

3rd Big Data Innovation Conference

7-8 June 2018
Frankfurt Germany

Improving Business Performance Through Big Data Benchmarking



Todor Ivanov (todor@dbis.cs.uni-frankfurt.de)
Frankfurt Big Data Lab, Goethe University



Atos



Institut "Jožef Stefan", Ljubljana, Slovenija



POLITECNICO
MILANO 1863

Frankfurt Big Data Lab
GOETHE UNIVERSITÄT

Benchmarking

- The Term Benchmark:
A benchmark is a measured „best-in-class“ achievement recognised as the standard of excellence for that business process. (APQC 1993)
- **Business Performance Benchmarking** – comparison of performance measures for the purpose of determining how good one’s own company is compared to others.
- **A software benchmark** is a program used for comparison of software products/tools executing on a pre-configured hardware environment.

Performance Metrics

- **Lord Kelvin defined KPIs as:** “When you can measure what are speaking about and measure it in numbers, you know something about it, when you cannot express it in numbers, your knowledge is of meager and unsatisfactory kind; it may be the beginning of knowledge but you have scarcely, in your thoughts advanced to the stage of science.”

[Arora Amishi, Kaur Sukhbir. *Performance assessment model for management educators based on KRA/KPI*. In: *International conference on technology and business management March*, vol. 23; 2015.]

- **Key Performance Indicators (KPIs)** - tell you what to do to highly increase performance.

- **Jim Gray** describes the benchmarking as: “This quantitative comparison starts with the definition of a benchmark or workload. The benchmark is run on several different systems, and the performance and price of each system is measured and recorded. Performance is typically a throughput metric (work/second) and price is typically a five-year cost-of-ownership metric. Together, they give a price/performance ratio.”

[Jim Gray (Ed.), *The Benchmark Handbook for Database and Transaction Systems*, 1992]

- TPC defines **complex performance metrics** for each standard benchmark:
 - Transaction per second (tps)
 - \$/tps

Example Use Case

System A	System B
4 Nodes (servers)	30 Nodes (servers)
4 hours (execution time)	4 minutes (execution time)
5500 Euro (total cost)	50 000 Euro (total cost)

- Which system will have **best price/performance** for my application?
→ ***Need to use a benchmark***
- What is more important the **minimal execution time or lower total cost?**
→ ***Depends on the business KPIs***
- The company decide to introduce Machine Learning model to improve product recommendations to customers. Different ML models have different Accuracy. How important is the accuracy? Should the company invest in improving the model accuracy?
→ ***Business decision***

Goals of DataBench

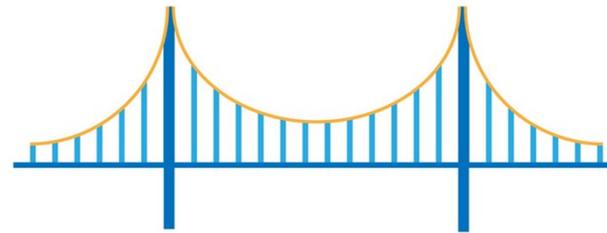
- ***... There is a lack of objective, evidence-based methods measuring the correlation between Big Data Technology benchmarks and Business benchmarks as well as Big Data Technology capability to impact the competitiveness and market success of organizations.***
- Objective:
... The objective is to provide a model which correlates technical benchmarks to performance and business needs of different sectors and domains. ...



Evidence Based Big Data Benchmarking
to Improve Business Performance

Building a bridge between technical and business benchmarking

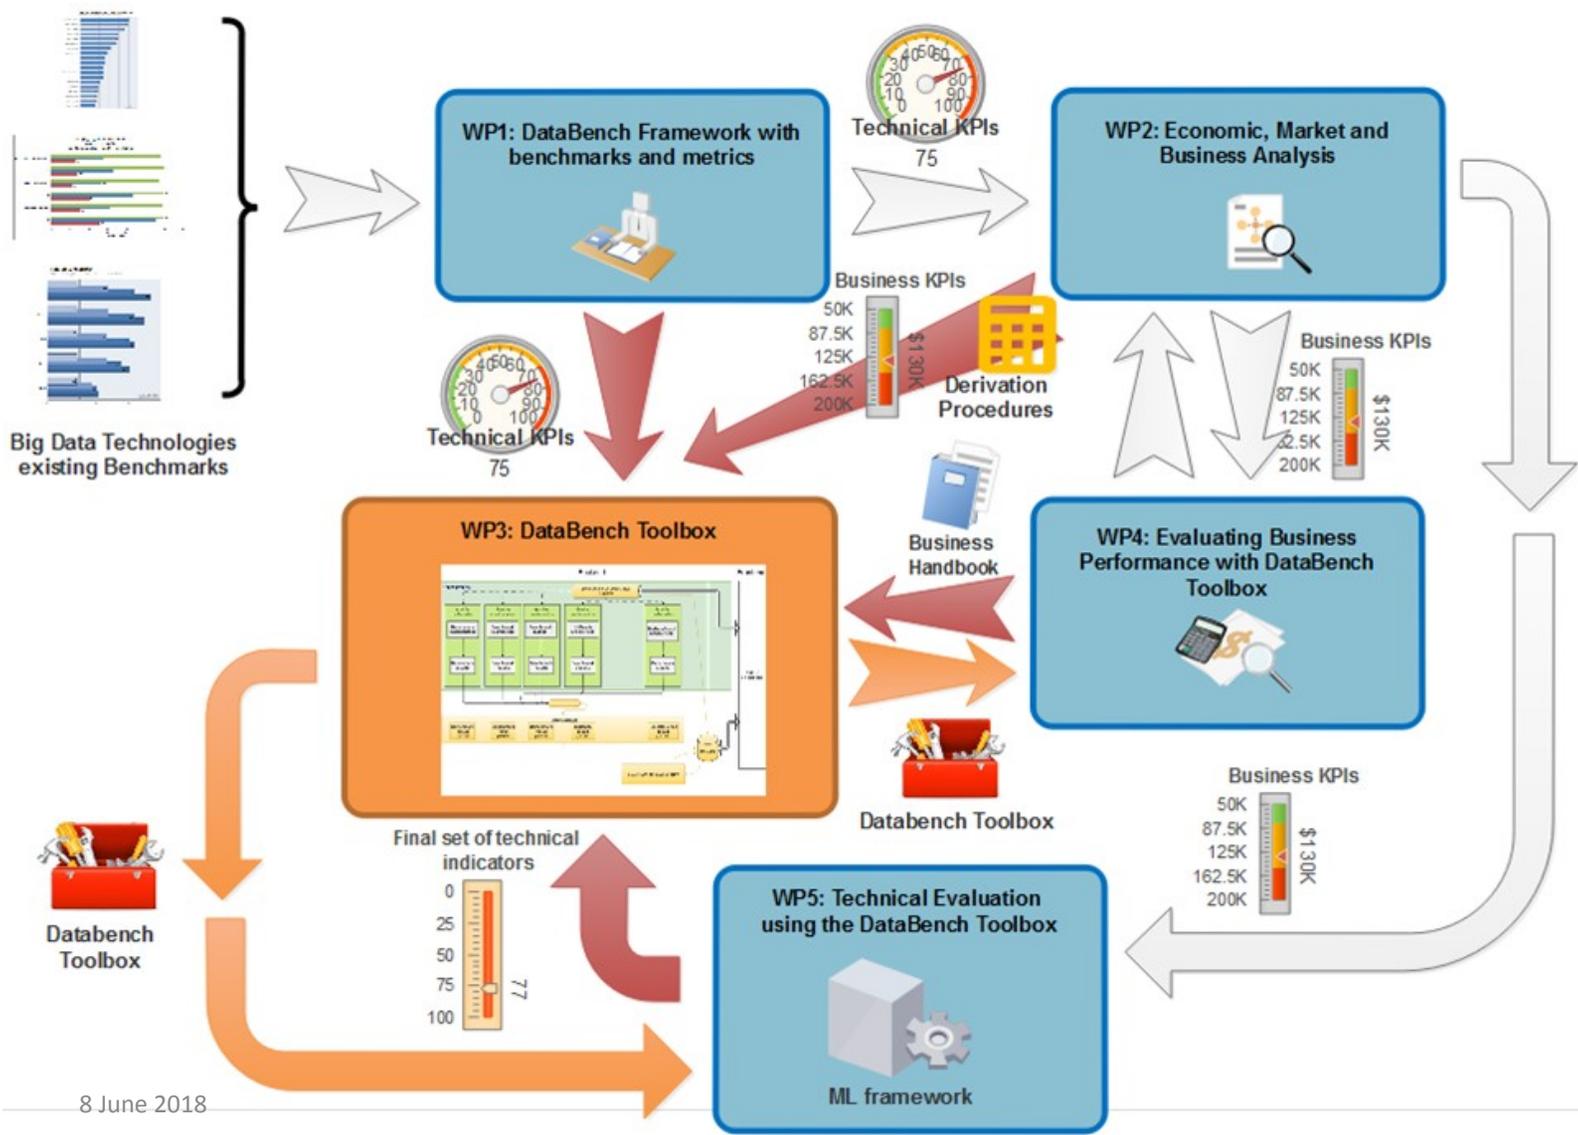
Mapping and
assessing
technical
benchmarks



Evaluating
business
performance and
benchmarks

Develop a Benchmarking Toolbox
and Handbook © IDC





DataBench ToolBox



Holistic benchmarking approach for big data

- The DataBench Toolbox will be a **component-based system** of both **vertical** (holistic/business/data type driven) **and horizontal** (technical area based) **big data benchmarks**, following the layered architecture provide by the **BDVA reference model**.

Not reinventing the wheel, but use wheels to build a new car

- It should be able to **work** or integrate **with existing benchmarking initiatives** and resources where possible.

Filling gaps

- The Toolbox will investigate **gaps of industrial significance** in the big data benchmarking field and contribute to overcome them.

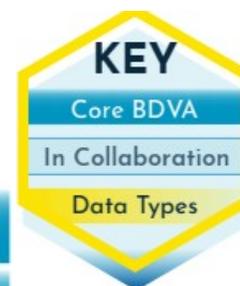
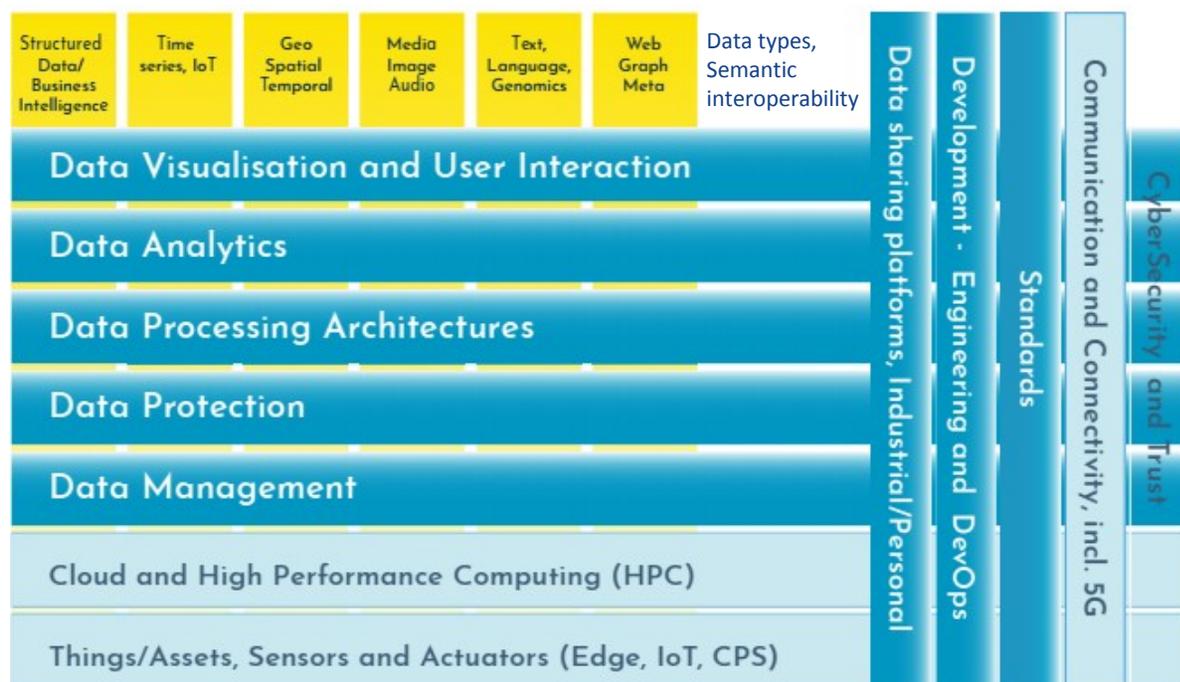
Homogenising metrics

- The Toolbox will implement ways to derive as much as possible **the DataBench technical metrics and business KPIs** from the metrics extracted from the integrated benchmarking.

Web user interface

- It will include a web-based visualization layer to **assist to the final users to specify their benchmarking requirements**, such as selected benchmark, data generators, workloads, metrics and the preferred data, volume and velocity, **as well as searching and monitoring** capabilities.

BDV – Big Data and Analytics/Machine Learning Reference Model



Source: http://bdva.eu/sites/default/files/BDVA_SRIA_v4_Ed1.1.pdf

BDVA Reference Model

Vertical, Incl. Data Types	Domain/Sector/Business solutions KPIs ...	Work-in-progress (WP1)																													
		30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	9	8	7	6	5	4	3	2	1	
Verticals, Incl. Data Types	Business	X	X					X	X		X		X		X	X	X	X			X										
	Transport			X																X											
	Manufacturing																													X	
	Energy																														
	Bioinformatics							X																							
	Health																														
	Telecommunication													X	X																
	Finance																														
	Social Media																														
	General Micro-benchmarks																														
	Standardized Benchmark	X	X																												
	MetaData																														
	Graph, Network																														
Text, NLP, Web																															
Image, Audio																															
Spatio Temp																															
Time Series, IoT																															
Structured, BI	X	X																													
Analytics, Processing, Data Management, Infra	Visual Analytics																														
	Industrial Analytics (Descriptive, Diagnostic, Predictive, Prescriptive)																														
	Machine Learning, AI, Data Science																														
	Streaming/ Realtime Processing																														
	Interactive Processing	X	X																												
	Batch Processing	X	X																												
	Data Privacy/Security																														
	Data Governance/Mgmt																														
	Data Storage	X	X																												
	Communication & Connectivity																														
	Cloud Services & HPC, Edge																														
	Benchmarks	MLPerf ABench																													
		IDEBench																													
DAWNPerf/MLPerf																															
Senka																															
OpenML benchmark suites																															
Penn machine learning benchmark (PMLB)																															
Sarasa																															
BigBench V2																															
TPC-HS v2																															
Hobbit Benchmark																															
RIOTBench																															
AdBench																															
Fathom																															
TensorFlow Benchmarks																															
DeepMark																															
DeepBench																															
Shenzhen Transportation System (SZT)																															
Yahoo Streaming Benchmark (YSB)																															
Graphalytics																															
TPC-BB																															
TPC-DS v2																															
BigFUN																															
TPC-V																															
SparkBench																															
TPC-HS																															
StreamBench																															
Social Network Benchmark																															
Sennan Publishing Benchmark (SPB)																															
PRIMEBALL																															
BigFrame																															
LambBench																															
BigDataBench																															
BigBench																															
AMPLab Big Data Benchmark																															
MERS																															
CloudSuite																															
PUMA Benchmark Suite																															
CloudRank-D																															
SWIM																															
YCSB																															
Hibench																															
CALDA																															
MRBench																															
PgMix																															
GridMix																															
Hadoop Workload Examples																															
Liner Road																															
TPC-DS v1																															
TPC-H																															

Some of the benchmarks to integrate (I)

Micro-benchmarks: stress test very specific part of the Big Data technologies typically utilizing synthetic data.

Year	Name	Type
2010	HiBench	Big data benchmark suite for evaluating different big data frameworks. 19 workloads including synthetic micro-benchmarks and real-world applications from 6 categories which are micro, machine learning, sql, graph, websearch and streaming .
2015	SparkBench	System for benchmarking and simulating Spark jobs . Multiple workloads organized in 4 categories.
2010	Yahoo! Cloud System Benchmark (YCSB)	Evaluates performance of different “ key-value ” and “ cloud ” servicing systems , which do not support the ACID properties. The YCSB++ , an extension, includes many additions such as multi-tester coordination for increased load and eventual consistency measurement.
2017	TPCx-IoT	Based on YCSB, but with significant changes. Workloads of data ingestion and concurrent queries simulating workloads on typical IoT Gateway systems . Dataset with data from sensors from electric power station(s).

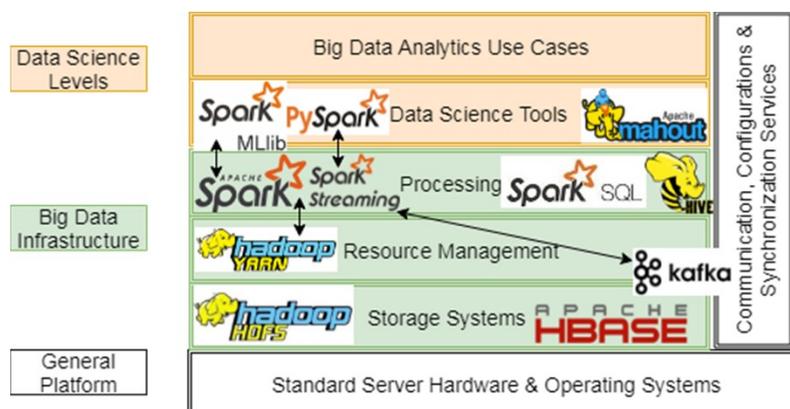
Some of the benchmarks to integrate (II)

Application-level benchmarks: cover a wider and more complex set of functionalities involving multiple Big Data technologies and typically utilizing real scenario data.

Year	Name	Type
2015	Yahoo Streaming Benchmark (YSB)	The Yahoo Streaming Benchmark is a streaming application benchmark simulating an advertisement analytics pipeline.
2013	BigBench/TPCx-BB	BigBench is an end-to-end, technology agnostic, application-level benchmark that tests the analytical capabilities of a Big Data platform. It is based on a fictional product retailer business model.
2017	BigBench V2	Similar to BigBench, BigBench V2 is an end-to-end, technology agnostic, application-level benchmark that tests the analytical capabilities of a Big Data platform.
2018	ABench (Work-in-Progress)	New type of multi-purpose Big Data benchmark covering many big data scenarios and implementations. Extends other benchmarks such as BigBench.

ABench (Work-in-Progress)

- Growing number of new **Big Data technologies** and **connectors** in the Big Data Stacks
 → Challenges for Solution Architects, Data Engineers, Data Scientist, Developers, etc.



ABench: Big Data Architecture Stack Benchmark.

Todor Ivanov, Rekha Singhal:
 Companion of the 2018 ACM/SPEC International Conference
 on Performance Engineering, ICPE 2018,
 Berlin, Germany, April 09-13, 2018.

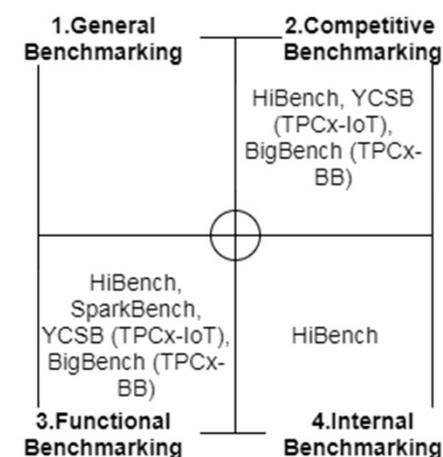
- Missing benchmarks for **each technology, connector** or a **combination of them**
- Consequence → **Increasing complexity in the Big Data Architecture Stacks**
- Our approach → **ABench: Big Data Architecture Stack Benchmark**

ABench Features

- Benchmark Framework
 - Data generators or plugins for custom data generators
 - Include data generator or public data sets to simulate workload that stresses the architecture
- Reuse of existing benchmarks
 - Case study using BigBench (on-going Streaming and Machine Learning extensions)
- Open source implementation and extendable design
- Easy to setup and extend
- Supporting and combining all four types of benchmarks in ABench

Benchmarks Types (adapted from Andersen and Pettersen)

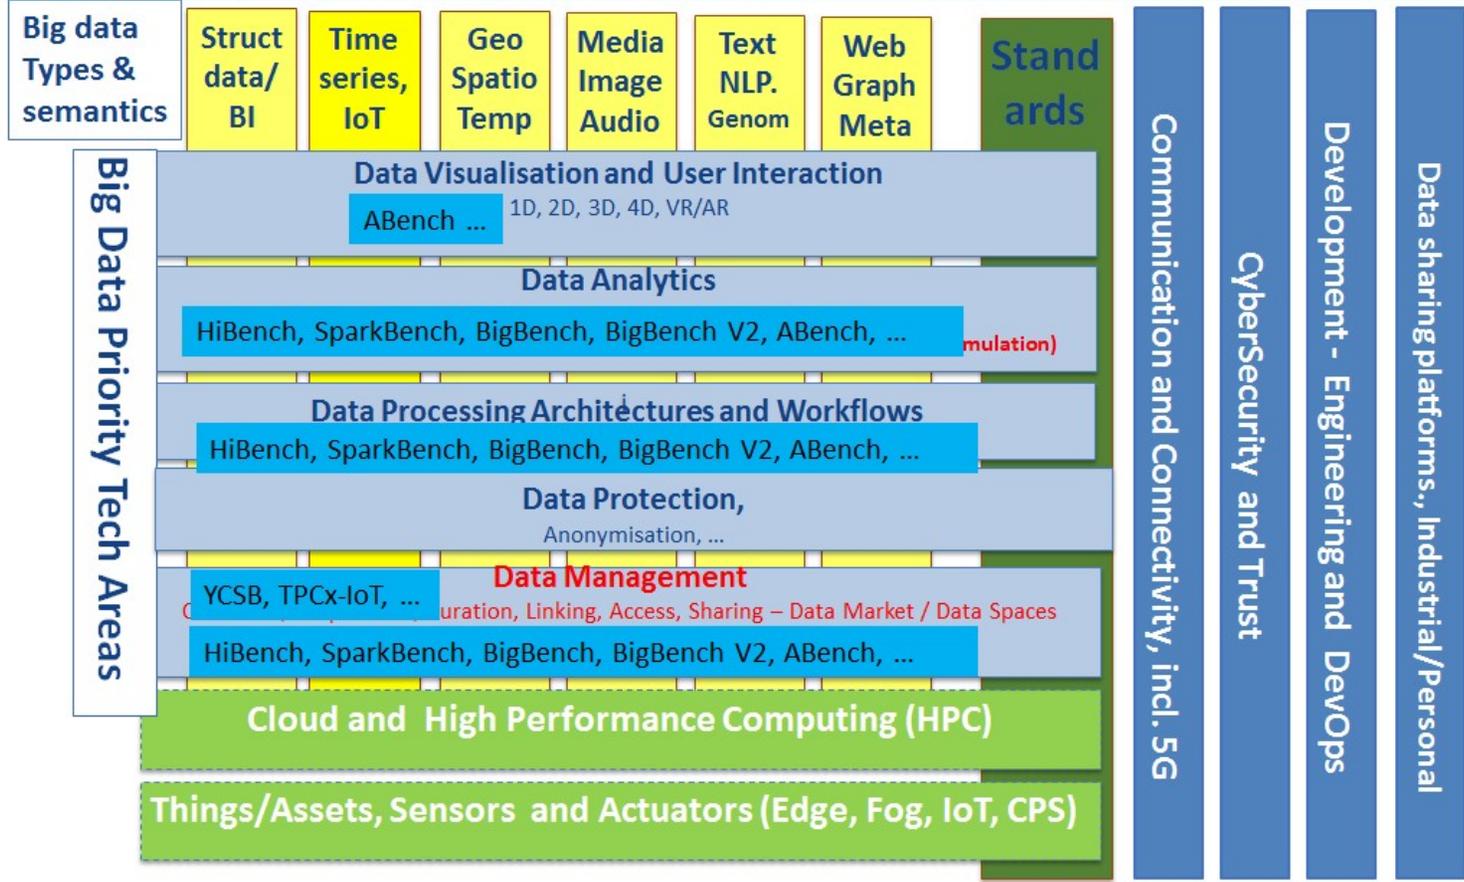
1. **Generic Benchmarking:** checks whether an implementation fulfills given business requirements and specifications (*Is the defined business specification implemented accurately?*).
2. **Competitive Benchmarking:** is a performance comparison between the best tools on the platform layer that offer similar functionality (*e.g., throughput of MapReduce vs. Spark vs. Flink*).
3. **Functional Benchmarking** is a functional comparison of the features of the tool against technologies from the same area. (*e.g., Spark Streaming vs. Spark Structured Streaming vs. Flink Streaming*).
4. **Internal Benchmarking:** comparing different implementations of a functionality (*e.g., Spark Scala vs. Java vs. R vs. PySpark*)



Category	Year	Name	Type	Domain	Data Type
Micro-benchmarks	2010	HiBench	Micro-benchmark Suite	Micro-benchmarks, Machine Learning, SQL, Websearch, Graph, Streaming Benchmarks	Structured, Text, Web Graph
	2015	SparkBench	Micro-benchmark Suite	Machine Learning, Graph Computation, SQL, Streaming Application	Structured, Text, Web Graph
	2010	YCSB	Micro-benchmark	cloud OLTP operations	Structured
	2017	TPCx-IoT	Micro-benchmark	workloads on typical IoT Gateway systems	Structured, IoT
Application Benchmarks	2015	Yahoo Streaming Benchmark	Application Streaming Benchmark	advertisement analytics pipeline	Structured, Time Series
	2013	BigBench/TPCx-BB	Application End-to-end Benchmark	a fictional product retailer platform	Structured, Text, JSON logs
	2017	BigBench V2	Application End-to-end Benchmark	a fictional product retailer platform	Structured, Text, JSON logs
	2018	ABench (Work-in-Progress)	Big Data Architecture Stack Benchmark	set of different workloads	Structured, Text, JSON logs

Big Data Value Reference Model

Applications/Solutions: Manufacturing, Health, Energy, Transport, BioEco, Media, Telco, Finance, EO, SE, ...



DataBench Summary

- A **framework for big data benchmarking** for all Big Data practitioners.
- Will provide **methodology** and **tools**.
- Added value:
 - An **umbrella to access to multiple benchmarks**
 - Homogenized **technical metrics**
 - Derived **business KPIs**
 - **Connecting** the different **benchmark communities**
- Open initiative: Public-private partnership (PPP) projects, industrial partners (BDVA and beyond) and benchmarking initiatives are **welcomed to work with us, either to use our framework or to add new benchmarks.**

Contacts



info@databench.eu



@DataBench_eu



DataBench



DataBench Project



DataBench