

BI4Dynamics BC Cloud Installation Requirements

Last update: May 2021

BI4Dynamics BC: 8.2

Document version: 1.3

Table of Contents

- 1 INTRODUCTION 2
 - 1.1 Implementation options 2
- 2 Implementation on Azure SQL Virtual Machine 3
 - 2.1 Permissions 3
 - 2.2 Azure SQL Virtual Machine - hardware 3
 - 2.3 Azure SQL Virtual Machine - software 4
 - 2.4 Other resources 4
- 3 On-Premises implementation 5
 - 3.1 Permissions 5
 - 3.2 On-Premises implementation - hardware 5
 - 3.3 On-Premises implementation - software 6
 - 3.4 Other resources 6
- 4 Best practice and recommendations 7
 - 4.1 SQL Server features and settings 7
 - 4.2 Faster processing and querying 8

1 INTRODUCTION

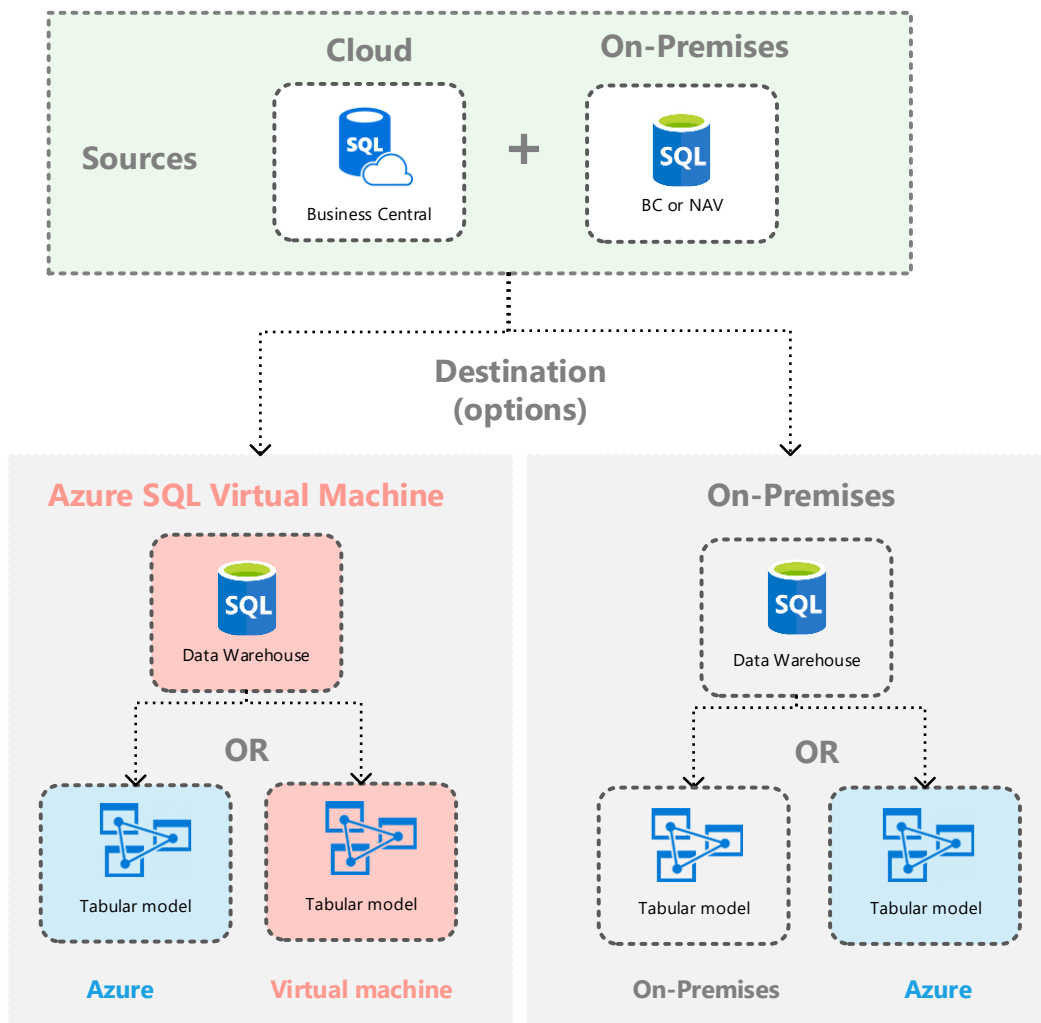
1.1 Implementation options

Source

BI4Dynamics can join **BC Cloud** with your legacy **BC or NAV** in one fully integrated BI solution. First selected source in BI4Dynamics must be BC Cloud. BC or NAV can be added optionally.

Destination

- Data warehouse can be implemented on **Azure SQL Virtual machine** or **On-Premises**.
 - Use VM when company strategy is cloud
 - Use OP when hardware is available on-premises
- Analytics can be implemented on the **same PC** (VM or On-Premises) or as **Azure service**.
 - Use same PC when PC is available or costs for PC are low
 - Use Azure when cloud performance, scalability, and flexibility matters



4 implementation option for BC Cloud

2 Implementation on Azure SQL Virtual Machine

2.1 Permissions

BC Cloud


- BC Cloud credentials: **Dynamics 365 administrator** role and **Application Developer** for app registration
- Same BC Extensions must be installed on Production and on Sandbox.

Azure

- Be local **administrator** on the Azure VM where BI4Dynamics is installed.
- **Azure Storage** info: Subscription, Resource group, Account name, Container name and Blob Key

2.2 Azure SQL Virtual Machine - hardware

Use Azure SQL option, that is VM with SQL 2019 Standard edition preinstalled.



Do not install Virtual machine, as it requires a separate SQL server installation.

Azure SQL Virtual Machine Size

VM processing time and consequently size mostly depend on BC database size. Here are 2 examples:

BC Cloud database = 5 GB

| VM Size | Family | vCPUs | RAM (GiB) | Data disks | Max IOPS | Temp Storage (GiB) | Premium disk | Cost per hour | Cost per month (based on 20 h) |
|---------|-----------------|-------|-----------|------------|----------|--------------------|--------------|----------------|--------------------------------|
| B4ms | General purpose | 4 | 16 | 8 | 2880 | 32 | Supported | € 0,51 \$ 0,61 | € 10 \$ 12 |
| D4s_v3 | General purpose | 4 | 16 | 8 | 6400 | 32 | Supported | € 0,70 \$ 0,78 | € 14 \$ 16 |

BC Cloud database = 20 GB

| VM Size | Family | vCPUs | RAM (GiB) | Data disks | Max IOPS | Temp Storage (GiB) | Premium disk | Cost per hour | Cost per month (based on 20 h) |
|---------|-----------------|-------|-----------|------------|----------|--------------------|--------------|----------------|--------------------------------|
| B8ms | General purpose | 8 | 32 | 16 | 4320 | 64 | Supported | € 1,03 \$ 1,13 | € 21 € 23 |
| D8s_v3 | General purpose | 8 | 32 | 16 | 12800 | 64 | Supported | € 1,39 \$ 1,55 | € 28 € 31 |

Calculations are based DW processing time of 1 hour, 20 times each month. Latest pricing is [here](#).

Attached disks: Azure SQL VM comes with 128 GB of C drive; OS and SQL files are using 40 GB. 80 GB is enough for DW processing of BC databases up to 30 GB. If BC database is bigger, add disk. Smaller disks are very slow, while bigger disks are much faster, so you may choose larger disk because of speed, not size.

Azure VM Setup: VM should be configure with **Public IP** and **DNS name**

2.3 Azure SQL Virtual Machine - software

Operating system: Azure SQL VM comes with Windows Server 2019.

SQL Server (software) - following features must be installed.

- SQL Server Database engine
- SQL Server Analysis Services Tabular (when implementing Analysis Services locally)
- Integration services
- PolyBase (PolyBase Query Services and Java connector for HDFS data sources)

PowerShell, NET Framework version 4.7.2 or higher

Azure CLI: set of commands that create and manage Azure resources, available [here](#).

2.4 Other resources

Every implementation

Azure Storage Account: Storage Account with Container used for storing BC tables that are exported from BC.

Azure Docker: specialized cloud engine that manages the export of BC tables to Blob Storage

Implementation option is Azure Analysis Services

Azure Analysis Services: in memory database available for user queries by any BI client (Excel, Power BI)

On-Premises Gateway: provides bridge data transfer between on-premises data and Microsoft cloud services

3 On-Premises implementation

Implementation on physical or virtual computer on-premises is the same.

3.1 Permissions

Local server

Permissions for SQL installation and processing: all accounts to be combined into one (1) Windows domain account:

- Be (local) administrator on the BI server where BI4Dynamics is installed.
- Create database and linked server permissions on SQL Server where the BI4Dynamics DW will be deployed.
- Granted Server role on Analysis Services where the BI4Dynamics tabular model will be deployed. Additionally, the Analysis Server service should run under BI4Dynamics domain account.
- SQL Agent job Service user account needs all the above-stated permissions.

BC Cloud

- BC Cloud credentials: **Dynamics 365 administrator** role and **Application Developer** for app registration
- Same BC Extensions must be installed on Production and on Sandbox.

Azure

- **Azure Storage** info: Subscription, Resource group, Account name, Container name and Blob Key

3.2 On-Premises implementation - hardware

Dedicated infrastructure for BI stack is recommended.

CPU: 2 – 4 processors at least 2GHz; more cores bring more parallelism.

RAM: used for processing data warehouse and for storing Analysis services database (tabular option).

Therefore, Tabular option requires more memory than with Multidimensional.

| Requested RAM (GB) | Analysis Services option | |
|--------------------|--------------------------|------------------|
| | Tabular | Multidimensional |
| < 50 GB | 20 | 16 |
| 100 to 250 GB | 40 | 32 |
| 250 to 500 GB | 75 | 64 |

DISK (SSD):

- 2 x size of all source databases, when using standard **Rowstore** storage
- 1 x size when using **Columnstore** feature

3.3 On-Premises implementation - software

Operating system

When using SQL 2016 and SQL 2017: Windows Server 2012, 2012 R2, 2016, 2019 or Win 8,10

When using SQL 2019: Windows 10 TH1 1507 or greater, Windows Server 2016 or greater

SQL Server database

SQL Server **Standard** Edition version **2016** or above. Recommended version **2019** due to new features and better user experience (showing dimensional: fact relationship as one of them).

Analysis services (when implementing Analysis Services locally)

- **Tabular option:** SQL version **2017** and above
- **Tabular option supporting Account Schedules functionality:** SQL version **2019**
- **Multidimensional option:** SQL version **2016** and above

SQL Server features

- Integration services
- PolyBase (PolyBase Query Services and Java connector for HDFS data sources)

Other programs

PowerShell, NET Framework version 4.7.2 or higher

Azure CLI: set of commands that create and manage Azure resources, available [here](#).

3.4 Other resources

Every implementation

Azure Storage Account: Storage Account with Container used for storing BC tables that are exported from BC.

Azure Docker: specialized cloud engine that manages the export of BC tables to Blob Storage

Implementation option is Azure Analysis Services

Azure Analysis Services: in-memory database available for user queries by any BI client (Excel, Power BI)

On-Premises Gateway: provides bridge data transfer between on-premises data and Microsoft cloud services

4 Best practice and recommendations

4.1 SQL Server features and settings

SQL Server editions

All BI4Dynamics features run on **Standard** SQL Edition. It also runs on Enterprise edition. **Per instance Limitation** of SQL Server Standard Edition:

- SQL Server engine Buffer pool 128 GB
- Analysis Services Tabular 16 GB

When data size requires more hardware resources, BI project should be **split into two instances**:

- **Finance** (Finance, Receivables, Payables, Fixed Assets, Bank Account)
- **Operations** (Sales, Inventory, Purchase, Retail, Manufacturing, Warehouse, Service)

SQL Server settings

Volumes, drives: Separate volume/drive for Data (#1) and Log & Temp (#2).

Tempdb: set 4 or 8 data files and 1 log to avoid the GAM/PFS page contention issue described [here](#).

Collation: the collation of the data warehouse server should match the collation of the BC database.

UAC should be **disabled**, or application needs to be installed outside Program Files folder.

Network availability:

- **Online servers:** The server is connected to the Internet with port 80 opened (if the server is under proxy, the application needs access to our authorization web service)
- **Offline servers:** special offline license is required; please contact sales@bi4dynamics.com

Integration services

Integration Services can speed up the processing Stage and Data warehouse from 50 to 150% by running stored procedures in parallel. If not installed or selected processing will be much slower.

PolyBase

Install PolyBase Query Services and Java connector for HDFS data sources. These features are a part of SQL server used for (among others) loading Blob storage files.

Columnstore option

BI4Dynamics supports Columnstore storage in staging (schema = stage) and Data warehouse area (schema = fact). Columnstore can achieve 10X compression. Dim schema is not supported with Columnstore as dimensional table are usually smaller and compression does not help performance. Loading may be increased by 10-20% due to building Columnstore indexes.

4.2 Faster processing and querying

Fast data warehouse processing

- **Fast disks:** disk speed is the most important part as data warehouse reads and writes a lot of data
- **Enough RAM:** RAM amount under the recommended size will significantly affect the processing speed.
- **Good CPU** with many cores

Fast Analysis Services processing and querying

While Multidimensional requirements are like Data warehouse requirements (fast disks, many cores for parallelism), Tabular requires fast CPU and fast memory. These are hardware components that affect performance for Tabular engine:

- **Fast CPU** speed is better and will affect query time (interaction between front end tools (Excel and Power BI) and Tabular engine when a user is analyzing data
- **Fast memory:** fast memory can be more expensive, but it is worth the investment
- Memory size should be 2X of sum of size of all Tabular databases; If there is not enough memory that data will be partly written to the disk and performance will be slow