kapacity

Kapacity Power BI Monitoring Solution Documentation

+45 7030 2919
info@kapacity.dk
www.kapacity.dk

København
Gærtorvet 3, 2. sal
1799 København V

Aarhus
Mariane Thomsens Gade 2F, 8. sal
8000 Aarhus C

Kolding
Trindholmsgade 4, 3. sal
6000 Kolding

kapacity

+45 7030 2919
info@kapacity.dk
www.kapacity.dk

København
Gærtorvet 3, 2. sal
1799 København V

Aarhus
Mariane Thomsens Gade 2F, 8. sal
8000 Aarhus C

Kolding
Trindholmsgade 4, 3. sal
6000 Kolding

kapacity

Dokumentinformation

Projektnavn: Monitoring

Version: 0.2

Forfatter: Marcel Jensen

Revisionshistorie

Version	Dato	Ændret af	Ændring
0.1	24/9	Marcel Jensen	Initial document
0.2	15/11	Niels Hundahl	Minor corrections
0.3	16/11	Marcel Jensen	Added architecture

+45 7030 2919
info@kapacity.dk
www.kapacity.dk

København
Gærtorvet 3, 2. sal
1799 København V

Aarhus
Mariane Thomsens Gade 2F, 8. sal
8000 Aarhus C

Kolding
Trindholmsgade 4, 3. sal
6000 Kolding

Contents

+45 7030 2919
info@kapacity.dk
www.kapacity.dk

København
Gærtorvet 3, 2. sal
1799 København V

Aarhus
Mariane Thomsens Gade 2F, 8. sal
8000 Aarhus C

Kolding
Trindholmsgade 4, 3. sal
6000 Kolding

1. About the solution

1.1 What is the purpose of the solution?

The purpose of Kapacity Power BI Monitor Soltion is to provide insight into the usage of Power BI across a whole tenant, over longer time intervals than provided in standard solutions, and provide insight into the ETL jobs in a Azure Datawarehouse orchestrated by Azure Data Factory (ADF).
It also provides insight into the analytical models in Azure Analysis Services (AAS).

1.2 How it works – Overall architecture

1.2.1 Architecture overview

Below is a visual overview of the architecture. For futher description, see section 1.2.2

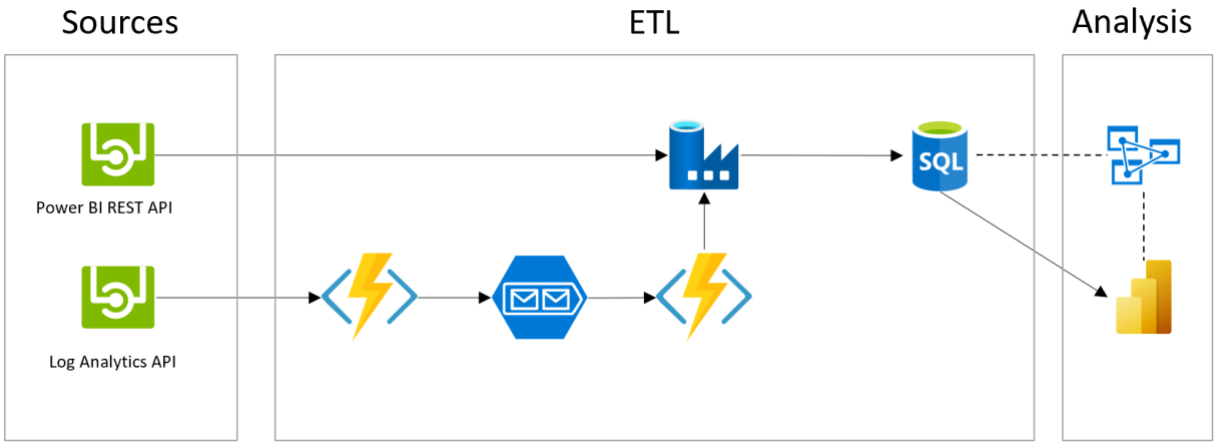

1.2.2 Description of architecture

- Data is extracted using the Power BI Management API and a Log Analytics workspace.
- API calls to the Power BI API, are handled by ADF pipelines.
- Data from Log Analytics are extracted using an Azure Function App. This contains 2 functions, one controller function to handle what data to get, and one worker function to retrieve data.
  The controller function passes messages to a storage message queue, which triggers the worker function once for each table to retrieve data from.
- All data is sent to the Azure SQL Database, which then handles the transform and load into a datawarehouse.
  There is an option to present the data in an Analysis Services tabular model.
- Everything is orchestrated by Azure Data Factory.
- Optional Analysis Services to analyse data.

2. Frontend

The solution features an out-of-the-box Power BI report, with a default design. The report is found in 2 different versions, one with an imported dataset, and one with a dataset based on an Azure Analysis Services-model.

Both reports contain a page for each of the monitored areas, as well as a frontpage.
The pages are:

1. Power BI Report Usage
2. Model Usage
3. ETL Performance

Frontpage

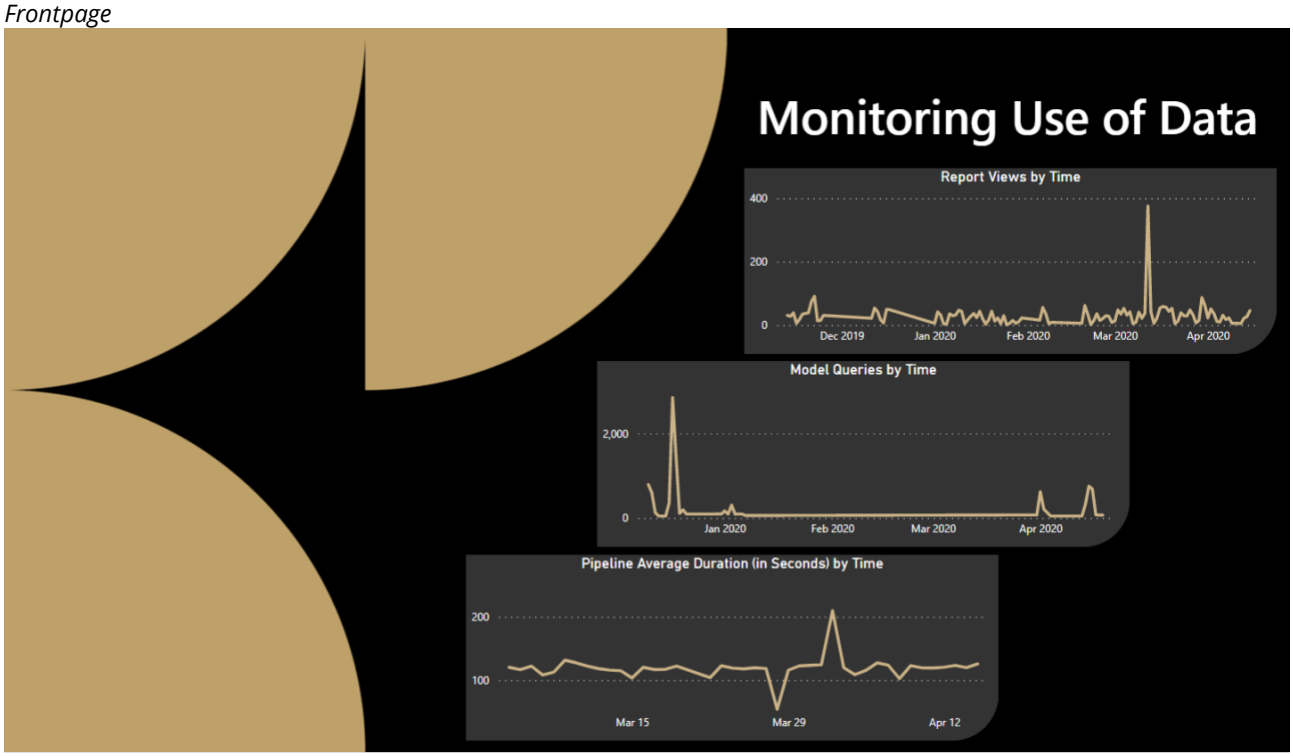

2.1 Import model report

The import model Power BI report, is based directly on the underlying Azure SQL Database.
Everything is modelled in Power BI, based on the following data:

- **Power BI Audit Log**: contains data about Power BI usage
- **Azure Analysis Services Query Log**: contains data about Data Model usage
- **Azure Data Factory Trigger/Pipeline/Activity Runs**: contains data about the ETL performance

In addition to the data itself, a number of measures based on this data is provided out-of-the-box.

2.2 Azure Analysis Services model report

This model contains the same data as the Import Model, but instead of being moddeled directly in Power BI, the model is made in Azure Analysis Services.
Both are modelled exactly the same.

2.3 Measures in the default report

The default report comes with some measures used for displaying a bit of status on the report:

- **Average Duration PBI**: A filtered version of the cube measure Average Duration, which only includes queries coming from Power BI Desktop or from the Power BI service.
- **Average Duration Not PBI**: A filtered version of the cube measure Average Duration, which only includes queries not coming from Power BI Desktop or from the Power BI service.
- **First AAS Log Observation DateTime**: The lowest timestamp of the cube queries within the given time filter. Can be used to show current data time-span.
- **Last AAS Log Observation DateTime**: The highest timestamp of the cube queries within the given time filter. Can be used to show current data time-span.
- **First PBI Log Observation DateTime**: The lowest timestamp of the Power BI user actions within the given time filter. Can be used to show current data time-span.
- **Last PBI Log Observation DateTime**: The highest timestamp of the Power BI user actions within the given time filter. Can be used to show current data time-span.

3. Backend

3.1 Overview of required Azure resources

- **Azure SQL Database**
- **Azure Function**
- **App Service plan (tied to Azure Function)**
- **Azure Storage**
- **Azure Data Factory**
- **Azure Key Vault**
- **Azure Analysis Services (optional)**

3.1.1 Azure SQL Database

Used to as the base of the data, provided for reporting. Also handles all transformation of data used.

3.1.2 Azure Function

Used to extract data from a Log Analytics workspace, to provide log data from Azure Analysis Serviced and Azure Data Factory.

3.1.3 App Service Plan

Used to provide compute power to the Azure Function

3.1.4 Azure Storage

Provides storage and message queue, for use in the Azure Function

3.1.5 Azure Data Factory

Acts as the orchestrator for the update of data in the solution, as well as extracting data from the Power BI Management API.

3.1.6 Azure Key Vault

Used to provide connection strings, secure storage of usernames and passwords, for the other Azure resources.

3.1.7 Azure Analysis Services (optional)

Used to provide the data for the report, if chosen. Optional, as there is also an option for a standalone import model report.

3.2 Data sources

3.2.1 Power BI Management API

Provides the solution with information about the Power BI Tenant, which is monitored, such as metadata about workspaces, reports, datasets, etc., as well as the Power BI usage log.

3.2.2 Log Analytics Workspace

Provides the solution with log data from the Azure Analysis Services and Azure Data Factory, which is monitored.

4. Setup of required access

The following is a brief overview of the required access, to enable the solution to work. For a more detailed description, see the document **Setup for Monitoring**.

4.1 Enabling Azure resource diagnostic logging

In order to provide the solution with log data from Azure Analysis Services and Azure Data Factory, **diagnostic logging** must to be enabled on these resources.
This is set up on the individual resources in the Azure Portal.

4.2 Azure Resource Group and access

In order to install the solution, we require **access to create Azure resources in an Azure Resource Group**.
This is handled automatically by a deployment script, that requires a **Service Principal** with privileges to create and mange Azure resources (Contributor access).

Two user assigned **Managed Identities** in the Azure tenant, is required for use in Azure Data Factory and the Azure function.
These Managed Identities need to be provided with access in the data sources, as described below.

4.3 Access to data sources

4.3.1 Power BI Management API

Requires a User assigned managed identity, which has been given **access to read data from the Power BI Management API**.
Access is provided through the Power BI portal.

4.3.2 Log Analytics Workspace

The Managed Identity which is used for the Azure Function, need to be provided with **access to the Log Analytics Workspace**, which contains the logs from Azure Analysis Services and Azure Data Factory.

4.4 Power BI Workspace

To publish the report, regardless of using import model or AAS model report, a Power BI Workspace with contributor access is required.

5. Known limitations

Depending on which version of the solution is implemented, there might be limitations to the amount of data monitored.

5.1 Import model report

The import model version of the report, is limited to **10 GB** of data, on a shared Power BI capacity tenant.
If running Premium Capacity, the above limitation is removed.

5.2 Initial Power BI workspaces monitored

The solution can only get data on a maximum of **5000** Power BI workspaces, in the initial extract of data.
After this, all changed workspaces since the last extract, will be updated

5.3 Data amounts

The amount of data, subsequently the number of days can be monitored in the report, is affected by the choice of pricing tier on the Azure resources used for the solution, mainly the Azure SQL Database and optionally the Azure Analysis Services.

6. Cost of running the solution

- Azure resources
  - SQL DB for hosting log-data - price dependent on pricing tier
  - SQL Analysis Services for hosting Analytic Model (optional)
  - Azure Functions -1 mio. free executions pr. month on “Consumption plan”
- Implementation
  - Solution set-up and test (x consultancy hours)
  - Customized report development (if requested)
- Updates
  - The solution is implemented as-is. If Azure services change or client’s data platform configurations changes in a way that requires changes to the Monitor solution, agreements regarding assistance from Kapacity can be made.

+45 7030 2919
info@kapacity.dk
www.kapacity.dk

København
Gærtorvet 3, 2. sal
1799 København V

Aarhus
Mariane Thomsens Gade 2F, 8. sal
8000 Aarhus C

Kolding
Trindholmsgade 4, 3. sal
6000 Kolding