeline step also shows very good performance as far as the real-life scenario is concerned. Following the conclusions, we also came up for the Farm Management pilot higher concurrency should be employed for this step to overcome high execution time as the volume of data increases. In terms of rdfizing the data, as we observed for the previous cases as well, the steps shows good performance using the respective command line tool. The extraction of the semantic enriched data is the one showing the poorest performance for this specific pilot. As we have also described in the previous cases, we consider the increase in terms of concurrency to help greatly in improving the performance of this step so that no bottleneck is observed. Finally, the extraction and storage of the satellite image processing dataset demonstrates great performance, keeping under consideration the low number of fields required to cover the pilot’s needs.

It should be noted at this point, as mentioned in the introductory section, that all of the described experimentation has been performed for the completed steps that are integrated into the BDG stack. Upon completion and integration of the rest, the experimentation will be updated to include the new additions.

3.4 WINE MAKING PILOT

This pilot deals with research in the fields of viticulture and oenology with a focus on the quality of the final product and its association with indicators collected from the field and the laboratory. The provided datasets include genetic data, weather and sensor data as well as laboratory experiments and field management activities. In the sections that follow, we analyze the provided datasets and heavily experiment on each of the identified data flow steps for this specific use case.

3.4.1 Dataset Analysis & Evaluation

As mentioned in the introductory section, this pilot has provided 8 different datasets. In the tables below we describe these datasets and evaluate them against the 4Vs of Big Data.

Table 18: Genetic data dataset

Dataset	Genetic data	
Metadata/Description	This dataset contains genetic data	
Provider	INRA	
Sample	Samples for these datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~2.6K rows, which is not a high value

	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 7 different data types, including metadata descriptions
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 19: Soil characteristics dataset

Dataset	Soil Characteristics	
Metadata/Description	This dataset contains data on soil characteristics	
Provider	INRA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~100 rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 13 different data types, along with metadata information
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 20: Plot management dataset

Dataset	Plot Management	
Metadata/Description	This dataset contains data on plot management	
Provider	INRA	

Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~400 rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 32 different data types, along with metadata information
	Veracity	In terms of veracity, we will follow the methodologies suggested in D7.1, BDGS and BigDataBench, for the generation of large amounts of synthetic data that follow the statistical distribution of the raw data provided

Table 21: Climatic data dataset

Dataset	Climatic Data	
Metadata/Description	This dataset contains climatic data	
Provider	INRA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~2.5K rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 58 different data types, along with metadata information
	Veracity	In terms of veracity, considering the high variety of this dataset we do not see the need for synthetic data generation

Table 22: Laboratory analysis dataset

Dataset	Laboratory Analysis	
Metadata/Description	This dataset contains the results of laboratory analysis	
Provider	INRA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~1.6K rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 95 different data types, along with metadata information, spread across 4 datasets
	Veracity	

Table 23: Winemaking activities dataset

Dataset	Winemaking Activities	
Metadata/Description	This dataset contains data on winemaking activities	
Provider	INRA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~550 rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity

	Variety	In terms of variety the provided dataset tracks 59 different data types, along with metadata information, spread across 4 datasets
	Veracity	In terms of veracity, considering the high variety of this dataset we do not see the need for synthetic data generation

Table 24: Sensory analysis dataset

Dataset	Sensory Analysis	
Metadata/Description	This dataset contains data from sensory analysis	
Provider	INRA	
Sample	Samples for this datasets can be found here	
Evaluation	Big Data Vs	
	Volume	In terms of volume the provided dataset has ~3.6K rows, which is not a high value
	Velocity	The provided dataset is a static file and does not show a high velocity rate. However, for the needs of our experimentation we can mock the desired velocity
	Variety	In terms of variety the provided dataset tracks 12 different data types, along with metadata information, spread across 2 datasets
	Veracity	In terms of veracity, considering the combination of volume and variety we do not see a high need for synthetic data generation

On an overall analysis of the provided data we observe a very high variety value, having a total of 276 different data types, spread across 14 datasets. The volume of the provided data is lacking, having a total of ~11K records. To ameliorate this, following the methodologies described in D7.1, we increased the volume of each dataset by a factor of 10, following the statistical distribution of each one so to carry out the experimentation with an aggregation of 2.7M data points, spread across 276 different data types and 14 datasets.

3.4.2 Data Flows

In this section we describe the workflows the Wine Making pilot follows inside the BDG stack, during which the experimentation will be performed. It should be noted that the identified steps are the currently completed and integrated ones into the BDG stack that support this specific use case.

Initially, similarly to the previous use cases the 14 datasets are uploaded into the BDG platform and are stored into our Elasticsearch instance. The next step for this pilot is to convert the data into RDF and store them into GraphDB for the semantic enrichment to take place. After that the enriched data is exported from GraphDB and ingested back to Elasticsearch. Following the storage of the data into the Elastic stack of BDG the data goes through 2 predictive and analytics steps, one is the uncertainty-aware visual component available [here](#) and the other one is the wine quality prediction available [here](#). The latter two however are not taken into consideration in the context of the current version of the experimentation since they have not been integrated into the BDG stack.

3.4.3 Experimentation

In this section we conduct our experimentations on the provided datasets for each of the identified steps of the data flows they follow inside the BDG stack. For our experimentation we use our synthetically generated datasets that follow the statistical distributions of the provided ones and will test three usage scenarios per step.

3.4.3.1 Dataset Upload

In this step all of the datasets provided are uploaded into the BDG stack. Similar to the approach described in section 3.3.3.1, we will showcase three usage scenarios. One in which 50% of the datasets will be uploaded sequentially, another where all of the datasets will be uploaded sequentially and a final one involving the upload of all of the datasets concurrently.

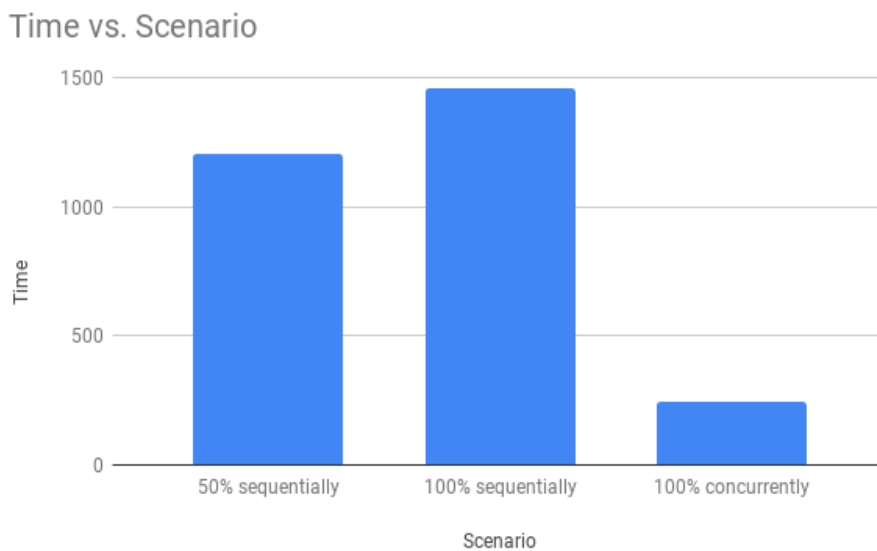


Figure 42: Completion time for dataset upload




Figure 43: CPU usage during dataset upload




Figure 44: Network usage during dataset upload




Figure 45: Memory usage during dataset upload

As shown in the figure XX above, the first scenario takes approximately 1,200 seconds, while for the second one we observe an increase of 300 seconds, in the time required for completion. The last scenario where all of the data is ingested concurrently shows a great decrease in terms of completion time, taking 245 seconds to finish. In terms of CPU we observe a steady increase across the 3 scenarios starting at 35% for the first, 80% for the second and resulting in a 250% for the last one. The same behaviour is also observed for the network traffic of the stack, with an average increase of 350KB/s per scenario, starting from 350KB/s and leading to 1MB/s for the last one. Finally, the memory usage shows a very steady value across all scenarios.

In conclusion to the experimentation for this step, we see that the performance of the last scenario with a high concurrency value and volume performs very well with the deployed BDG stack, resulting, as is the case with other use cases for this specific step as well, to a high parallelization of the involved components.

3.4.3.2 Rdfization & Semantic Enrichment

In section 3.3.3.3 we thoroughly experiment on this step with higher volume datasets so do not conduct the same experiment here with the provided datasets.

3.4.3.3 Ingestion of Semantically Enriched Data

This step concerns the extraction of the semantically enriched data from the GraphDB instance of the stack and storage of this data into the Elasticsearch one. We split this experimentation into 3 scenarios. In the first one we will monitor the performance concurrent ingestion of 10% of the synthetically generated data, in the next one we will ingest 50% of it and on the final one we will extract and ingest all of the 2.7M data points. All of the scenarios were executed using 100 concurrent requests with batches of 10000.

Time vs. # of Records

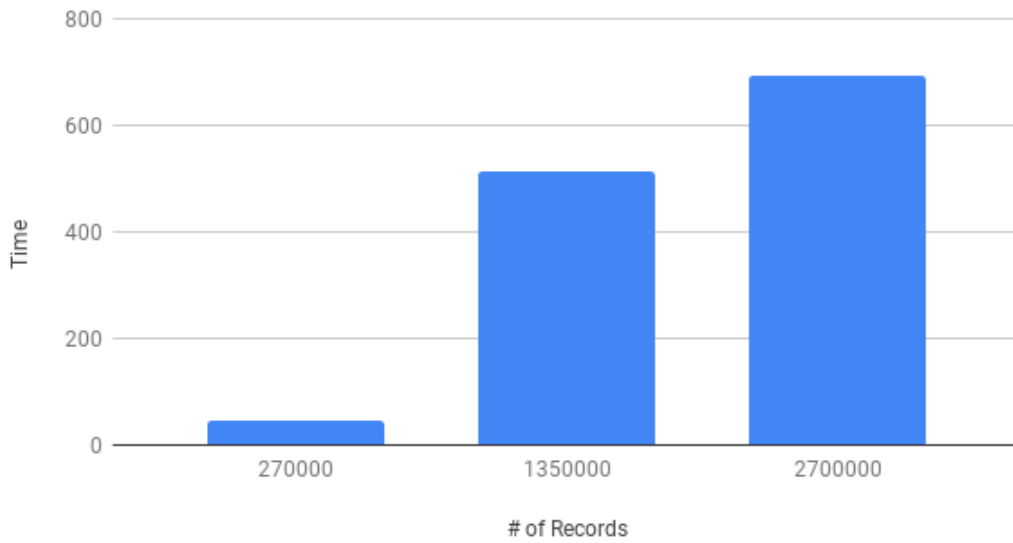


Figure 46: Completion time for the ingestion of semantically enriched data




Figure 47: CPU usage during the ingestion of semantically enriched data




Figure 48: Network usage during the ingestion of semantically enriched data




Figure 49: Memory usage during ingestion of semantically enriched data

The figures shown above present the behaviour of the monitored metrics for the 3 chosen scenarios for this specific step. The first one, where 270K of records are extracted from GraphDB and stored into Elasticsearch, was completed in 46 seconds with an average CPU usage of 265% and a network traffic of 15MB/s. For the ingestion of 1.35M of records, 512 seconds were required for the completion, with a CPU usage peak of 440% and we also observe a minor increase in terms of network traffic, leading to an average of 17MB/s for this scenario. Finally, the last scenario, where all of the data are extracted and inserted into Elasticsearch, 692 seconds are required for the completion, while a small increase of CPU usage is observed (on average 470%). In terms of network traffic this scenario shows similar behaviour to the previous, having 18-20MB/s throughout this experimentation. The memory usage remains on the same levels for all the chosen scenarios.

Similarly, to our conclusion for this step in the different use cases, we observe very high usage of all the metrics throughout the BDG stack, which leads us to create background tasks that frequently sync the data between the 2 frameworks.

3.4.4 End to End Report

In this section we experimented on the datasets provided by the Wine Making pilot and the datasets it has provided. In terms of the data flows specific for this pilot we have identified that the initial step the dataset upload step, presents great performance in respect to the completion time as well as the CPU, network and memory usage. It is a step that can be easily made with a high degree of concurrency without seriously affecting the rest of the stack. Following the experimentation, we performed on the rdfization step, we consider this step to have a good performance regardless of the volume of the data, using the command line tool developed for this step. Interestingly for the step, that of the semantically enriched data extraction and storage into Elasticsearch, we observed a slightly different behaviour than the other pilots. This difference can be explained considering the differences in the data model for each pilot and the better performance of this step is directly affected by the work done in D3.1.

It should be noted at this point, as mentioned in the introductory section, that all of the described experimentation has been performed for the completed steps that are integrated into the BDG stack. Furthermore, due to the lack of Big Data coverage for the provided datasets, the experimentation was performed with high amounts of synthetic data, which although follow the statistical distribution of the provided may decrease the accuracy of the results.

4 CONCLUSIONS

In this document we performed rigorous experimentation of the provided datasets using the BDG stack. To measure the performance of the experimentation we used a number of metrics, as suggested in D7.1 and described in the respective section of this document. We moved on analyzing each of the provided datasets against Big Data Vs and generated synthetic data following the same statistical distribution where needed. For each pilot we identified and abstractly described the data flows its datasets follow inside the BDG stack, with respect to the desired outcomes for each. Every completed and integrated step was then experimented upon using 3 different usage scenarios in an attempt to stress test the stack. Finally, for each pilot we created an end to end report on the observed performance of each of the steps and components it utilizes inside the stack.

We should note that due to the current version of the BDG stack not being the final one, we plan to further expand our experimentation to cover the components not currently completed and integrated into the stack. Once all of the components are integrated and usable in the stack, we also plan to further experiment on the provided datasets by having more concurrent requests made to the stack covering multiple pilots, datasets and steps of the data flows.

Table 25: Experimentation outcomes

Dataflow Step	Experimentation Outcomes	Suggestions
Sensor data ingestion / Data streams pipeline	Highest performance was observed as the parallelization was increased.	We suggest a high parallelization value for this step having as upper limit the number of data types collected.
Satellite image processing data ingestion	Highest performance was observed by having an increased concurrency value.	We suggest an increased concurrency value, up to the number of fields.
Dataset upload	Performance increased by lowering the volumes of data handled or by increasing concurrency.	We suggest either increasing the parallelization value or splitting the uploaded datasets into small batches.
Rdfization & Semantic enrichment	Lowest completion time observed with the usage of the command line tool.	We suggest this step to be scheduled as a cron job handling the creation or update/delete of data.
Ingestion of semantically enriched data	Best performance observed with low volumes of data.	We suggest this step to be scheduled as a cron job handling small batches of data.
Correlation of data	Performance was increased as the volume increased.	We suggest the split of this step into 2 substeps, triggering the correlation through offline jobs.

