LDBC

Collaborative Project

FP7 – 317548

D5.5.6 Showcase

Coordinator: Ioan Toma (UIBK)
With contributions from: Alice Carpentier (UIBK), Serge Tymaniuk (UIBK), Damaris Coll (UPC), Ricard Tapias (UPC), Iliya Enchev (ONTO)
1st Quality reviewer: Josep L. Larriba (UPC)
2nd Quality reviewer: Peter Boncz (VUA)

Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (PU)
Contractual delivery date:	31.03.2015
Actual delivery date:	08.04.2015
Version:	1.0
Total number of pages:	15
Keywords:	Showcase, dissemination, community building

Abstract

This deliverable supplies a public, Web-enabled Showcase that can be used to disseminate and raise awareness about the final LDBC results. The Showcase includes pointers to publications, deliverables, presentations, screencasts, software, datasets and actions towards building a community produced by the project. All this information is provided as part of the LDBC organization Web portal at ldbcouncil.org.

Executive summary

This deliverable supplies a public, Web-enabled Showcase that can be used to disseminate and raise awareness about the final LDBC results. The Showcase includes pointers to publications, deliverables, presentations, screencasts, software, datasets and actions toward building a community produced by the project. All this information is provided as part of the LDBC organization Web portal at ldbcouncil.org.

Document Information

IST Project Number	FP7 - 317548	**Acronym**	LDBC
Full Title	LDBC		
Project URL	http://ldbcouncil.org/ and http://www.ldbc.eu/		
Document URL			
EU Project Officer	Carola Carstens		

Deliverable	**Number**	D5.5.6	**Title**	Showcase
Work Package	**Number**	WP5	**Title**	Community building and dissemination

Date of Delivery	**Contractual**	M30	**Actual**	M30
Status	version 1.0		final ■	
Nature	prototype □ report ■ dissemination □			
Dissemination level	public ■ consortium □			

Authors (Partner)				
Responsible Author	**Name**	Ioan Toma	**E-mail**	ioan.toma@sti2.at
	Partner	UIBK	**Phone**	+43 512 507 53721

Abstract (for dissemination)	This deliverable supplies a public, Web-enabled Showcase that can be used to disseminate and raise awareness about the final LDBC results. The Showcase includes pointers to publications, deliverables, presentations, screencasts, software, datasets and actions toward building a community produced by the project. All this information is provided as part of the LDBC organization Web portal at ldbcouncil.org.
Keywords	Showcase, dissemination, community building

Version Log			
Issue Date	**Rev. No.**	**Author**	**Change**
25.03.2015	0.1	Ioan Toma	First complete version.
06.04.2015	0.2	Ioan Toma	Include content from UPC. Address Larri's comments.
08.04.2015	1.0	Ioan Toma	Include 2nd input from Larri, address comments.

Table of Contents

1 Introduction

One of the main objectives of LDBC work package 5 (WP5) is to setup and coordinate the dissemination of project activities. All project partners are contributing to this objective that includes a wide range of activities such as the creation and maintenance of the project Web portal, the generation of promotional materials, presentations, posters etc. Among the last LDBC dissemination activities carried out as part of the project, we have gathered all relevant dissemination materials and provide that as part of the LDBC showcase. This deliverable provides information about the showcase that can be used by the Commission to disseminate and raise awareness about LDBC results. The Showcase includes pointers to publications, deliverables, presentations, software and datasets produced by the project. All this information is provided as part of the LDBC organization Web portal at ldbcouncil.org.

The structure of this deliverable is as follows. In Section 2 we provide pointers to the organization of workshops, publications, deliverables and presentations. Section 3 includes pointers to software and datasets developed by LDBC. Section 4 provides all the actions towards community building by the consortium. Section 5 summarizes the actions to be carried in the future as continuation of the project. Finally, Section 6 summarizes the deliverable.

2 Publications, Deliverables, Presentations, Screencasts and Organized Workshops

This section contains pointers to LDBC publications, deliverables, presentations, screencasts and organized workshops.

2.1 Publications

The list of LDBC publications, related publications and LDBC project deliverables is available at: http://ldbcouncil.org/publications.

2.1.1 LDBC Publications

1. The LDBC Social Network Benchmark: Interactive Workload. O. Erling, A. Averbuch, J.L. Larriba-Pey, H. Chafi, A. Gubichev, A. Prat, M.-D. Pham, P. Boncz. To be published in SIGMOD'15, Melbourne.
2. The linked data benchmark council: a graph and RDF industry benchmarking effort. R. Angles, P. A. Boncz, J.-L. Larriba-Pey, I. Fundulaki, T. Neumann, O. Erling, P. Neubauer, N. Martínez-Bazan, V. Kotsev, I. Toma: SIGMOD Record 43(1): 27-31 (2014)
3. Benchmarking Linked Open Data Management Systems. R. Angles Rojas, P. Minh Duc, and P. A. Boncz, (2014), Benchmarking Linked Open Data Management Systems. ERCIM News.
4. Introduction to Graph Databases. J.-L. Larriba-Pey, N. Martínez-Bazan, D. Domínguez-Sal: Reasoning Web 2014: 171-194
5. High quality, scalable and parallel community detection for large real graphs.A. Prat-Pérez, D. Dominguez-Sal, J.-L. Larriba-Pey: WWW 2014: 225-236
6. Using semijoin programs to solve traversal queries in graph databases. N. Martínez-Bazan, D. Domínguez-Sal. GRADES 2014
7. How community-like is the structure of synthetically generated graphs? A. Prat-Pérez, D. Domínguez-Sal. GRADES 2014
8. Exploiting the query structure for efficient join ordering in SPARQL queries. Andrey Gubichev, and Thomas Neumann, (2014), Exploiting the query structure for efficient join ordering in SPARQL queries. Proc. 17th International Conference on Extending Database Technology EDBT, Athens, Greece, March 24-28, 2014., page 439--450. OpenProceedings.org.
9. Graph Pattern Matching -- Do We Have to Reinvent the Wheel? Andrey Gubichev, and Manuel Then. (2014), Graph Pattern Matching -- Do We Have to Reinvent the Wheel? Second International Workshop on Graph Data Management Experiences and Systems, GRADES 2014, co-loated with SIGMOD/PODS 2014, Snowbird, Utah, USA, June 22, 2014, CWI/ACM.
10. How to generate query parameters in RDF benchmarks? Andrey Gubichev, Renzo Angles, and Peter A. Boncz, (2014), How to generate query parameters in RDF benchmarks? Workshops Proceedings of the 30th International Conference on Data Engineering Workshops, ICDE 2014, Chicago, IL, USA, March 31 - April 4, 2014, page 272--274. IEEE, http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6818339.
11. MonetDB/RDF: Discovering and Exploiting The Emergent Schema Of RDF Data. P. Minh Duc, and P. A. Boncz, (2014), MonetDB/RDF: Discovering And Exploiting The Emergent Schema Of RDF Data ERCIM News.
12. Parameter Curation For Benchmark Queries. A. Gubichev, and P. A. Boncz, (2014), Parameter Curation For Benchmark Queries. Proceedings of the TPC Technology Conference on Performance Evaluation; Benchmarking TPCTC, 2014.

13. Benchmarking database systems for social network applications. Renzo Angles, Arnau Prat-Pérez, David Dominguez-Sal, and Josep-Lluis Larriba-Pey, Grades 2013.
14. Experiences With Virtuoso Cluster RDF Column Store. P. A. Boncz, Orri Erling, and P. Minh Duc, (2013), Experiences With Virtuoso Cluster RDF Column Store. Linked Data Management: Principles and Techniques, page -. CRC Press.
15. On the Formulation of Performant SPARQL Queries. Antonis Loizou, and Paul T. Groth, (2013), On the Formulation of Performant SPARQL Queries. CoRR.
16. Partitioning Graph Databases - A Quantitative Evaluation. Alex Averbuch, and Martin Neumann, (2013), Partitioning Graph Databases - A Quantitative Evaluation. CoRRvolume abs/1301.5121.
17. SPARQLing Kleene -- Fast Property Paths in RDF-3X. Andrey Gubichev, Srikanta Bedathur, and Stephan Seufert. (2013) SPARQLing Kleene -- Fast Property Paths in RDF-3X.First International Workshop on Graph Data Management Experiences and Systems, New York, NY, USA, ACM.
18. The Linked Data Benchmark Council Project. P. A. Boncz, I Fundulaki, A. Gubichev, J. Larriba Pey, and T. Neumann, (July 2013), The Linked Data Benchmark Council Project. Datenbank-Spektrum 13(2): 121-129 (2013)
19. Time-varying social networks in a graph database: a Neo4j use case. Ciro Cattuto, Marco Quaggiotto, André Panisson, and Alex Averbuch, (July 2013), Time-varying social networks in a graph database: a Neo4j use case. Proceeding GRADES '13 First International Workshop on Graph Data Management Experiences and Systemsvolume 1. page 11.
20. TPC-H Analyzed: Hidden Messages And Lessons Learned From An Influential Benchmark. P. A. Boncz, T. Neumann, and O Erling. (August 2013) TPC-H Analyzed: Hidden Messages And Lessons Learned From An Influential Benchmark. Proceedings of the TPC Technology Conference on Performance Evaluation $1amp; Benchmarking TPCTC, 2013.
21. Fast approximation of steiner trees in large graphs. Andrey Gubichev, and Thomas Neumann, (2012), Fast approximation of Steiner trees in large graphs. CIKM, page 1497-1501. ACM.

2.1.2 LDBC Related Publications

1. A Practical Query Language for Graph DBs. Renzo Angles, Pablo Barceló, and Gonzalo Ríos, (2013), A Practical Query Language for Graph DBs. 7th Alberto Mendelzon International Workshop on Foundations of Data Management AMW.
2. On benchmarking online social media analytical queries. Haixin Ma, Jinxian Wei, Weining Qian, Chengcheng Yu, Fan Xia, and Aoying Zhou. (2013) On benchmarking online social media analytical queries. GRADES, page 10.
3. Self-Organizing Structured RDF In MonetDB. P. Minh Duc. (April 2013) Self-Organizing Structured RDF In MonetDB. Proceedings of ICDE/PHD Symposium 2013.
4. TPC-H Analyzed: Hidden Messages And Lessons Learned From An Influential Benchmark. P. A. Boncz, T. Neumann, and O Erling, (August 2013), TPC-H Analyzed: Hidden Messages And Lessons Learned From An Influential Benchmark. Proceedings of the TPC Technology Conference on Performance Evaluation Benchmarking TPCTC, 2013.
5. Virtuoso, a Hybrid RDBMS/Graph Column Store. Orri Erling. (2012) Virtuoso, a Hybrid RDBMS/Graph Column Store. IEEE Data Eng. Bull.volume 35. page 3-8.
6. SRBench: A Streaming RDF/SPARQL Benchmark. Ying Zhang, P. Minh Duc, O. Corcho, and J. P. Calbimonte. (November 2012) SRBench: A Streaming RDF/SPARQL Benchmark. Proceedings of International Semantic Web Conference 2012.
7. Heuristics-based query optimisation for SPARQL. Petros Tsialiamanis, Lefteris Sidirourgos, Irini Fundulaki, Vassilis Christophides, and Peter Boncz. (2012) Heuristics-based query optimisation for SPARQL.

8. Linked Stream Data Processing: Facts And Figures. L. P. Danh, D. T. Minh, P. Minh Duc, P. A. Boncz, E. Thomas, and F. Michael. (November 2012) Linked Stream Data Processing: Facts And Figures. Proceedings of International Semantic Web Conference 2012.
9. Robust Runtime Optimization and Skew-Resistant Execution Of Analytical SPARQL Queries On Pig. S Kotoulas, J Urbani, P. A. Boncz, and P Mika. (November 2012) Robust Runtime Optimization and Skew-Resistant Execution Of Analytical SPARQL Queries On Pig. Proceedings of International Semantic Web Conference 2012, Springer, PDF
10. S3G2: A Scalable Structure-Correlated Social Graph Generator. P. Minh Duc, P. A. Boncz, and O Erling. (2012) S3G2: A Scalable Structure-Correlated Social Graph Generator. Proceedings of TPC Technology Conference on Performance Evaluation & Benchmarking 2012.
11. iGraph in action: performance analysis of disk-based graph indexing techniques. W.S. Han, M.D. Pham, J. Lee, R. Kasperovics, and J.X. Yu. (2011) iGraph in action: performance analysis of disk-based graph indexing techniques. Proceedings of the 2011 international conference on Management of data, page 1241--1242.

2.1.3 LDBC Project Deliverables:

Nr	Name	Workpackage	Lead	Level	Date
D1.1.1	Overview and analysis of existing benchmark frameworks	WP1	FORTH-ICS	PU	March, 2013
D1.1.2	Benchmark principles and methods	WP1	VUA	PU	September, 2013
D1.1.3	Initial benchmark development portal	WP1	UIBK	PU	March, 2013
D1.1.4	Improved version of benchmark development portal	WP1	UIBK	PU	September, 2013
D1.1.5	Final benchmark development portal	WP1	UIBK	PU	September, 2014
D1.1.6	Initial benchmarks integration and release	WP1	ONTO	PU	September, 2014
D1.1.7	Final benchmarks integration and release	WP1	ONTO	PU	March, 2015
D2.2.1	Analysis and Classification of Choke Points	WP2	TUM	PU	March, 2013
D2.2.2	Data Generator	WP2	FORTH-ICS	PU	September, 2013
D2.2.3	Benchmarking transactions	WP2	TUM	PU	September, 2014
D2.2.4	Benchmarking complex queries	WP2	TUM	PU	September, 2014
D3.3.1	Use case analysis and choke point analysis	WP3	NEO	PU	June, 2013
D3.3.2	Graph database infrastructure and language	WP3	UPC	PU	September,

Nr	Name	Workpackage	Lead	Level	Date
	expressivity				2013
D3.3.34	Benchmark design for navigational pattern matching benchmarking	WP3	UPC, NEO	PU	September, 2014
D4.4.1	Use Case Analysis and Classification of Choke Points	WP4	FORTH-ICS	PU	September, 2013
D4.4.2	Benchmark Design for Reasoning	WP4	FORTH-ICS	PU	September, 2014
D4.4.3	Benchmark Design for Instance Matching	WP4	FORTH-ICS	PU	September, 2014
D4.4.4	Benchmark Design for Mapping Management	WP4	VUA	PU	September, 2014
D5.5.1	Dissemination Report v1	WP5	VUA	PU	September, 2013
D5.5.2	Dissemination Report v2	WP5	VUA	PU	September, 2014
D5.5.3	First version of Portal	WP5	UIBK	PU	December, 2012
D5.5.4	Second version of Portal	WP5	UIBK	PU	March, 2014
D5.5.5	Project Fact Sheet	WP5	VUA	PU	October, 2012
D5.5.6	Showcase	WP5	UIBK	PU	March, 2015
D6.6.1	Incorporation of the LDBC foundation	WP6	VUA	CO	September, 2013
D6.6.2	IPR Plan	WP6	OGL	PU	September, 2013
D6.6.3	Auditor Training	WP6	OGL	PU	September, 2014
D6.6.4	Standardisation report	WP6	ONTO	PU	March, 2015

2.2 Presentations

The list of presentations is available at: http://ldbcouncil.org/talks and also on Slideshare at http://www.slideshare.net/ldbcproject/ and includes:

1. Peter Boncz – LDBC: Industry-strength benchmarks for Graph Data Management - Keynote Speaker at IDEAS 2013, http://confsys.encs.concordia.ca/IDEAS/ideas13/

2. Peter Boncz – Benchmarking Graph Data Management Systems - Keynote Speaker at EDBT/ICDT 2014, http://www.edbticdt2014.gr/index.php/keynotes
3. Peter Neubauer- The Linked Data Benchmark council - LDBC FOSDEM 2013, https://archive.fosdem.org/2013/schedule/event/ldbc/
4. Josep-L. Larriba-Pey, N. Martinez-Bazán, D. Dominguez-Sal. Graph Databases and their Applications. Lecturers at the Reasoning Web Summer School, 2014 http://rw2014.di.uoa.gr/?q=Lectures

2.3 Screencasts

Screencasts that showcase the LDBC benchmarks are available at:

1. Semantic Publishing Benchmark http://www.ldbcouncil.org/benchmarks/spb
2. Social Network Benchmark http://www.ldbcouncil.org/benchmarks/snb

The consortium has also published the lectures and presentations of the last TUC meetings in the LDBC channel of Youtube. This can be found in: https://www.youtube.com/channel/UC6HbzfJ4016Vez-2HKNeDag

2.4 Workshops

LDBC has been very active in creating a community of user and technology companies, practitioners and scientists that are active in the area of graph management and benchmarking. The events organized around those activities is:

1. GRADES. This workshop, organized in conjunction with the SIGMOD conference has been run in three different occasions:
    a. **Grades'13**. New York, Sunday June 23rd 2013, with 30 papers submitted and 16 accepted. The number of attendees was more than 50. See: http://event.cwi.nl/grades2013/
    b. **Grades'14.** Snowbird, Sunday June 22nd 2014, with 25 papers submitted and 12 accepted. The number of attendees was more than 50. See: http://event.cwi.nl/grades2014/
    c. **Grades'15.** Melbourne Sunday May 31st 2015, not celebrated yet. See the web page of the workshop for more details: http://event.cwi.nl/grades2015/
2. BeRSys. This workshop has been organized in conjunction with ESWC and VLDB conferences and it is a forum for the community around benchmarking RDF systems. The workshop has been organized the following two times:
    a. **BeRSys'13**. Montpellier, May 26th 2013. See: http://www.ldbc.eu/events/bersys-2013/
    b. **BeRSys'14**. Hangzou, September 5th, 2014. See: http://events.sti2.at/bersys2014/
3. Graph-TA. This is a one day workshop organized in Barcelona by DAMA-UPC aiming at the community of graph scientists and practitioners. The workshop is attended by a large community of more than 50 people. The web site of the events is:
    a. **First Graph-TA**, February 19th 2013. See the link to the event here: http://www.dama.upc.edu/seminars/1st-graph-ta
    b. **Second Graph-TA**, February 21st 2014. See the link to the event here: http://www.dama.upc.edu/seminars/2nd-graph-ta
    c. **Third Graph-TA**, March 18th 2015. See the link to the event here: http://www.dama.upc.edu/seminars/3rd-graph-ta

3 Software and Datasets

The software produced in LDBC is made available on github as follows:

1. Semantic Publishing Benchmark. This can be found in different sections of GitHub:
    a. **Documentation**. The documentation can be found in: https://github.com/ldbc/ldbc_snb_docs
    b. **Implementations**. The vendor implementations of the benchmark can be found in: https://github.com/ldbc/ldbc_snb_implementations
    c. **Data generator**. The data generator can be found in: https://github.com/ldbc/ldbc_snb_datagen_0.2
    d. **Driver**. The driver of the Interactive workload can be found in: https://github.com/ldbc/ldbc_driver
    e. **Validation sets**. The validation code for the benchmark can be found in: https://github.com/ldbc/ldbc_snb_interactive_validation
2. Social Network Benchmark. This can also be found in different sections of GitHub:
    a. **First version**. The first version of the SPB can be found in: https://github.com/ldbc/ldbc_spb_bm
    b. **Second version**. The second version of the benchmark can be found in: https://github.com/ldbc/ldbc_spb_bm_2.0
    c. **Ontology**. The benchmark ontology can be found in: https://github.com/ldbc/ldbc_bm_ontology

4 Actions towards community building

In this respect, LDBC has organized six Technical User Community meetings that have been very fruitful for the dissemination of the work of the project among the community of users, technology companies and practitioners in the area, as well as for the collection of requirements from the User Community. The schedule of those events and the slides of the talks can be found here: http://wiki.ldbcouncil.org/display/TUC/Events

In addition, there have been a series of actions through Social Networks, like Twitter, LinkedIn and Facebook. It is possible to find all the activity of LDBC in those Social Networks in:

1. **Twitter**. https://twitter.com/ldbcouncil
2. **Facebook**. https://www.facebook.com/ldbcouncil
3. **LinkedIn**. https://www.linkedin.com/grp/home?gid=4955240
4. **Google+.** https://plus.google.com/105060163998563565976/videos

In addition, we have created a Chanel in **Youtube**. In this channel, one may find the videos of the last TUC meetings: https://www.youtube.com/channel/UC6HbzfJ4016Vez-2HKNeDag

Although this is not the matter of this Deliverable, we have created a set of reports on the progress of the Social activity and the progress in terms of audience. Those reports are kept private to the consortium.

5 Future actions

LDBC continues its action through the post-FP7 project organization. This allows the addition of new companies and a structure that allows for the management of its actions. The structures that give structure and rule the LDBC organization are:

1. **Task Forces**. The life of LDBC is organized around the Task Forces. A TF is formed by a group of persons from different organizations or personally, to carry a specific task like designing and implementing a Benchmark, or workload within a benchmark.
2. **Board of Directors**. Each company member of LDBC is allowed to have a Director in the organization. The Directors meet every 6 months (the first meetings took place on Nov 2014 with the occasion of the 5th TUC meeting in Athens, and March 2015 with the occasion of the 6th TUC meeting in Barcelona) and take decisions about the life of the organization.

   Some decisions taken in the last BoD meetings are:

   a. To support a Task Force for the creation of the Analytics Workload.

   b. To support a Task Force to work towards the creation of a standard Graph Query Language.
3. **TUC meetings**. LDBC will continue to organize TUC meetings. The objective of those TUC meetings is the engagement of new companies and organizations in the LDBC. The following TUC meetings will be:

   a. The seventh TUC will be organized in the TJ Watson Research Center in NY, and hosted by IBM next November 2015.

   b. The eighth TUC meeting will probably be organized in Europe.

   c. The ninth TUC meeting will be organized in the Oracle Headquarters, in the US west coast.

6 Conclusions

In this deliverable we provided a set of pointers to publications, deliverables, presentations, screencasts, software and datasets, that can be used to showcase LDBC. All this information is provided as part of the LDBC organization Web portal.

We believe that LDBC has attained its dissemination objectives and the showcase described in this report is a clear example. In particular, the dissemination of the project was multiplied by a significant factor due to the very constructive comments of the Reviewers and Project Officer, who motivated a great deal of creativity to make a series of actions a real and tangible asset for the project. Those are:

1. The creation of a better portal for the project, www.ldbcouncil.org.
2. The creation of periodic content with added value information for the users and practitioners of benchmarking.
3. The creation of accounts in the most significant Social Networks, that are now part of our everyday relationship with this segment of valuable knowledge.

Those actions, together with the work of the consortium members, have allowed us to:

1. Make the benchmarking community conscious of the need for an industry fuelled organization.
2. Involve important companies in the organization such as Oracle and IBM.
3. Start thinking of the future big, with the creation of events related to Graph technologies and benchmarking.