NoSQL Databases and Unstructured Data Management Training Course



14 - 18 Dec 2026

Vienna

5800 € (Per Person)

Ref: #DM5180_91404

Course Introduction / Overview:

In today's data-driven world, the limitations of traditional relational databases have become increasingly apparent, paving the way for the rise of NoSQL technologies designed to handle the volume, velocity, and variety of modern unstructured data. This comprehensive training course provides a deep dive into the world of NoSQL databases, moving beyond theoretical concepts to impart practical, applicable skills. We will explore the fundamental principles that differentiate NoSQL from SQL, guided by the insights of thought leaders like Martin Fowler and his seminal work, "NoSQL Distilled". Participants will gain a thorough understanding of the four main types of NoSQL databases: document, key-value, column-family, and graph. The curriculum is meticulously designed to cover everything from core concepts like the CAP theorem and data modeling techniques to hands-on interaction with leading platforms such as MongoDB, Redis, Cassandra, and Neo4j. BIG BEN Training Center is proud to offer this program, which equips professionals with the expertise to design, implement, and manage scalable, high-performance data solutions, ensuring they can effectively harness the power of unstructured data to drive innovation and business value. This course is the definitive guide for anyone looking to master the technologies that underpin modern big data applications.

Target Audience / This training course is suitable for:

- Database Administrators and Architects.
- Software Developers and Engineers.
- Data Scientists and Analysts.
- IT Managers and Team Leaders.
- System Integrators and Consultants.
- Business Intelligence Professionals.
- Professionals transitioning into big data roles.
- Solutions Architects.

Target Sectors and Industries:

- Information Technology and Software Services.
- E-commerce and Retail.
- Financial Services and FinTech.
- Healthcare and Life Sciences.
- Telecommunications and Media.
- Government and Public Sector Agencies.
- Manufacturing and IoT.
- Logistics and Supply Chain Management.

Target Organizations Departments:

- Information Technology (IT) and Infrastructure.
- Software Development and Engineering.
- Data Science and Analytics.
- Research and Development (R&D).
- Business Intelligence and Reporting.
- Product Management and Development.
- Digital Transformation and Innovation Units.
- Operations and Systems Management.

Course Offerings:

By the end of this course, the participants will have able to:

- Differentiate between SQL and NoSQL databases and identify their respective use cases.
- Explain the principles of the CAP theorem and its implications for distributed systems.
- Classify and describe the four main types of NoSQL databases.
- Design and implement data models for various unstructured data scenarios.
- Perform CRUD (Create, Read, Update, Delete) operations in a document database like MongoDB.
- Utilize key-value stores like Redis for caching and session management.
- Understand the architecture and data distribution model of column-family databases like Cassandra.
- Model and query complex relationships using a graph database like Neo4j.
- Evaluate and select the appropriate NoSQL database for specific business requirements.
- Develop strategies for scaling and ensuring high availability of NoSQL solutions.
- Integrate NoSQL databases within a broader big data ecosystem.

Course Methodology:

The training methodology at BIG BEN Training Center is designed to be immersive, practical, and highly interactive, ensuring that participants not only learn the theory but also gain hands-on experience. This course moves beyond traditional lectures by incorporating a blended learning approach. Each session combines expert-led instruction with practical lab exercises, allowing participants to work directly with different NoSQL databases and solve real-world problems. We utilize a series of carefully curated case studies from various industries to demonstrate the application and impact of NoSQL technology in business contexts. Group discussions and collaborative workshops are integral to the learning process, encouraging participants to share insights, tackle complex data modeling challenges together, and learn from peer experiences. Our instructors facilitate these sessions, providing continuous feedback and personalized guidance to ensure every participant achieves the learning objectives. The focus is on building practical, job-ready skills, empowering attendees to confidently apply their new knowledge to design, deploy, and manage robust NoSQL solutions within their own organizations immediately after the course.

Course Agenda (Course Units):

Unit One: Foundations of NoSQL and Unstructured Data

- Introduction to the Big Data Era.
- Limitations of Relational Database Management Systems (RDBMS).
- Defining NoSQL and its core principles.
- The CAP Theorem (Consistency, Availability, Partition Tolerance) explained.
- Exploring the different types of NoSQL databases (Document, Key-Value, Column-Family, Graph).
- Use cases and real-world applications of NoSQL.
- Comparing SQL and NoSQL data architectures.

Unit Two: Deep Dive into Document Databases with MongoDB

- Introduction to document-oriented databases.
- Understanding JSON, BSON, and data structures.
- Core concepts of MongoDB: Collections, Documents, and Schemas.
- Performing CRUD (Create, Read, Update, Delete) operations.
- Indexing strategies for performance optimization.
- The Aggregation Framework for data analysis.
- Schema design and data modeling best practices in MongoDB.

Unit Three: Mastering Key-Value and Column-Family Databases

- Exploring the principles of Key-Value stores.
- Practical applications with Redis: caching, session management, and message queues.
- Introduction to Column-Family (Wide-Column) databases.
- Understanding the architecture of Apache Cassandra.
- Data modeling and schema design in Cassandra.
- The Cassandra Query Language (CQL).
- Strategies for achieving high availability and fault tolerance.

Unit Four: Understanding Graph Databases with Neo4j

- Introduction to graph theory and graph databases.
- Core concepts: Nodes, Relationships, and Properties.
- The property graph model explained.
- Introduction to Cypher, the declarative graph query language.
- Modeling complex, interconnected data.
- Performing advanced queries and pathfinding.
- Common use cases: social networks, recommendation engines, and fraud detection.

Unit Five: Advanced NoSQL Concepts and Strategic Implementation

- Advanced data modeling techniques for unstructured data.
- Scalability strategies: Sharding, Replication, and Clustering.
- Security considerations in NoSQL databases.
- Integrating NoSQL with existing enterprise systems and data pipelines.
- Polyglot Persistence: Using multiple databases in one system.
- Evaluating and choosing the right NoSQL database for your project.
- Capstone Project: Designing a NoSQL solution for a real-world business problem.

FAQ:

Qualifications required for registering to this course?

There are no requirements.

How long is each daily session, and what is the total number of training hours for the course?

This training course spans five days, with daily sessions ranging between 4 to 5 hours, including breaks and interactive activities, bringing the total duration to 20 - 25 training hours.

Something to think about:

Considering the CAP theorem's trade-offs, how might the choice between consistency and availability in a NoSQL database fundamentally alter a global e-commerce platform's business model and user experience?

What unique qualities does this course offer compared to other courses?

This course distinguishes itself by offering a holistic and comparative perspective on the NoSQL ecosystem, rather than focusing on a single technology. While many programs provide expertise in one database, our curriculum is strategically designed to cover the four primary types of NoSQL databases, empowering participants with the critical ability to evaluate and select the optimal tool for any given business problem. The emphasis is not just on technical execution but on architectural and strategic thinking. We delve deep into the "why" behind each technology, exploring the trade-offs dictated by principles like the CAP theorem and how they impact system design and business outcomes. The methodology is intensely practical, moving beyond theory with hands-on labs for each database type—MongoDB, Redis, Cassandra, and Neo4j. This ensures participants develop tangible skills and a versatile understanding of data modeling, querying, and management across different paradigms. Ultimately, this course provides a future-proof skill set, preparing professionals to design and implement sophisticated, scalable, and resilient data solutions in a world of increasingly complex and unstructured information.