

MCSA: SQL 2016 BI Development – Skills Measured

NOTE: The bullets that appear below each of the skills measured are intended to illustrate how we are assessing that skill. This list is not definitive or exhaustive.

NOTE: In most cases, exams do NOT cover preview features, and some features will only be added to an exam when they are GA (General Availability).

Exam 70-767: Implementing a Data Warehouse using SQL

Design, implement, and maintain a data warehouse (35–40%)

Design and implement dimension tables

- design shared and conformed dimensions, determine support requirements for slowly changing dimensions, determine attributes, design hierarchies, determine star or snowflake schema requirements, determine the granularity of relationship by using fact tables, determine auditing or lineage requirements, determine keys and key relationships for a data warehouse, implement dimensions, implement data lineage of a dimension table

Design and implement fact tables

- identify measures, identify dimension table relationships, create composite keys, design a data warehouse that supports many-to-many relationships, implement semi-additive measures, implement non-additive measures

Design and implement indexes for a data warehouse workload

- design an indexing solution; select appropriate indexes; implement clustered, non-clustered, filtered, and columnstore indexes

Design storage for a data warehouse

- design an appropriate storage solution, including hardware, disk, and file layout

Design and implement partitioned tables and views

- design a partition structure to support a data warehouse, implement sliding windows, implement partition elimination, design a partition structure that supports the quick loading and scale-out of data

Extract, transform, and load data (40–45%)

Design and implement an extract, transform, and load (ETL) control flow by using a SQL Server Integration Services (SSIS) package

- design and implement ETL control flow elements, including containers, tasks, and precedence constraints; create variables and parameters; create checkpoints, sequence and loop containers, and variables in SSIS; implement data profiling, parallelism, transactions, logging, and security

Design and implement an ETL data flow by using an SSIS package

- implement slowly changing dimension, fuzzy grouping, fuzzy lookup, audit, blocking, non-blocking, and term lookup transformations; map columns; determine the appropriate transform object for a given task; determine appropriate scenarios for Transact-SQL joins versus SSIS lookup; design table loading by using bulk loading or standard loading; remove extra rows or bad rows by using deduplication

Implement an ETL solution that supports incremental data extraction

- design fact table patterns, enable Change Data Capture, create a SQL MERGE statement

Implement an ETL solution that supports incremental data loading

- design a control flow to load change data, load data by using Transact-SQL Change Data Capture functions, load data by using Change Data Capture in SSIS

Debug SSIS packages

- fix performance, connectivity, execution, and failed logic issues by using the debugger; enable logging for package execution; implement error handling for data types; implement breakpoints; add data viewers; profile data with different tools; perform batch clean-up

Deploy and configure SSIS packages and projects

- create an SSIS catalog; deploy packages by using the deployment utility, SQL Server, and file systems; run and customize packages by using DTUTIL

Build data quality solutions (15–20%)

Create a knowledge base

- create a Data Quality Services (DQS) knowledge base, determine appropriate use cases for a DQS knowledge base, perform knowledge discovery, perform domain management

Maintain data quality by using DQS

- add matching knowledge to a knowledge base, prepare a DQS for data deduplication, create a matching policy, clean data by using DQS knowledge clean data by using the SSIS DQS task, install DQS

Implement a Master Data Services (MDS) model

- install MDS; implement MDS; create models, entities, hierarchies, collections, and attributes; define security roles; import and export data; create and edit a subscription; implement entities, attributes, hierarchies, and business rules

Manage data by using MDS

- use MDS tools, use the Master Data Services Configuration Manager, create a Master Data Manager database and web application, deploy a sample model using MDSModelDeploy.exe, use the Master Data Services web application, use the Master Data Services Add-in for Excel, create a Master Data Management hub, stage and load data, create subscription views

Exam 70-768: Developing SQL Data Models

Design a multidimensional business intelligence (BI) semantic model (25–30%)

Create a multidimensional database by using Microsoft SQL Server Analysis Services (SSAS)

- design, develop, and create multidimensional databases; select a storage model

Design and implement dimensions in a cube

- select an appropriate dimension model, such as fact, parent-child, roleplaying, reference, data mining, many-to-many, and slowly changing dimension; implement a dimension type; define attribute relationships

Implement measures and measure groups in a cube

- design and implement measures, measure groups, granularity, calculated measures, and aggregate functions; define semi-additive behavior

Design a tabular BI semantic model (20–25%)

Design and publish a tabular data model

- design measures, relationships, hierarchies, partitions, perspectives, and calculated columns; create a time table; publish from Microsoft Visual Studio; import from Microsoft PowerPivot; select a deployment option, including Processing Option, Transactional Deployment, and Query Mode

Configure, manage, and secure a tabular model

- configure tabular model storage and data refresh, configure refresh interval settings, configure user security and permissions, configure row-level security

Develop a tabular model to access data in near real time

- use DirectQuery with Oracle, Teradata, Excel, and PivotTables; convert in-memory queries to DirectQuery

Develop queries using Multidimensional Expressions (MDX) and Data Analysis Expressions (DAX) (15–20%)

Create basic MDX queries

- implement basic MDX structures and functions, including tuples, sets, and TopCount

Implement custom MDX solutions

- create custom MDX or logical solutions for pre-prepared case tasks or business rules, define a SCOPE statement

Create formulas by using the DAX language

- use the EVALUATE and CALCULATE functions, filter DAX queries, create calculated measures, perform data analysis by using DAX

Configure and maintain SQL Server Analysis Services (SSAS) (30–35%)

Plan and deploy SSAS

- configure memory limits, configure Non-Union Memory Architecture (NUMA), configure disk layout, determine SSAS instance placement

Monitor and optimize performance

- monitor performance and analyze query plans by using Extended Events and Profiler, identify bottlenecks in SSAS queries, monitor processing and query performance, resolve performance issues, configure usability limits, optimize and manage model design

Configure and manage processing

- configure partition processing; configure dimension processing; use Process Default, Process Full, Process Clear, Process Data, Process Add, Process Update, Process Index, Process Structure, and Process Clear Structure processing methods; configure Parallel, Sequential, and Writeback processing settings

Create Key Performance Indicators (KPIs) and translations

- configure KPI options, including Associated measure group, Value Expression, Goal Expression, Status, Status expression, Trend, Trend expression, and Weight; create KPIs in multidimensional models and tabular models; create and develop translations