



Data Warehousing with BigQuery: Storage Design

Course Introduction:

Data Warehousing with BigQuery: Storage Design provides learners with an in-depth understanding of how to design efficient and scalable storage solutions. This course covers the fundamental concepts and best practices of using Google BigQuery for data warehousing. Participants will explore the architecture, learn to optimize storage, and leverage the unique capabilities of BigQuery to manage large datasets effectively. This course is ideal for data professionals looking to harness the power of BigQuery for their storage needs.

Module 1: Introduction to Data Warehousing and BigQuery

- Understanding Data Warehousing Concepts

Gain insights into the principles and architecture of data warehousing and its significance in data management.

- Overview of BigQuery

Explore the core functionalities and key features of Google BigQuery that make it suitable for modern data warehousing.

- The Role of BigQuery in the Data Ecosystem

Examine how BigQuery integrates with other tools and technologies to form a comprehensive data ecosystem.

Module 2: Architecture and Storage Considerations

- BigQuery Architecture Deep Dive

Delve into the internal architecture of BigQuery, understanding how it processes and stores data efficiently.

- Designing for Performance and Scalability

Learn how to design storage solutions in BigQuery that cater to performance needs and scalable growth.

- Storage Options and Trade-offs

Explore different storage options in BigQuery and the trade-offs involved in choosing the right



strategy.

Module 3: Data Modeling and Schema Design

- Principles of Data Modeling

Understand the best practices in data modeling and how they apply to BigQuery environments.

- Designing Schemas for BigQuery

Learn to create efficient and flexible schemas that optimize query performance and data storage.

- Handling Semi-Structured Data

Explore techniques for managing and querying semi-structured data, such as JSON, within BigQuery.

Module 4: Data Partitioning and Clustering

- Introduction to Partitioning in BigQuery

Understand how partitioning can improve query performance and manage large datasets effectively.

- Implementing Clustering for Optimized Queries

Learn how clustering can further enhance query performance by organizing data based on specific columns.

- Best Practices for Partitioning and Clustering

Discover strategies for combining partitioning and clustering to achieve optimal performance and cost efficiency.

Module 5: Data Compression and Optimization Techniques

- Leveraging Data Compression

Explore various data compression techniques in BigQuery to save on storage costs without



sacrificing performance.

- Query Optimization Strategies

Learn to optimize queries in BigQuery for faster execution and better resource management.

- Using Materialized Views

Understand the benefits of using materialized views to improve query performance and streamline data processing.

Module 6: Security and Data Governance

- Securing Data in BigQuery

Examine the security features of BigQuery, including access controls and encryption, to protect sensitive data.

- Implementing Data Governance Policies

Learn how to implement effective data governance practices to ensure data integrity and compliance.

- Auditing and Monitoring BigQuery

Discover tools and techniques for auditing and monitoring BigQuery usage to maintain a secure and efficient environment.

Module 7: Cost Management and Optimization

- Understanding BigQuery Pricing Model

Gain insights into the pricing structure of BigQuery to better manage and predict costs.

- Techniques for Cost Optimization

Explore various strategies to optimize costs while using BigQuery, such as query and storage management.

- Monitoring and Controlling Costs



Learn how to use monitoring tools to track BigQuery usage and implement controls to prevent cost overruns.

Module 8: Real-World Applications and Case Studies

- Case Studies of BigQuery Implementations

Analyze real-world examples of how organizations have successfully implemented BigQuery for their data warehousing needs.

- Lessons Learned and Best Practices

Learn from the experiences of others, focusing on lessons learned and best practices in BigQuery storage design.

- Future Trends in Data Warehousing with BigQuery

Explore emerging trends and future directions in data warehousing and BigQuery's role in the evolving landscape.

Course Conclusion:

The course concludes by reinforcing the key concepts covered and emphasizing the skills acquired in designing efficient storage solutions using BigQuery. Participants will be equipped with the knowledge to implement best practices in their data warehousing projects, ensuring they can leverage BigQuery's full potential for their organization's data storage needs.