

TECHNOLOGY TRANSFER PRESENTS

# RICK VAN DER LANS

---

## THE LOGICAL DATA WAREHOUSE

---

### DESIGN, ARCHITECTURE AND TECHNOLOGY

---

**MAY 28-29, 2018**  
RESIDENZA DI RIPETTA - VIA DI RIPETTA, 231  
ROME (ITALY)



info@technologytransfer.it  
www.technologytransfer.it

## ABOUT THIS SEMINAR

Business Intelligence has changed dramatically the last years. The time-to-market for new reports and analysis has to be shortened, new data sources have to be made available to business users more quickly, self-service BI and Data Science must be supported, more and more users want to work with zero-latency data, adoption of new technologies, such as Hadoop, Spark, and NoSQL, must be easy, and analysis of streaming data and Big Data is required.

The classic Data Warehouse architecture has served many organizations well. But it's not the right architecture for this new world of BI. It's time for organizations to migrate gradually to a more flexible architecture: the Logical Data Warehouse architecture. This architecture, introduced by Gartner, is based on a decoupling of reporting and analyses on the one hand, and data sources on the other hand.

Classic Data Warehouse architectures are made up of a chain of databases. This chain consists of numerous databases, such as the staging area, the central Data Warehouse and several data marts, and countless ETL programs needed to pump data through the chain. Integrating self-service BI products with this architecture is not easy and certainly not if users want to access the source systems. Delivering 100% up-to-date data to support operational BI is difficult to implement. And how do we embed new storage technologies into this architecture?

With the Logical Data Warehouse architecture new data sources can be hooked up to the Data Warehouse more quickly, self-service BI can be supported correctly, operational BI is easy to implement, the adoption of new technology is much easier, and in which the processing of Big Data is not a technological revolution, but an evolution.

The technology to create a Logical Data Warehouse is available, and many organizations have already completed the migration successfully; a migration that is based on a step-by-step process and not on a full rip-and-replace approach.

In this practical seminar, the architecture is explained and products will be discussed. It discusses how organizations can migrate their existing architecture to this new one. Tips and design guidelines are given to help make this migration as efficient as possible.

### LEARNING OBJECTIVES

- What are the practical benefits of the Logical Data Warehouse Architecture and what are the differences with the classical architecture
- How can organizations step-by-step and successfully migrate to this flexible Logical Data Warehouse Architecture?
- You will learn about the possibilities and limitations of the various available products
- How do data virtualization products work?
- How can big data be added transparently to the existing BI environment?
- How can self-service BI be integrated with the classical forms of BI?
- How can users be granted access to 100% up-to-date data without disrupting the operational systems?
- What are the real-life experiences of organizations that have already implemented a Logical Data Warehouse?

# OUTLINE

## 1. Challenges for the Classic Data Warehouse

- Integrating Big Data with existing data and making it available for reporting and analytics
- Supporting self-service BI and self-service data preparation
- Faster time-to-market for reports
- Polyglot persistency – processing data stored in Hadoop and NoSQL systems
- Operational Business Intelligence, or analyzing of 100% up-to-date data

## 2. The Logical Data Warehouse

- The essence: decoupling of reporting and data sources
- From batch-integration to on-demand integration of data
- The impact on flexibility and productivity – an improved time-to-market for reports
- Examples of organizations operating a Logical Data Warehouse
- Can a logical data warehouse really work without a physical data warehouse?

## 3. Implementing a Logical Data Warehouse with data virtualization servers

- Why data virtualization?
- Market overview: AtScale, Cirro Data Hub, Cisco Information Server, Data Virtuality UltraWrap, Denodo Platform, RedHat JBoss Data Virtualization, Rocket DV, and Stone Bond Enterprise Enabler
- Importing non-relational data, such as XML and JSON documents, Web Services, NoSQL, and Hadoop data
- The importance of an integrated business glossary and centralization of metadata specifications

## 4. Improving the query performance of data virtualization servers

- How does caching really work
- Which virtual tables should be cached?
- Query optimization techniques and the explain feature
- Smart drivers/connectors can help improve query performance
- How can SQL-on-Hadoop engines speed up query performance?
- Working with multiple data virtualization servers in a distributed environment to minimize network traffic

## 5. Migrating to a Logical Data Warehouse

- An A to Z roadmap
- Guidelines for the development of a Logical Data Warehouse
- Three different methods for modeling: outside-in, inside-out, and middle-out
- The value of a canonical data model
- Considerations for security aspects
- Step by step dismantling of the existing architecture
- The focus on sharing of metadata specifications for integration, transformation, and cleansing

## 6. Self-Service BI and the Logical Data Warehouse

- Why self-service BI can lead to “report chaos”
- Centralizing and reusing metadata specifications with a Logical Data Warehouse
- Upgrading self-service BI into managed self-service BI
- Implementing Gartner’s BI-modal environment

## **7. Big Data and the Logical Data Warehouse**

- New data storage technologies for Big Data, including Hadoop, MongoDB, Cassandra
- The appearance of the polyglot persistent environment; or each application its own optimal database technology
- Design rules to integrate Big Data and the Data Warehouse seamlessly
- Big Data is too “big” to copy
- Offloading cold data with a logical Data Warehouse

## **8. Physical Data Lakes or virtual Data Lakes?**

- What is a Data Lake?
- Is developing a physical Data Lake realistic when working with Big Data?
- Developing a virtual Data Lake with data virtualization servers
- Can the logical Data Warehouse and the virtual Data Lake be combined?

## **9. Implementing Operational BI with a Logical Data Warehouse**

- Examples of operational reporting and operational analytics
- Extending a Logical Data Warehouse with operational data for real-time analytics
- “Streaming” data in a Logical Data Warehouse
- The coupling of data replication and data virtualization

## **10. Making Data Vault more flexible with a Logical Data Warehouse**

- What exactly is Data Vault?
- Using a Logical Data Warehouse to make data in a Data Vault available for reporting and analytics
- The structured SuperNova design technique to develop virtual data marts
- SuperNova turns a Data Vault in a flexible database

## **11. The Logical Data Warehouse and the Environment**

- Design principles to define data quality rules in a Logical Data Warehouse
- How data preparation can be integrated with a logical Data Warehouse
- Shifting of tasks in the BICC
- Which new development and design skills are important?
- The impact on the entire design and development process

## **12. Summary and Conclusions**

### **WHO SHOULD ATTEND**

This seminar is intended for everyone who needs to be aware of developments in the field of Business Intelligence and Data Warehousing, such as BI architects, business analysts, Data Warehouse and database designers, database experts, consultants, technology planners, project managers, and system analysts. Some knowledge of the classical Data Warehouse architecture is required.

# INFORMATION

<p><b>PARTICIPATION FEE</b></p> <p>€ 1300</p> <p>The fee includes all seminar documentation, luncheon and coffee breaks.</p> <p><b>VENUE</b></p> <p>Residenza di Ripetta Via di Ripetta, 231 Rome (Italy)</p> <p><b>SEMINAR TIMETABLE</b></p> <p>9.30 am - 1.00 pm 2.00 pm - 5.00 pm</p>	<p><b>HOW TO REGISTER</b></p> <p>You must send the registration form with the receipt of the payment to: TECHNOLOGY TRANSFER S.r.l. Piazza Cavour, 3 - 00193 Rome (Italy) Fax +39-06-6871102</p> <p><b>within May 14, 2018</b></p> <p><b>PAYMENT</b></p> <p>Wire transfer to: Technology Transfer S.r.l. Banca: Cariparma Agenzia 1 di Roma IBAN Code: IT 03 W 06230 03202 000057031348 BIC/SWIFT: CRPPIT2P546</p>	<p><b>GENERAL CONDITIONS</b></p> <p><b>DISCOUNT</b></p> <p>The participants who will register 30 days before the seminar are entitled to a 5% discount.</p> <p>If a company registers 5 participants to the same seminar, it will pay only for 4.</p> <p>Those who benefit of this discount are not entitled to other discounts for the same seminar.</p> <p><b>CANCELLATION POLICY</b></p> <p>A full refund is given for any cancellation received more than 15 days before the seminar starts. Cancellations less than 15 days prior the event are liable for 50% of the fee. Cancellations less than one week prior to the event date will be liable for the full fee.</p> <p><b>CANCELLATION LIABILITY</b></p> <p>In the case of cancellation of an event for any reason, Technology Transfer's liability is limited to the return of the registration fee only.</p>
--	--	--

**RICK VAN DER LANS**  
**THE LOGICAL DATA WAREHOUSE**  
**DESIGN, ARCHITECTURE AND**  
**TECHNOLOGY**

May 28-29, 2018  
Residenza di Ripetta  
Via di Ripetta, 231  
Rome (Italy)

Registration fee:  
€ 1300

*If registered participants are unable to attend, or in case of cancellation of the seminar, the general conditions mentioned before are applicable.*

first name .....

surname .....

job title .....

organisation .....

address .....

postcode .....

city .....

country .....

telephone .....

fax .....

e-mail .....



Stamp and signature

Send your registration form with the receipt of the payment to:  
**Technology Transfer S.r.l.**  
Piazza Cavour, 3 - 00193 Rome (Italy)  
Tel. +39-06-6832227 - Fax +39-06-6871102  
info@technologytransfer.it  
www.technologytransfer.it



## **SPEAKER**

**Rick van der Lans** is an independent consultant, author and lecturer. He specializes in Business Intelligence, Data Warehousing, and Service Oriented Architectures. He is managing director of R20/Consultancy. Mr. van der Lans is an internationally acclaimed lecturer. For the last 12 years, he has been presenting professionally, and has lectured in many of the European countries, South America, the USA, and in Australia. He has presented many keynote speeches at international events. Mr. van der Lans is chairman of the Database Systems Show (organised annually in The Netherlands since 1984), he is columnist for two major newspapers in the Benelux, called *Computable* and *DataNews*. Additionally, he is advisor for magazines such as Software Release Magazine and Database Magazine. His popular books, including “**Introduction to SQL**” and “**The SQL Guide to Oracle**”, have been translated into many languages and have sold over 100,000 copies. Recently, he has published a very successful book on presentation skills.