

Updated on 24/07/2025

Sign up

Exasol High-Speed Analytics Training

2 days (14 hours)

Presentation

Master Exasol High-Speed Analytics in its expert dimension with this comprehensive training course, designed for data engineers, analysts and architects wishing to harness the full power of this high-performance in-memory data warehouse.

The course begins with an exploration of Exasol's technical architecture, intelligent memory management, parallelization and compression mechanisms. You'll learn how to install, configure and secure a cluster, whether in the cloud or on-premise.

You will then be trained in advanced analytical modeling, high-performance ingestion via files, connectors or Kafka feeds, and the exploitation of Exasol's enriched SQL, with analytical functions, scripts and UDFs in Python or Java for customized processing.

Performance optimization is at the heart of the program: execution plans, partitioning, query tuning, real-time supervision and good design practices will be studied in detail, with concrete case studies.

As with all our training courses, this one will be presented with the latest [exasol](#) updates.

Objectives

- Understand Exasol's massively parallel, in-memory architecture (MPP),
- Design, model, partition and optimize data schemas for real-time analysis.
- master methods for ingesting massive data in batch or streaming mode
- Be able to secure, administer and supervise an Exasol cluster,
- Industrialize queries, SQL scripts and UDFs using orchestration tools
- Take full advantage of Exasol's analytical engine with advanced SQL.

Target audience

- Data architects
- SQL developers

Prerequisites

- Basic knowledge of SQL language

Exasol High-Speed Analytics training program Introduction to

Exasol

- Positioning in the analytics market
- Typical use cases
- Comparison with Snowflake, Redshift, BigQuery
- Hybrid memory architecture
- Massively parallel processing
- Data compression & automatic indexing
- In-Memory Query Engine
- Standard & analytical SQL support
- BI and Data Science connectors

Installation and configuration

- On-premise vs. cloud vs. ExaCloud
- Deployment via Docker/VM images
- Network configuration and high availability
- Adding nodes & load balancing
- Critical system parameters
- Monitoring and alerts
- User and role management
- Encryption of data at rest and in transit
- Security audits and logs

Data modeling and ingestion

- Connectors

- Import with EXALoader, EXAPARTNER and EXAplus
- Mass loading techniques
- Star schema vs. Snowflake schema
- Denormalization and performance
- Table partitioning
- Data distribution between nodes
- Use of compression
- Optimized data types

SQL language in Exasol

- Analytical functions
- CTE
- GROUP BY ROLLUP, CUBE
- SQL Scripts
- Loops, conditions, cursors
- Error handling and logs
- Creation in Python, Lua, Java
- UDF debugging and logging
- Use cases: ML scoring, complex transformations

Performance optimization

- Using EXPLAIN PLAN
- Reading performance metrics
- Detecting bottlenecks
- Invisible index principle
- Automatic access optimization
- Horizontal partitioning of large tables
- Avoiding costly joins
- Optimizing sub-queries and views
- Use of column statistics

Integration with the analytics ecosystem

- Connection with Power BI, Tableau, Qlik, Looker
- DirectQuery vs. Import
- High-frequency Dashboard on Exasol
- Exasol-Python: pandas, scikit-learn, seaborn
- Jupyter Notebooks with Exasol
- Exasol R plugin and native R calls
- Integration with Airflow, Jenkins, dbt
- Deployment of versioned scripts

- Automating SQL query testing

Practical examples and case studies

- Workflow management via Kafka
- Real-time sales aggregation
- Dynamic visualization with Power BI
- Feature engineering in Exasol
- ML model integration
- Production application via UDF Python
- Optimizing a Talend/Exasol workflow
- Using Bulk Insert mode
- Performance monitoring via logs

Administration and supervision

- EXAoperation interface
- Monitoring slow requests
- Access history and execution times
- Incremental vs. full backup
- ExaBackup & ExaRestore
- Automatic restore tests
- Hot addition of new nodes
- Partition rebalancing
- Measuring horizontal scalability

Companies concerned

This training course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced IT technology, or to acquire specific business knowledge or modern methods.

Positioning on entry to training

Positioning at the start of training complies with Qualiopi quality criteria. As soon as registration is finalized, the learner receives a self-assessment questionnaire which enables us to assess his or her estimated level of proficiency in different types of technology, as well as his or her expectations and personal objectives for the forthcoming course, within the limits imposed by the selected format. This questionnaire also enables us to anticipate any connection or security difficulties within the company (intra-company or virtual classroom) which could be problematic for the follow-up and smooth running of the training session.

Teaching methods

Practical training: 60% hands-on, 40% theory. Training material distributed in digital format to all participants.

Organization

The course alternates theoretical input from the trainer, supported by examples, with brainstorming sessions and group work.

Validation

At the end of the session, a multiple-choice questionnaire is used to check that skills have been correctly acquired.

Certification

A certificate will be awarded to each trainee who completes the training course.