

TECHNOLOGY TRANSFER PRESENTS

# RICK VAN DER LANS

---

## **INCORPORATING BIG DATA, HADOOP AND NOSQL IN DW AND BI SYSTEMS**

---

ONLINE LIVE STREAMING

NOVEMBER 3-4, 2021



info@technologytransfer.it  
www.technologytransfer.it

## **ABOUT THIS SEMINAR**

Big Data, Hadoop, in-memory analytics, Spark, Kafka, self-service BI, data warehouse automation, analytical database servers, data virtualization, data vault, operational intelligence, predictive analytics, and NoSQL are just a few of the new technologies and techniques that have become available for developing BI systems. Most of them are very powerful and allow for development of more flexible and scalable BI systems. But which ones do you pick?

Due to this waterfall of new developments, it's becoming harder and harder for organizations to select the right tools. Which technologies are relevant? Are they mature? What are their use cases? These are all valid but difficult to answer questions.

This seminar gives a clear, extensive, and critical overview of all the new developments and their inter-relationships. Technologies and techniques are explained, market overviews are presented, strengths and weaknesses are discussed, and guidelines and best practices are given.

The biggest revolution in BI is evidently Big Data. Therefore, considerable time in the seminar is reserved for this intriguing topic. Hadoop, Spark, MapReduce, Kafka, Hive, NoSQL, SQL-on-Hadoop are all explained. In addition, the relation with analytics is discussed extensively.

This seminar gives you a unique opportunity to see and learn about all the new BI developments. It's the perfect update for those interested in knowing how to make BI systems ready for the coming ten years.

### **WHAT YOU WILL LEARN**

- Learn about the trends and the technological developments related to Business Intelligence, Analytics, Data Warehousing, and Big Data
- Discover the value of Big Data and Analytics for organizations
- Learn which products and technologies are winners and which ones are losers
- Learn how new and existing technologies, such as Hadoop, NoSQL and NewSQL, will help you create new opportunities in your organization
- Learn how more Agile Data Business Intelligence systems can be designed
- Learn how to embed Big Data and Analytics in existing Business Intelligence architectures

### **WHO SHOULD ATTEND**

- Business Intelligence Specialists
- Data Warehouse Designers
- Business Analysts
- Technology Planners
- Technical Architects
- Enterprise Architects
- IT Consultants
- IT Strategists
- Systems Analysts
- Database Developers
- Database Administrators
- Solutions Architects
- Data Architects
- IT Managers

# OUTLINE

## 1. The Changing World of Business Intelligence

- Big Data: hype or reality?
- Operational Intelligence: does it require online Data Warehouses?
- Fast data is the next frontier of big data
- Data Warehouses in the Cloud
- Self-service BI
- The business value of analytics

## 2. Hadoop explained

- The relationship between Big Data and analytics
- The Hadoop software stack explained, including HDFS, MapReduce, YARN, Kudu, Hive, Impala, Storm, Sqoop, Flume, and HBase
- The balancing act: productivity versus scalability
- Making Big Data available to a larger audience with SQL-on-Hadoop engines, such as Apache Drill, Apache Hive, Apache Impala, Apache Phoenix, Dremio, HP Vertica, IBM BigSQL, JethroData, MemSQL, SparkSQL and Splice Machine

## 3. Spark explained

- Spark is about in-memory analytical processing
- The interfaces: SQL, R, Scala, Python
- Does Spark need Hadoop?
- The relationship between Spark and data science
- Examples of use cases of Spark

## 4. NoSQL explained

- Classification of NoSQL database servers: key-value stores, document stores, column-family stores and graph data stores
- Market overview: CouchDB, Cassandra, Cloudera, MongoDB, and Neo4j
- Strong consistency or eventual consistency?
- Why an aggregate data model?
- Use case of NoSQL products
- How to analyze data stored in NoSQL databases

## 5. Overview of Analytical SQL Database Servers

- Are classic SQL database servers more suitable for Data Warehousing?
- Important performance improving features: column-oriented storage, in-database analytics
- The new generation of GPU-based database servers: BlazingDB, Kinetica, MapD, SQream
- Market overview of analytical SQL database servers: Amazon Athena and Redshift, Google BigQuery, IBM PureData Systems for Analytics, InfoBright, Kognitio WX2, Microsoft, SAP HANA, SnowflakeDB, Teradata Vantage, and Vertica

## 6. Technologies for fast data and streaming analytics

- The key use case for fast data: the Internet of Things (IoT)
- IoT implies streaming data and fast analysis of data - analytics at the speed of business
- IoT devices: Smartphones (watches), RFID sensors, machines, general sensors, cameras, pace makers, and so on
- The challenge: real-time reactions on streaming data
- The difference between big data and fast big data
- Technologies for streaming data: Apache Kafka, Apache ActiveMQ, Amazon Kinesis, Kestrel, RabbitMQ, and ZeroMQ
- Differences between these new technologies and traditional message queuing products
- Products for big data streaming: Apache Storm and Flink, IBM InfoSphere Streams, Informatica for Streaming Analytics, Software AG Apama, and Spark Streaming
- How to integrate fast data with the enterprise data warehouse?

## 7. New Forms of Reporting and Analytics

- Data Virtualization offers on-demand data integration
- Seamlessly integrating big data and the data warehouse
- Market overview: AtScale, DataVirtuality UltraWrap, Denodo Platform, IBM Data Virtualization Manager, RedHat JBoss Data Virtualization, Stone Bond Enterprise Enabler, and Tibco Data Virtualization
- Importing non-relational data, such as XML documents, web services, NoSQL and Hadoop data, and unstructured data
- Differences between data virtualization and with data blending
- Are SQL-on-Hadoop engines becoming data virtualization tools?

## 8. New Business Intelligence Architectures

- Discussion of different BI architectures, including Kimball's Data Warehouse Bus, Architecture, Inmon's Corporate Information Factory, DW 2.0, the Federated Architecture, the Centralized Warehouse Architecture, the Data Virtualization Architecture, and the BI in the Cloud Architecture
- Do we still need Data Marts?
- What is the role of master Ddata Mmanagement in BI architectures?
- Using data vault to create more flexible data warehouses
- Data Warehouse automation to create data warehouses and data marts faster

## 9. NewSQL and Translytical Database Servers

- NewSQL stands for high-performance transactional SQL database servers
- Simpler transaction mechanisms to implement scale-out
- What does the term geo-compliance mean?

- Market overview: Clustrix, GenieDB, memSQL, NuoDB, and VoltDB
- Combining transactions and analytics = Translytical database server
- Market overview includes MemSQL, SAP HANA, and SpliceMachine

## 10. Summary and Conclusions

- Explanation of non-relational concepts, such as column families, hierarchies, sets, and lists
- Is storing unstructured and semi-structured data really more flexible?
- The differences between schema-on-read and schema-on-write
- Rules for transforming classic data models to NoSQL concepts
- Application needs influence database design

## 11. Closing Remarks

# INFORMATION

<p><b>PARTICIPATION FEE</b></p> <p>€ 1100</p> <p>The fee includes all seminar documentation.</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: info@technologytransfer.it</p> <p>TECHNOLOGY TRANSFER S.r.l. Piazza Cavour, 3 - 00193 Rome (Italy)</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 INCORPORATING BIG DATA, AND NOSQL IN DW AND BI SYSTEMS

November 3-4, 2021

Registration fee:  
€ 1100

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 a highly-respected independent analyst, consultant, author, and internationally acclaimed lecturer specializing in data warehousing, business intelligence, big data, and database technology.

He has presented countless seminars, webinars, and keynotes at industry-leading conferences. He also helps clients worldwide to design their data warehouse, big data, and business intelligence architectures and solutions and assists them with selecting the right products. He has been influential in introducing the new logical data warehouse architecture worldwide which helps organizations to develop more agile business intelligence systems.

Over the years, Rick has written hundreds of articles and blogs for newspapers and websites and has authored many educational and popular white papers for a long list of vendors. He was the author of the first available book on SQL, entitled including Introduction to SQL, which has been translated into several languages with more than 100,000 copies sold. More recently, he published his book **Data Virtualization for Business Intelligence Systems**.

He presents seminars, keynotes, and in-house sessions on Big data and analytics, data virtualization, the logical data warehouse, data warehousing and business intelligence.