SQL KONFERENZ 2026

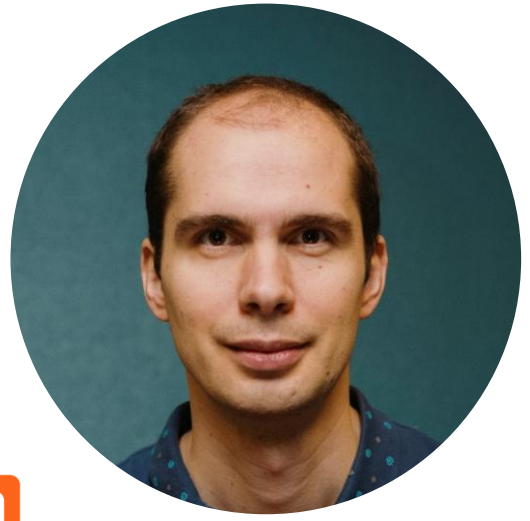


The Power of F2

What You can (and can’t) do with Microsoft Fabric’s smallest capacity

Speaker info

- Budapest, Hungary
- BI Architect @ VISUAL LABS



SQL KONFERENZ 2026

Fabric Adoption
Where Are You on the journey?



1. I've used MS Fabric at least once (trial, demo or real project).
2. I've deployed a Fabric solution that other users rely on.
3. I'm familiar with the difference between Pay-as-You-Go and reserved capacity.
4. I have experienced Fabric slowing down or being throttled due to capacity limits.
5. I've personally worked with F2 capacity.
6. I currently have an F2 capacity running in production.

+1 I have experience and already built solutions relying on ADF or Synapse.

AGENDA



1. 🔍 **Fabric and F2 high level intro**
2. 📈 **When F2 works well:** real-world usage patterns
3. ⚖️ **F2 vs other Azure services**
4. 🛠️ **Building on F2:** development reality and team pitfalls
5. 🚧 **Where F2 hits its limits**
6. 🧱 **Architecture recommendations for F2**

SQL
KONFERENZ
2026



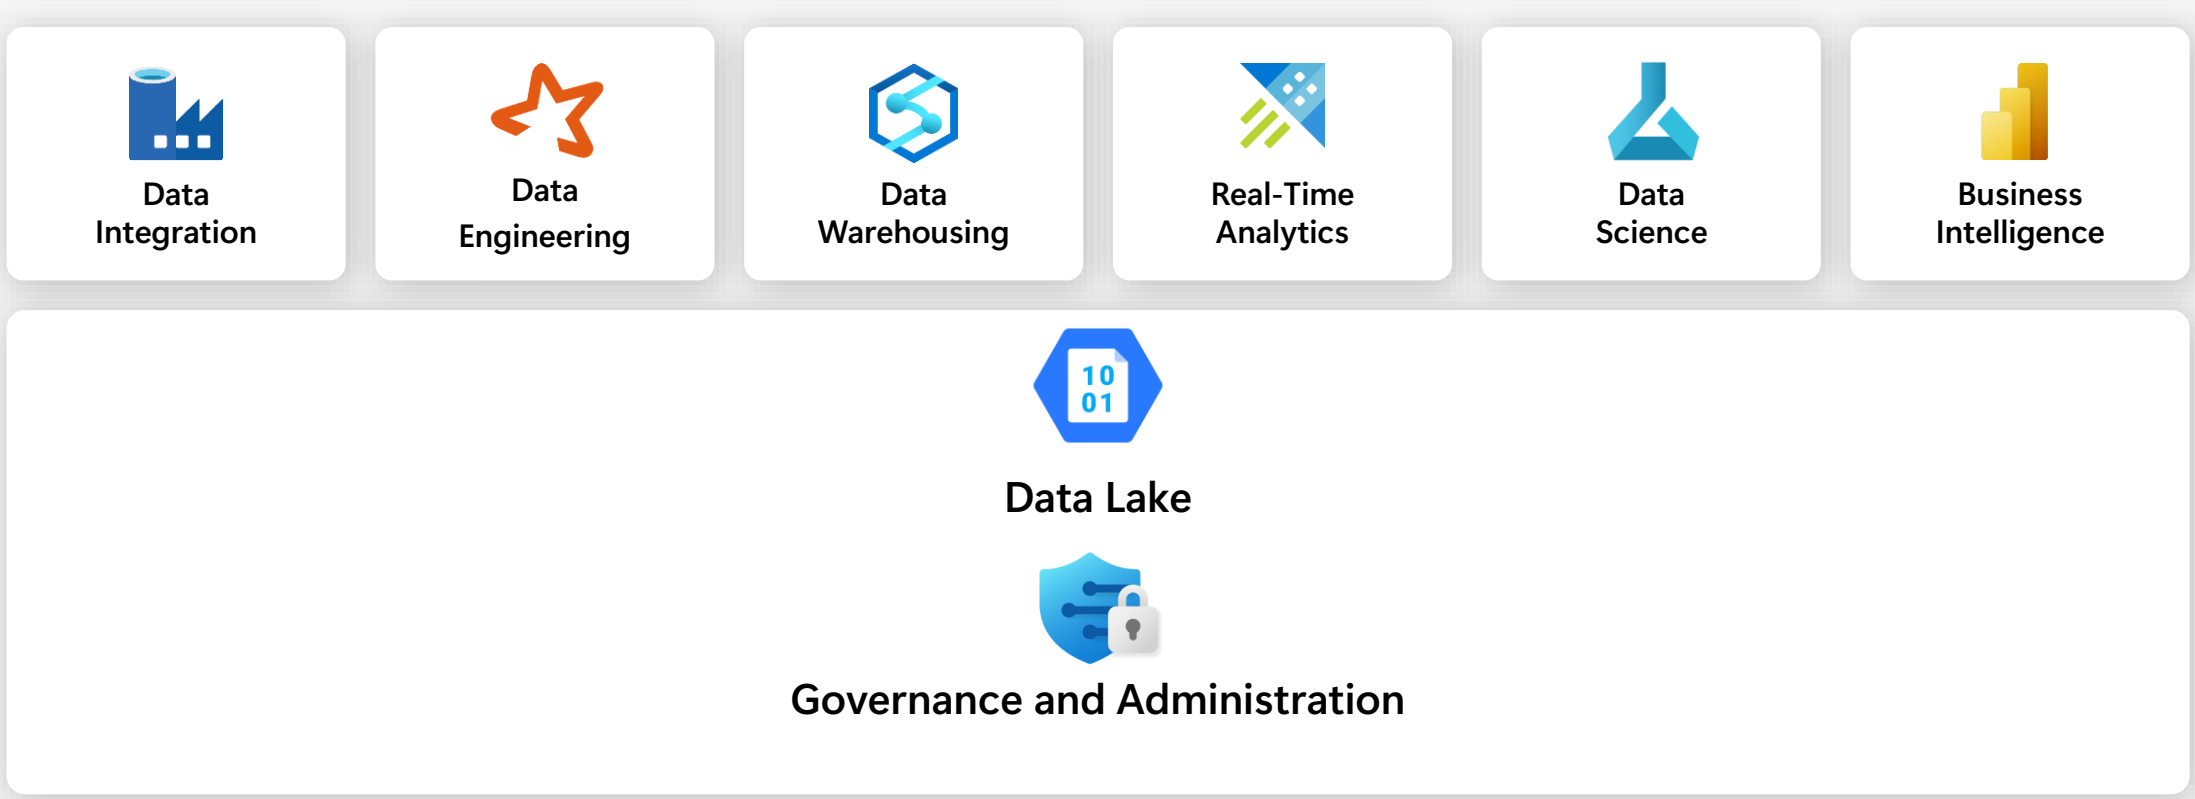
🔍 Fabric and F2 Explained

Analytics today – predictable pattern



Slide credit: Microsoft

Analytics today
MS got you covered already



Slide credit: Microsoft

Analytics today – But it's too complex



Many Products

Different Experiences 

Proprietary and Open

Dedicated and Serverless

PaaS and SaaS

Different Business Models 

Steep Learning Curves 

Deep Expertise Needed

High Integration Effort 


Purview


Power BI


Kusto


Data Factory


Azure AI


Synapse DW


Synapse Spark

Slide credit: Microsoft

Fabric is/might be/can be the SaaS solution



OneLake | Data Factory | Synapse Data Engineering | Synapse Data Warehouse | Synapse Data Science | Synapse Real-Time Analytics | Power BI | Data Activator

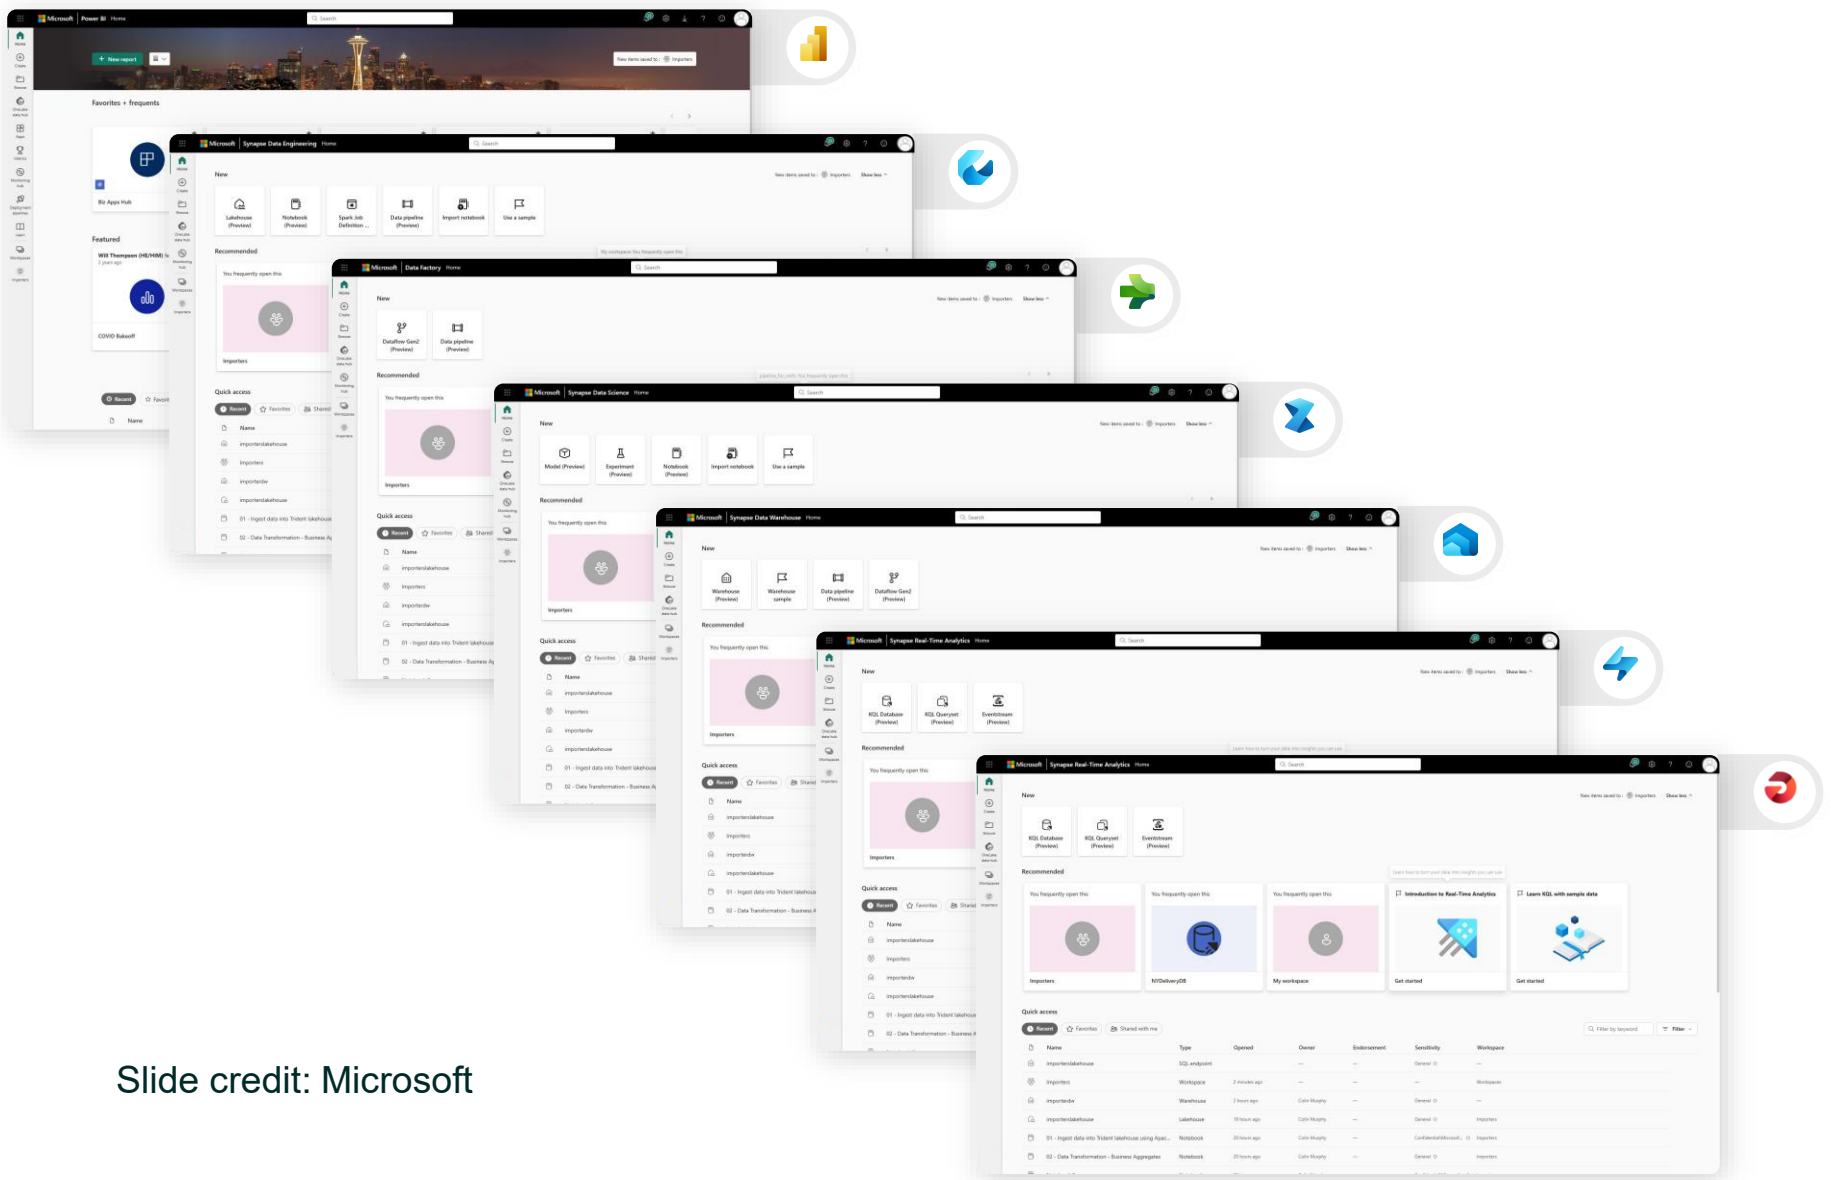
Lake-centric and open | Role-specific tools | Empower every Microsoft 365 user | Persistent governance and security

Single...

- Onboarding and trials
- Sign-on
- Navigation model
- UX model
- Workspace organization
- Collaboration experience
- Data Lake
- Storage format
- Data copy for all engines
- Security model
- CI/CD
- Monitoring hub
- Data Hub
- Governance & compliance

Slide credit: Microsoft

Unified, familiar & intuitive



Slide credit: Microsoft

SQL KONFERENZ 2026

Basics - Fabric capacities



SQL KONFERENZ 2026

1. 💳 What exactly you pay for
2. ⚖️ How much you can use at once
3. ⚙️ What consumes capacity

SQL KONFERENZ 2026

💳 When and how you pay



Pay-as-You-Go



Reserved



! The bill is growing even if you are not „using” Fabric !

SQL KONFERENZ 2026

Billed based on uptime not on resources used

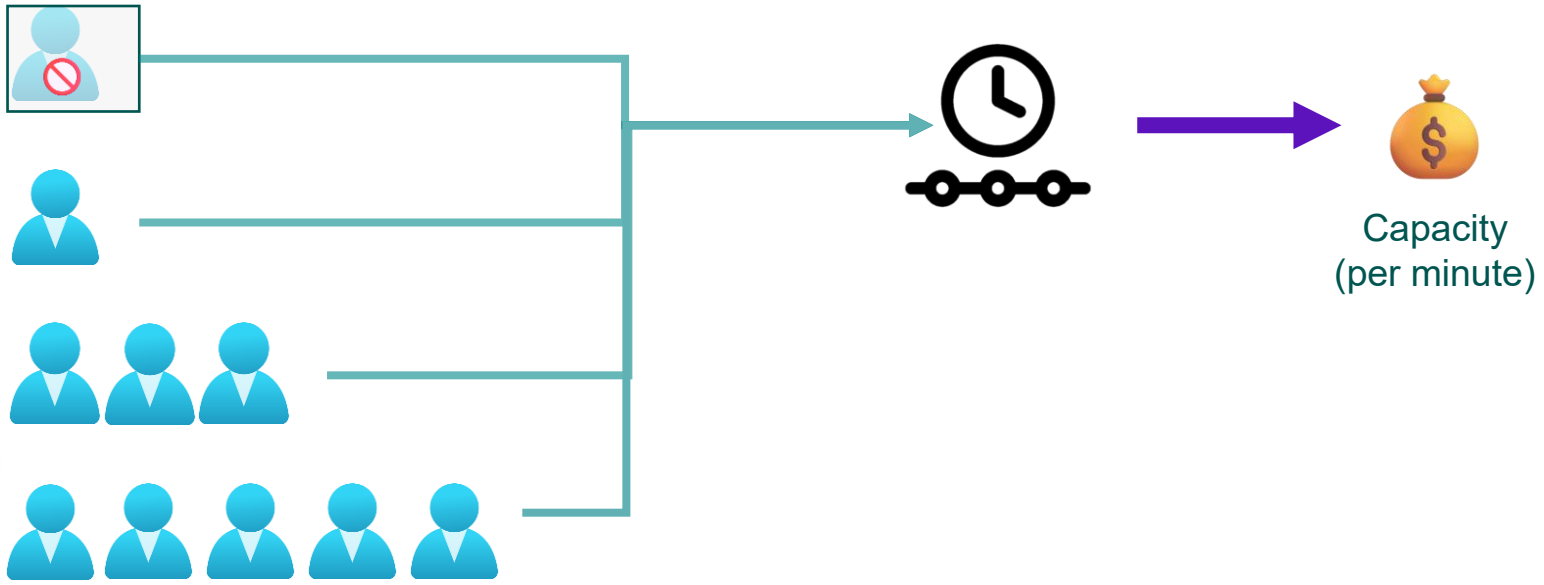




>

SQL KONFERENZ 2026

Billed based on uptime not users/WSs





SQL KONFERENZ 2026

💳 Not billed based on # of created items





SQL KONFERENZ 2026

⚖️ How much you can use at once

SKU	CU / sec	Max memory (GB)	Live connection / sec	Max memory per query (GB)	Model refresh parallelism	Direct Lake rows per table (million)
F2	**2**	**3**	**2**	**1**	**1**	**300**
F4	4	3	2	1	2	300
...	...	...	...	...	...	...
F64	64	25	30	10	40	1500
...	...	...	...	...	...	...
F2048	2048	400	960	40	1,280	24,000



SQL KONFERENZ 2026

How much CUs you can use?



It depends

Current + future
workload (type & complexity)

F2 gives you the least amount of resources

What consumes capacity

Interactive operations

Users expect immediate reponse

Examples:

- **Direct user actions (opens, clicks)**
- **DAX queries generated by report**
- **SSMS**



Background operations

Users are expected to come back to check the result

Examples:

- **Dataflow refresh**
- **Semantic model refresh**
- **Scheduled jobs**



SQL KONFERENZ 2026

SQL
KONFERENZ
2026



Questions before next section?

SQL
KONFERENZ
2026



📈 When F2 works well

Real-World examples

📈 Use case – local FMCG

💰 €25+ M

🔗 Legacy ERP

🎯 Power BI reports: vendor and product performance

▦ 8 input tables

🏁 Medallion, central datastore, 2 semantic models & 3 reports

SQL
KONFERENZ
2026



Guess the avg monthly bill!

📈 Use case – global enterprise



💰 €150+ M

🔗 Legacy ERP systems and D365 F&O

🎯 Power BI reports & Centralized data platform

▦ 180 input tables from D365 (Synapse Link) & 10+ external

Disclaimer

🏁 Medallion, CDC, multiple semantic models & reports

F2 can be perfectly fine



- Predictable, scheduled workloads
- Low or no concurrency requirements
- Batch-oriented processing patterns with „quite windows”
- Clear separation between DEV and PROD
- Separated Fabric and Power BI environment – you need to have PRO license(s) anyway

SQL
KONFERENZ
2026

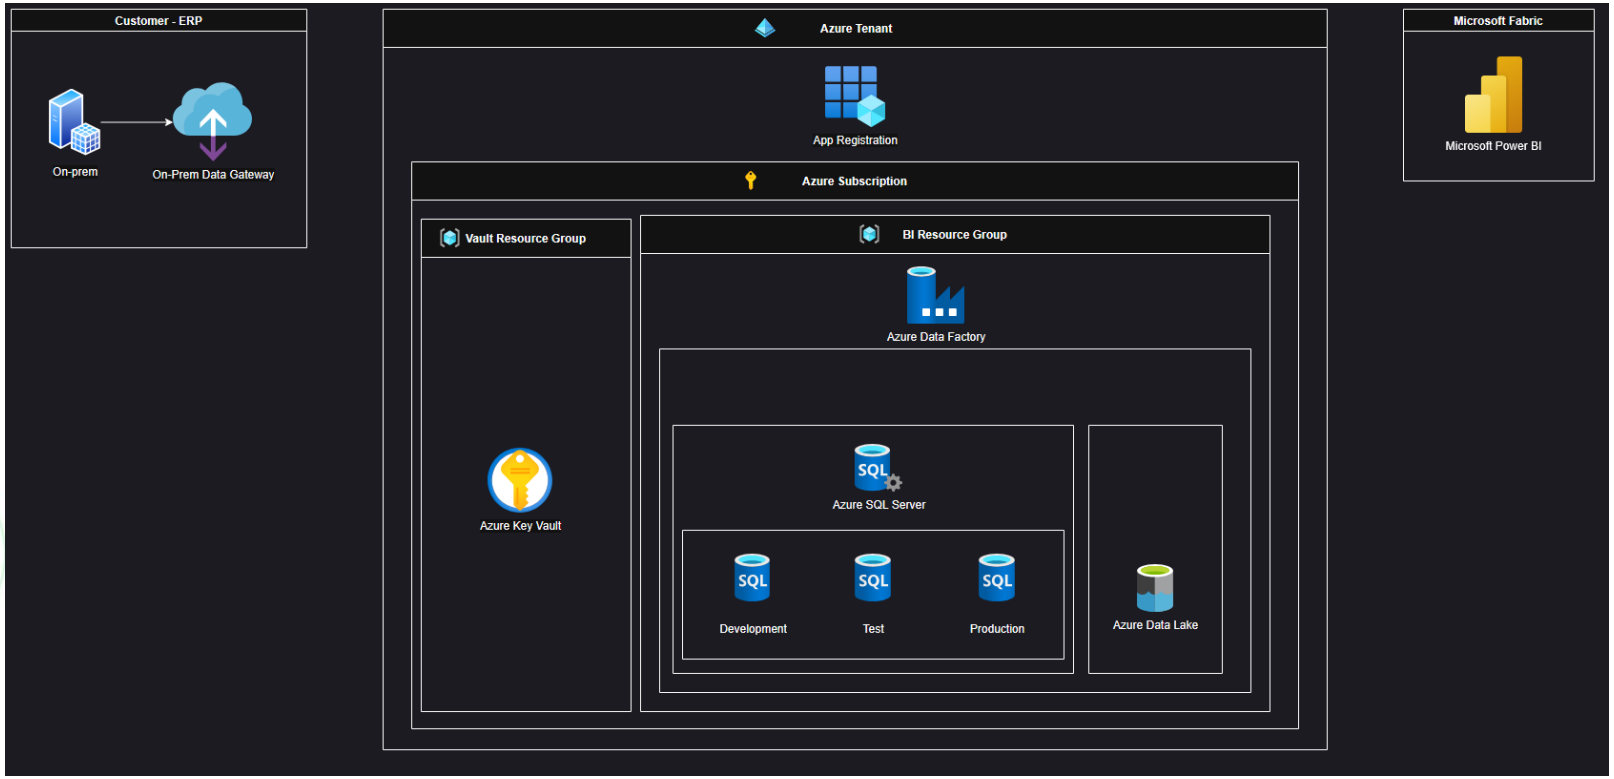


⚖️ F2 vs other Azure services

SQL KONFERENZ 2026

High level architecture on Azure





SQL KONFERENZ 2026

Consideration when choosing Fabric over Azure (all SKUs)

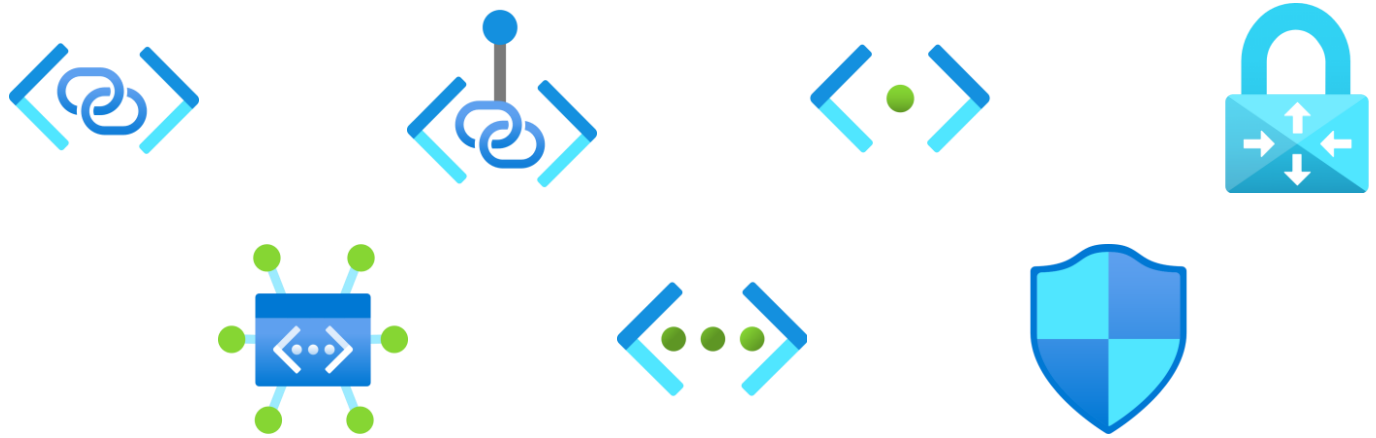


- Fabric billing is still easier to explain
- Unified UI & UX can be a huge advantage (users & developers)
- Better for smaller, code-first porojects
- It can be (see example) but not necessarily cheaper or not significantly cheaper - *Take development and maintenence cost into consideration!*

SQL KONFERENZ 2026

Unannounced or still-missing networking features can make classic Azure components the safer choice





SQL
KONFERENZ
2026



1st Demo Scenario

Fabric vs Azure resources

SQL KONFERENZ 2026

Set up



- Contoso 1M dataset
- Fact tables - duplicated 12x time
- Incremental load (only facts) – every 3 hours
- Approx **3M** rows / load to 🥉 layer
- 🥈 🥇 layers need to handle increasing data volume
- Cleanse, transform, aggregate steps

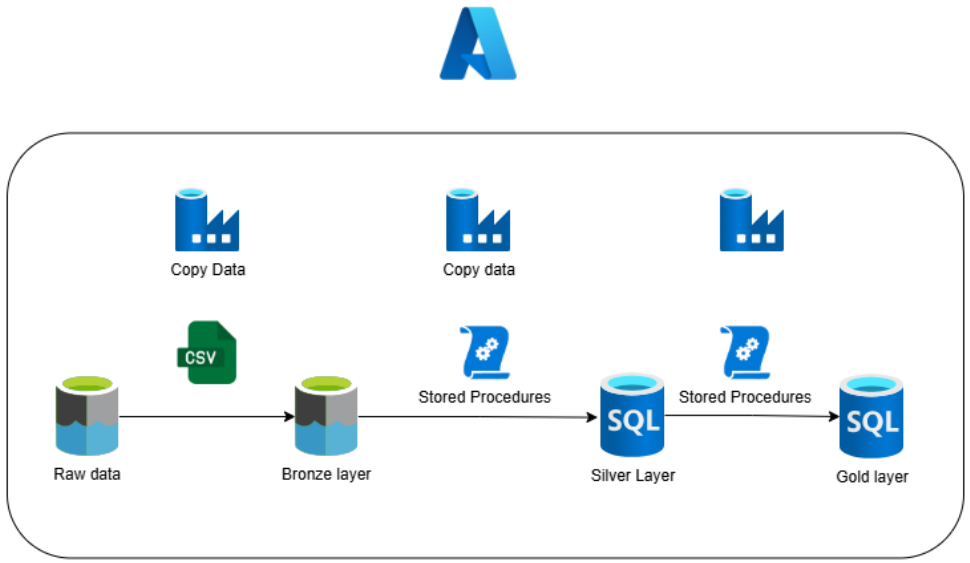
SQL KONFERENZ 2026

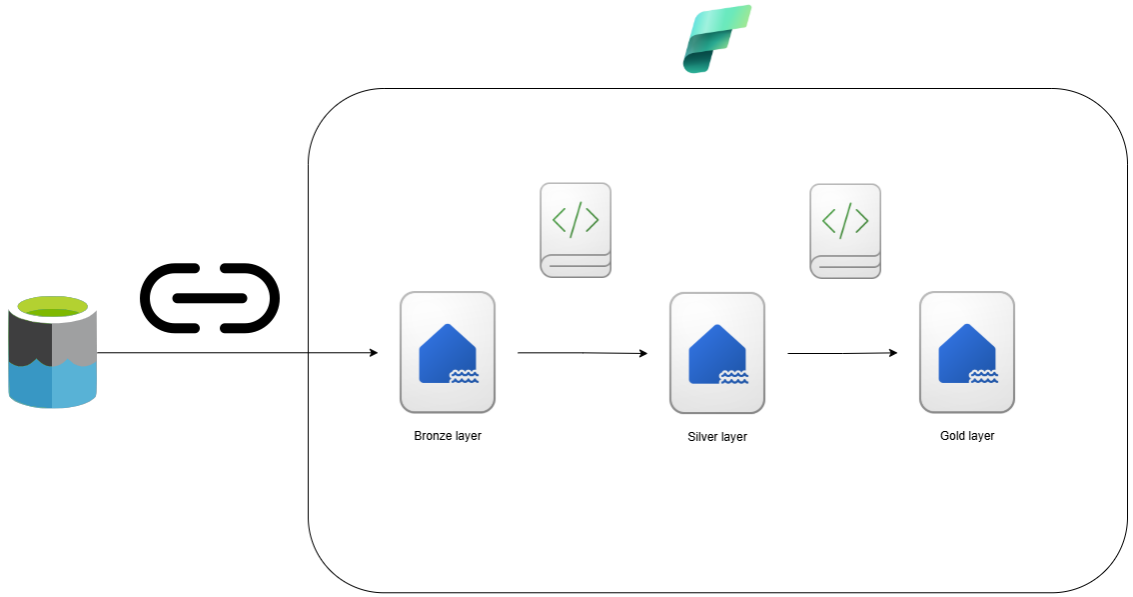
Set up



- Silver: joins facts & dims
- Add calculated columns
- Convert everything to USD
- Create 8 different gold tables from “executive KPI” to product performance or customer churn risk
- Run OPTIMIZE and VACUUM
- Write logs

1st Demo scenario
Azure vs Fabric



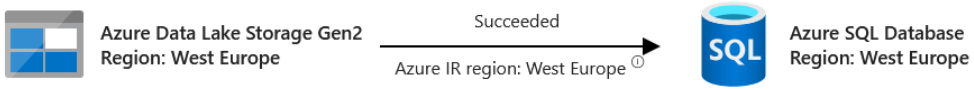


SQL KONFERENZ 2026

Copy duraiton – ADF Copy data activity (ADLS to SQL)

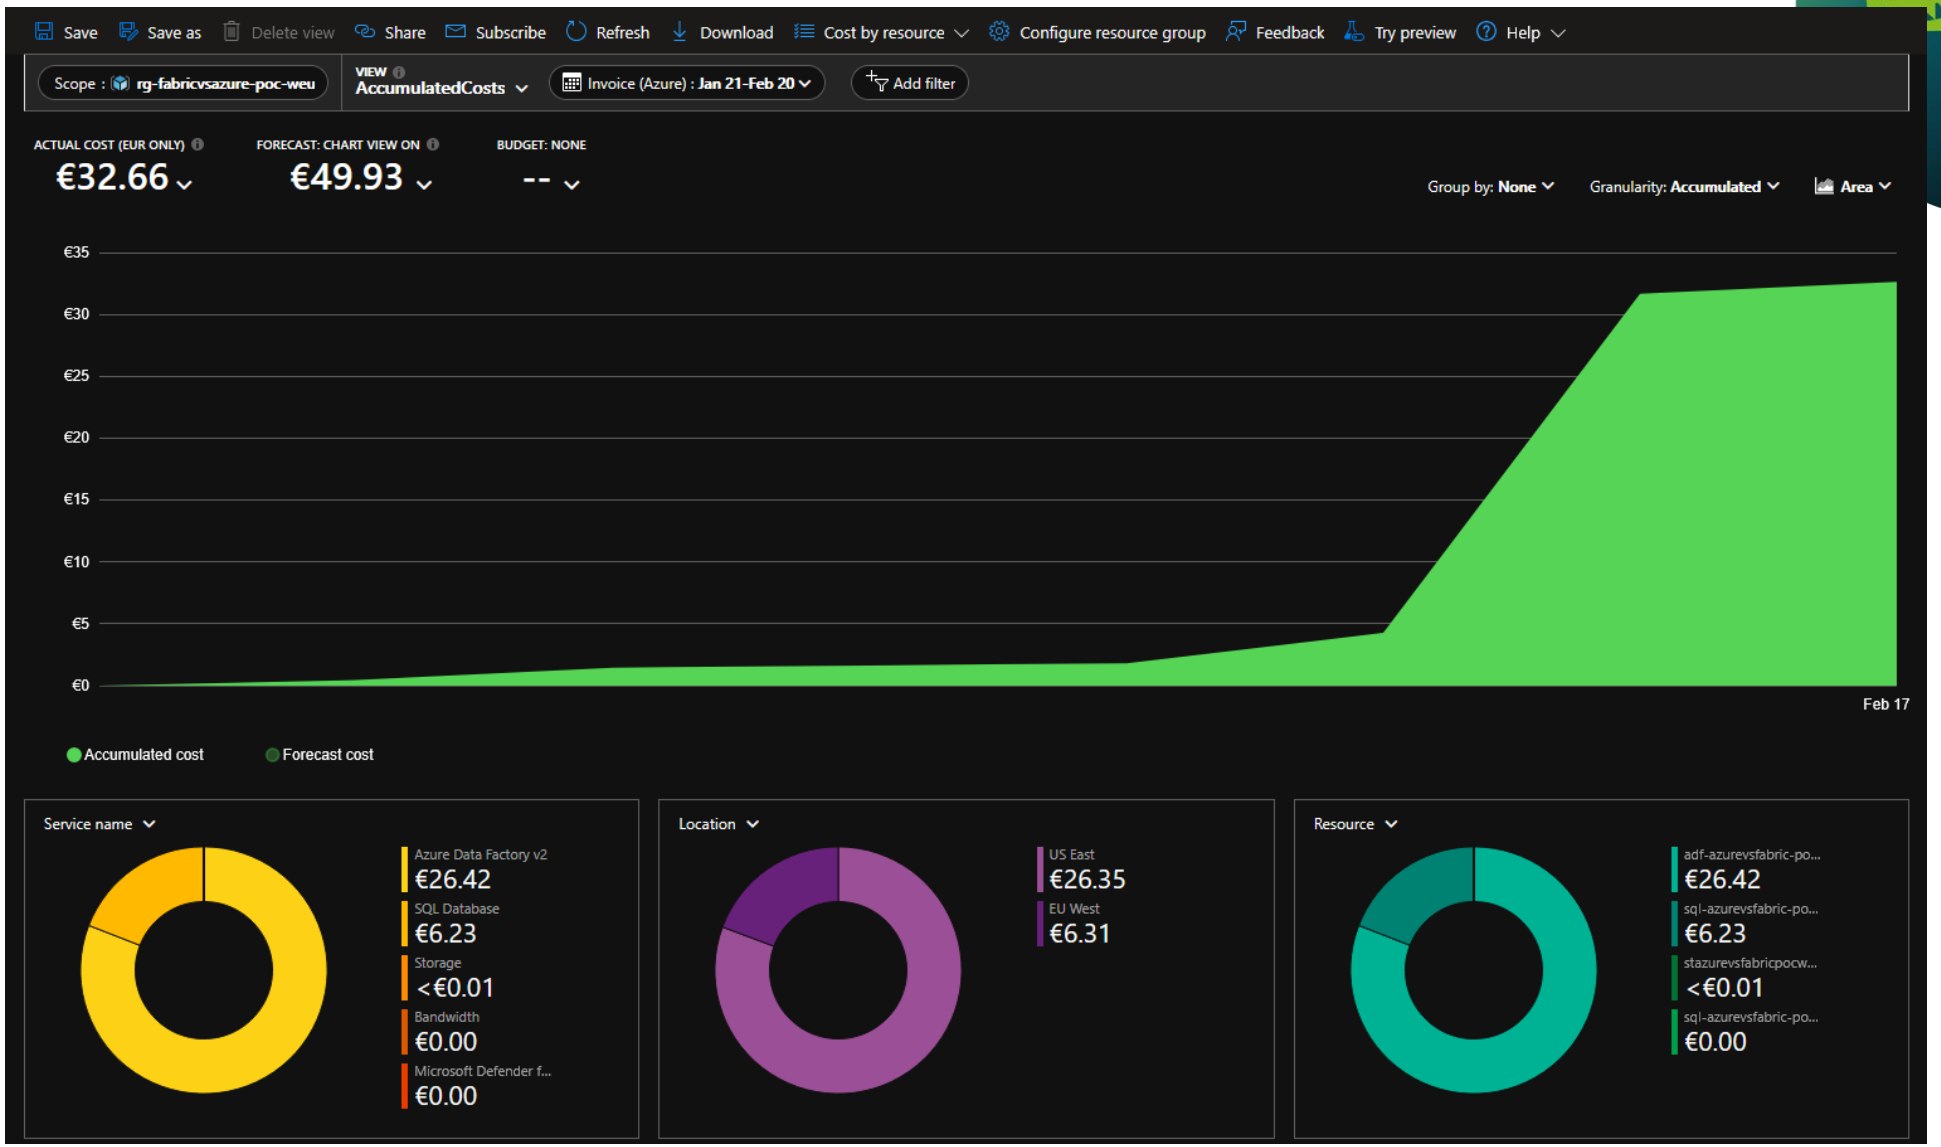


Activity run id: 6d12bd05-f8ef-4e0c-93d3-f9d57a709512

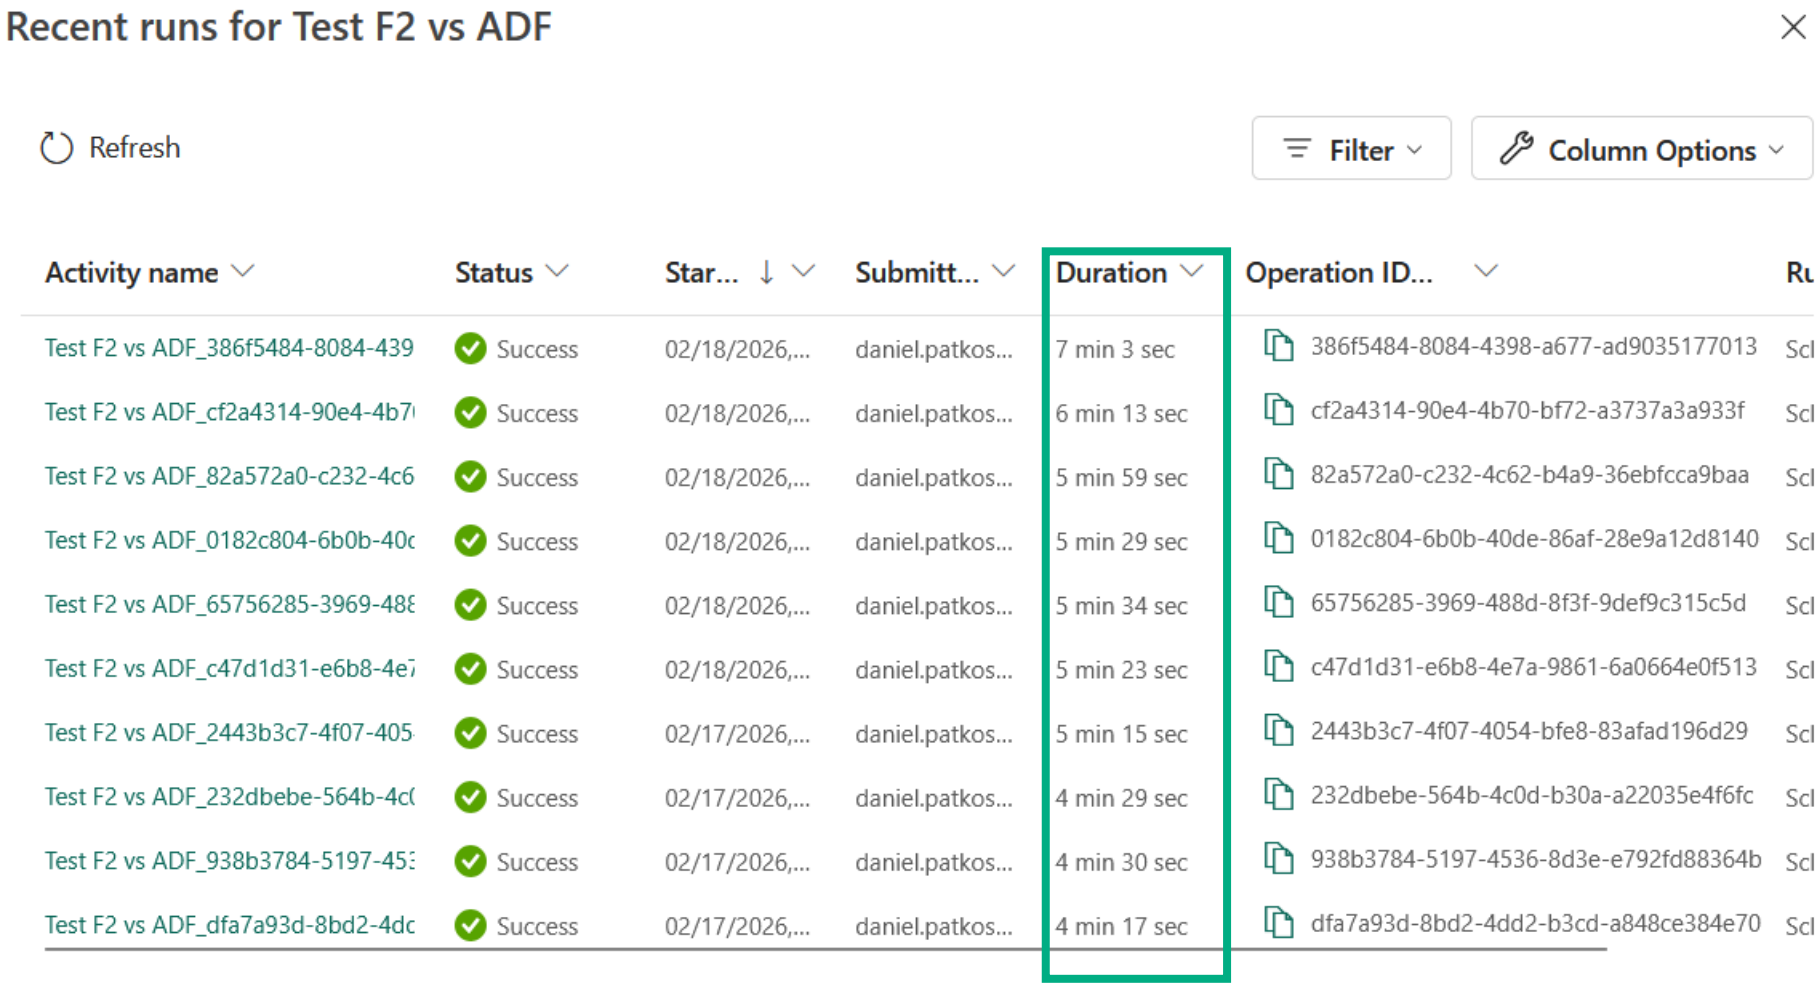


Run ID	Copy duration
1	04:00
7	27:17
10	53:13

ADF cost after 10th run



Fabric – Notebook job duration

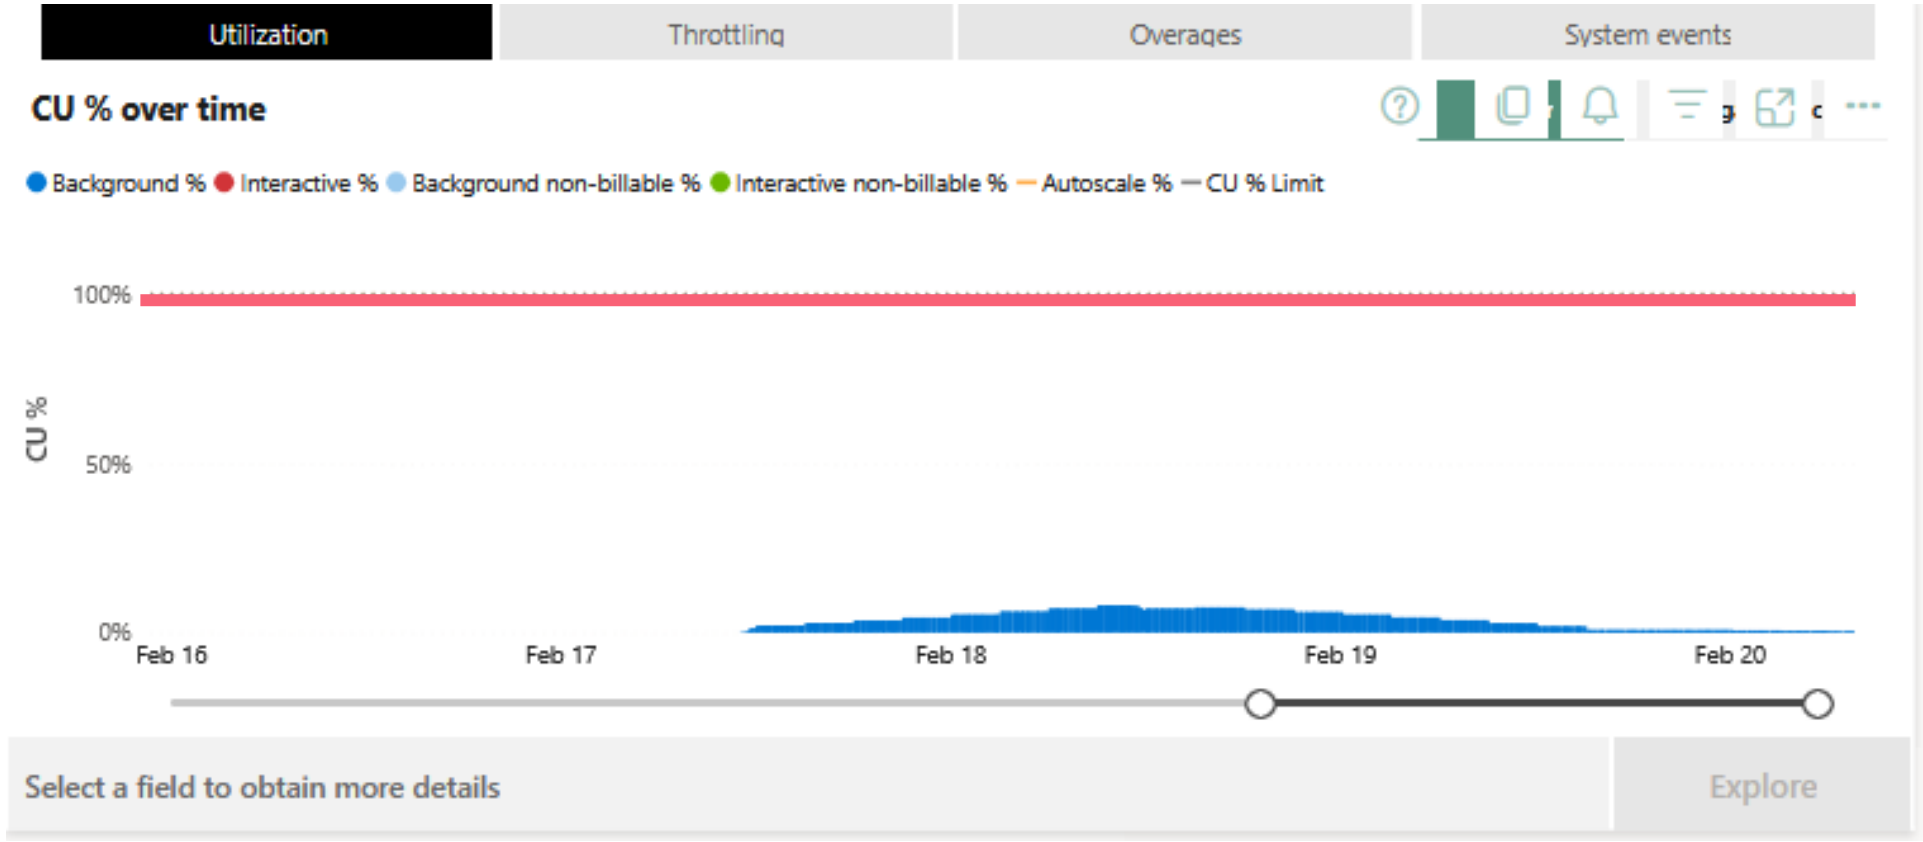


SQL
KONFERENZ
2026



Guess the CU consumption %!

CU Usage - Fabric



SQL KONFERENZ 2026



How far can we push F2?

How to kill a capacity?

SQL KONFERENZ 2026

Set up

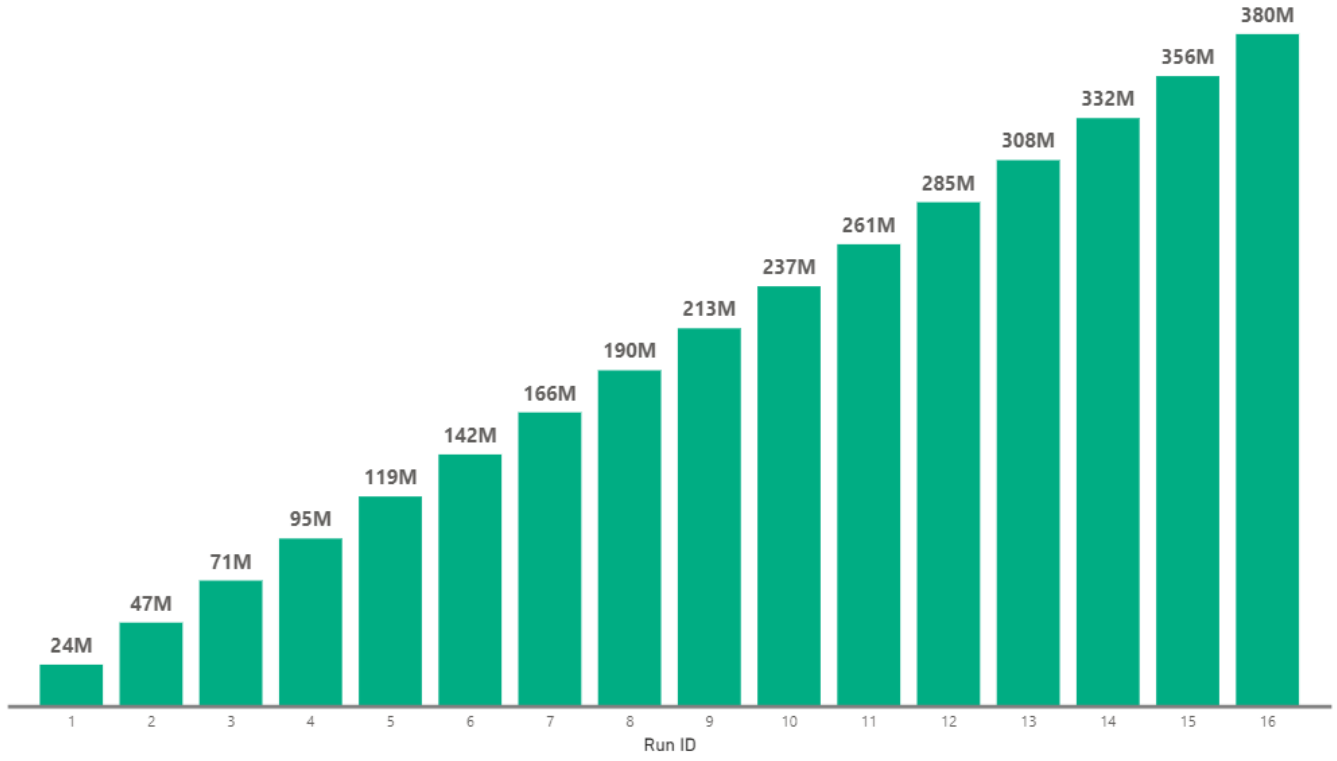


- Contoso 10M dataset
- Fact tables - duplicated 50x time
- Incremental load (only facts) – **every 2 hours**
- Approx **23M** rows / load to 🥉 layer
- 🥈 🥇 layers need to handle increasing volume
- Cleanse, transform, aggregate steps

SQL KONFERENZ 2026

Rows in Silver Layer



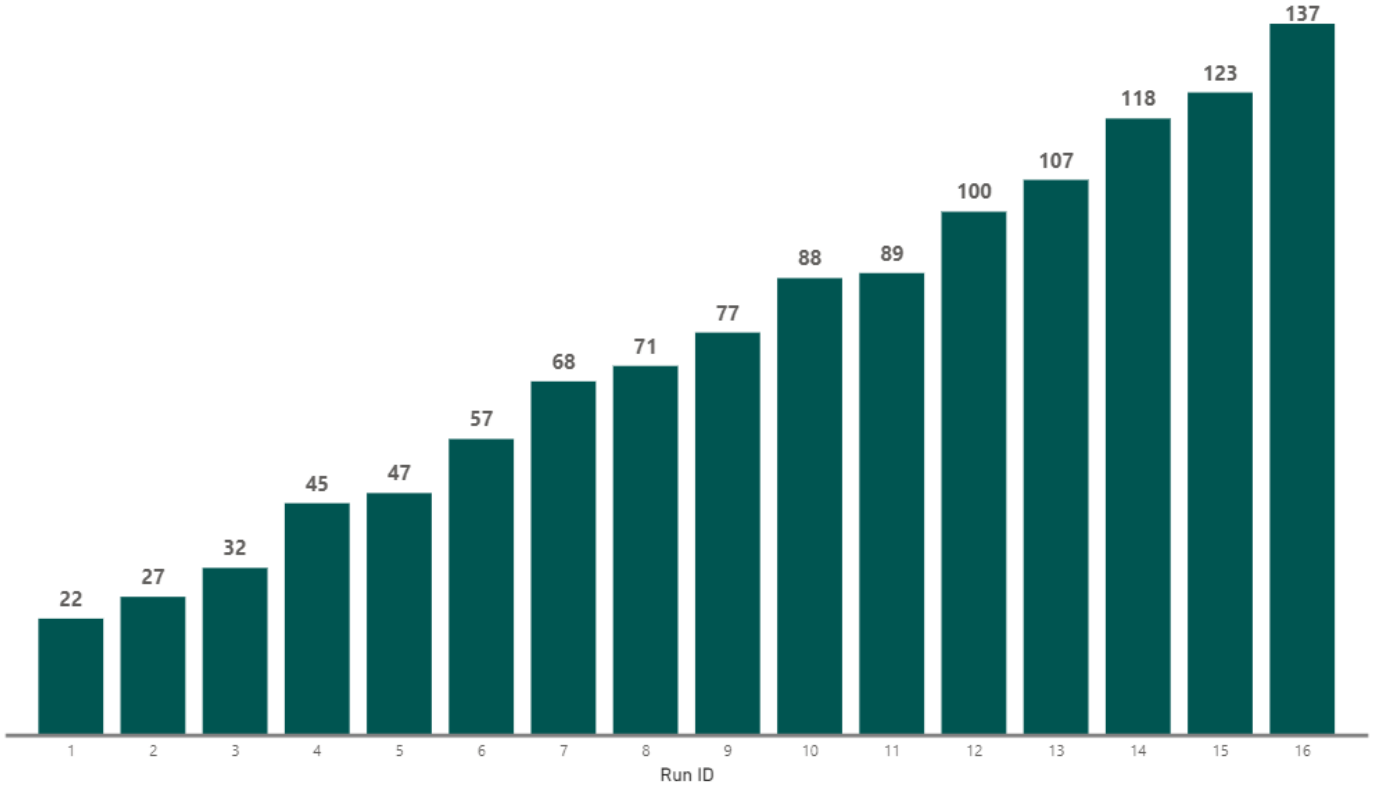


SQL KONFERENZ 2026

Runtime



Duration by batch loads in minutes



SQL
KONFERENZ
2026

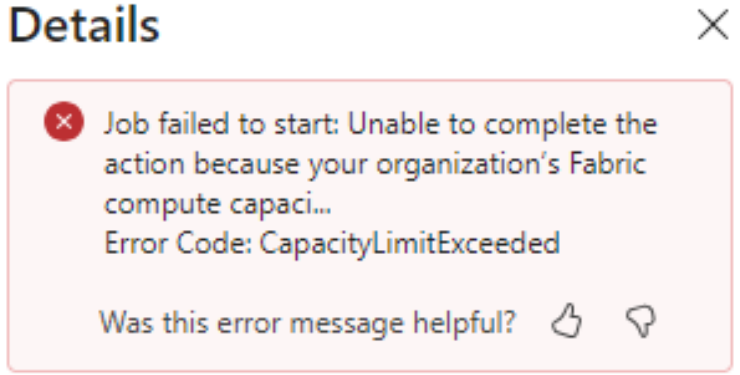


Which was the 1st load triggering rejection?

SQL
KONFERENZ
2026



Which was the 1st load to be fully rejected?



Monitor

View and track the status of the activities across all the workspaces for which you have permissions within Microsoft Fabric.

Refresh

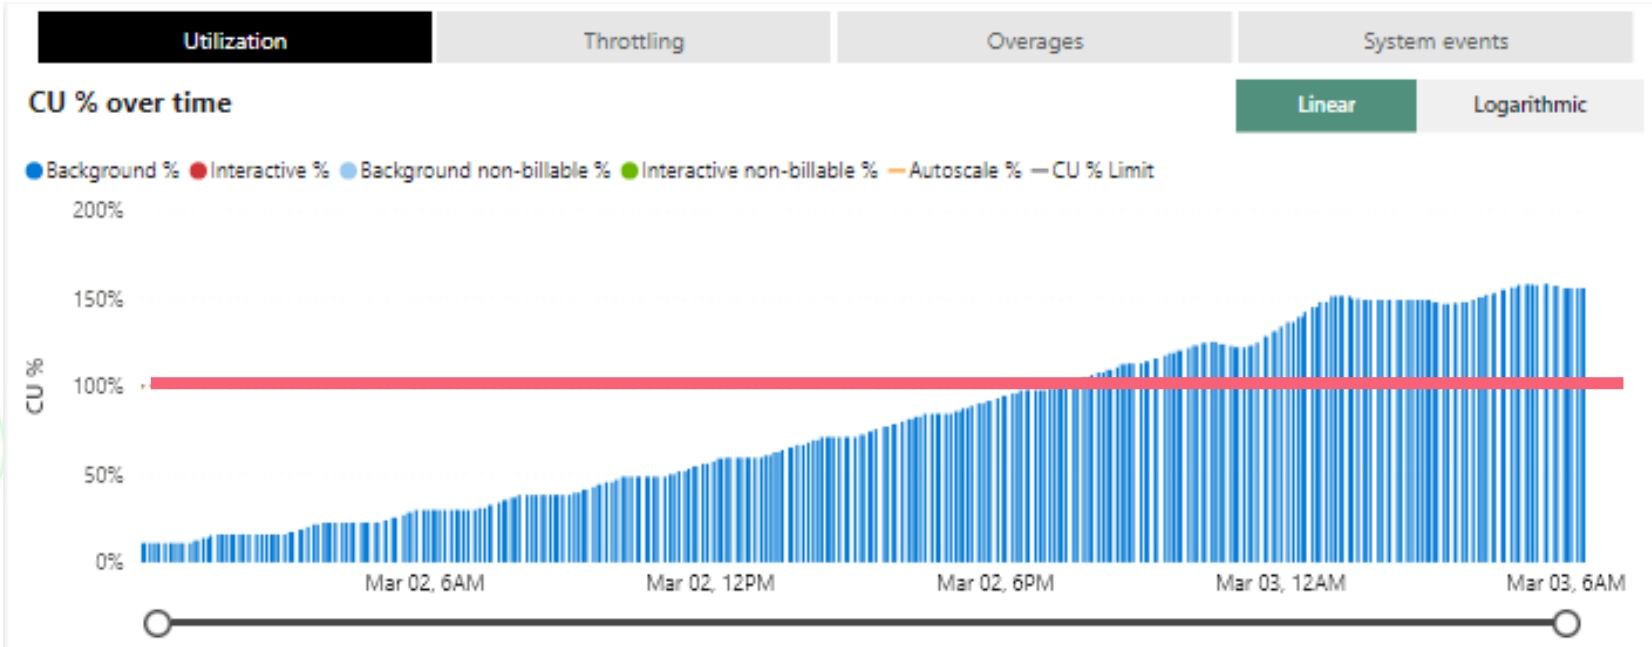
Clear all To apply filters, select the values from the Filter dropdown menu.

Activity name	Status	Item type	Start time	Submitted by	Location
contoso_notebook	Failed	Notebook	03/03/2026, 5:15 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_91c0e542-bf60-4436-b0de-e...	Succeeded	Notebook	03/03/2026, 3:15 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook	Failed	Notebook	03/03/2026, 1:15 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_140fc4de-ed3d-4f3f-bcc6-9...	Succeeded	Notebook	03/02/2026, 11:15 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_24b4a121-17b9-4058-982b-...	Succeeded	Notebook	03/02/2026, 11:00 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_93d2d004-7ada-4f3e-a3ee-d...	Succeeded	Notebook	03/02/2026, 9:01 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_685efac9-7cf9-4d62-947e-c...	Succeeded	Notebook	03/02/2026, 7:01 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_711c617c-d8a5-457d-991f-3...	Succeeded	Notebook	03/02/2026, 5:00 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_fc50cd0f-9dfc-4e78-b1ee-7f...	Succeeded	Notebook	03/02/2026, 3:00 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_27fc963d-d930-4265-bee0-9...	Succeeded	Notebook	03/02/2026, 1:00 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_0224557a-0f34-4bca-bbcd-3...	Succeeded	Notebook	03/02/2026, 11:00 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_a2ec8c76-2f52-4351-8ecd-1...	Succeeded	Notebook	03/02/2026, 9:00 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_ecc59816-30cd-4989-9c08-e...	Succeeded	Notebook	03/02/2026, 7:01 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_71eb5b4c-85c1-49b5-b5f2-a...	Succeeded	Notebook	03/02/2026, 5:01 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_51d6e7cd-1f56-4324-b951-1...	Succeeded	Notebook	03/02/2026, 3:00 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_071243a6-71c9-4d62-8687-...	Succeeded	Notebook	03/02/2026, 1:00 AM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_946a4278-77e2-423c-a005-e...	Succeeded	Notebook	03/01/2026, 11:00 PM	Dániel Patkós	F2 Demo Workspace
contoso_notebook_1297ee2a-1fe1-47e6-939b-4...	Succeeded	Notebook	03/01/2026, 10:31 PM	Dániel Patkós	F2 Demo Workspace

SQL KONFERENZ 2026

Result – WOW! You can do this?!

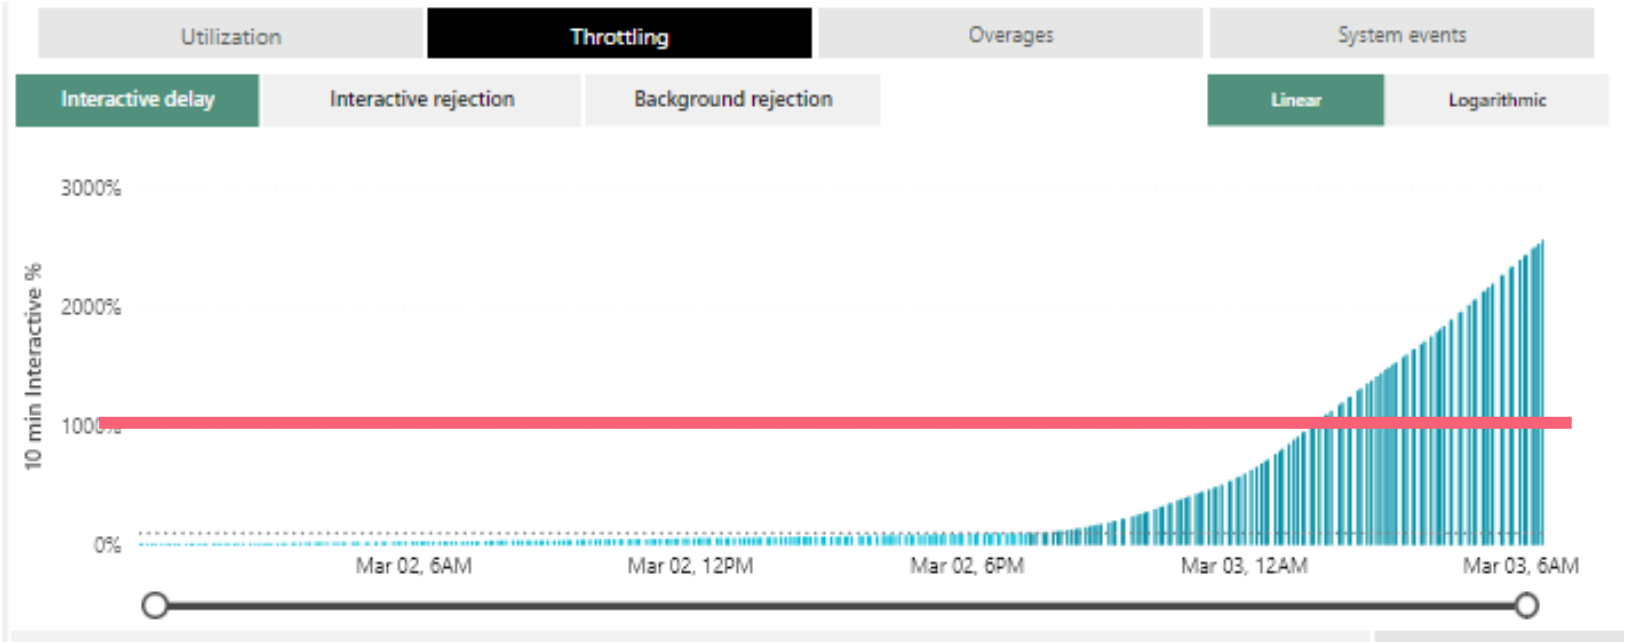




SQL KONFERENZ 2026

Can =/= Should





Home Replication Security Help

SQL database

Manage SQL security | Manage SQL Auditing

Explorer

Search

Capacity limit exceeded

Queries



Capacity limit exceeded

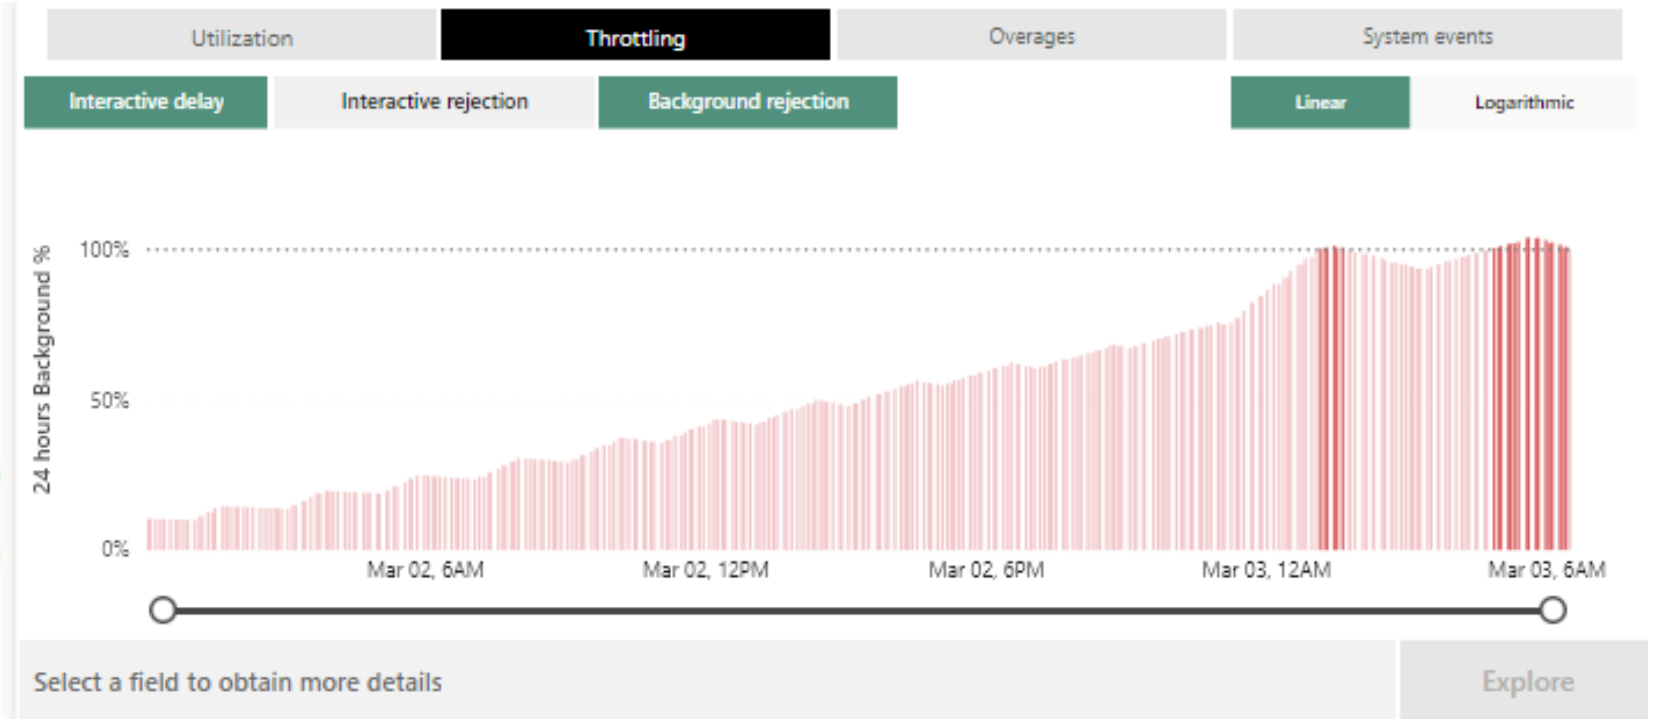
SQL database failed to create. Your organization's Fabric compute capacity has exceeded its limits. Try again later.

Refresh

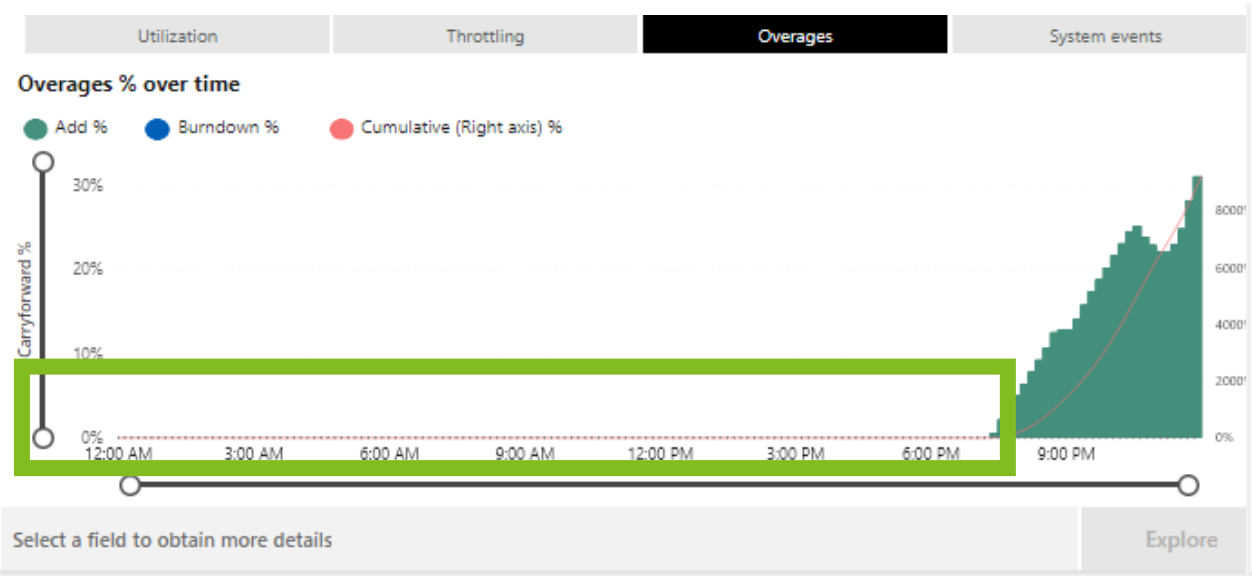
SQL KONFERENZ 2026

Can =/= Should

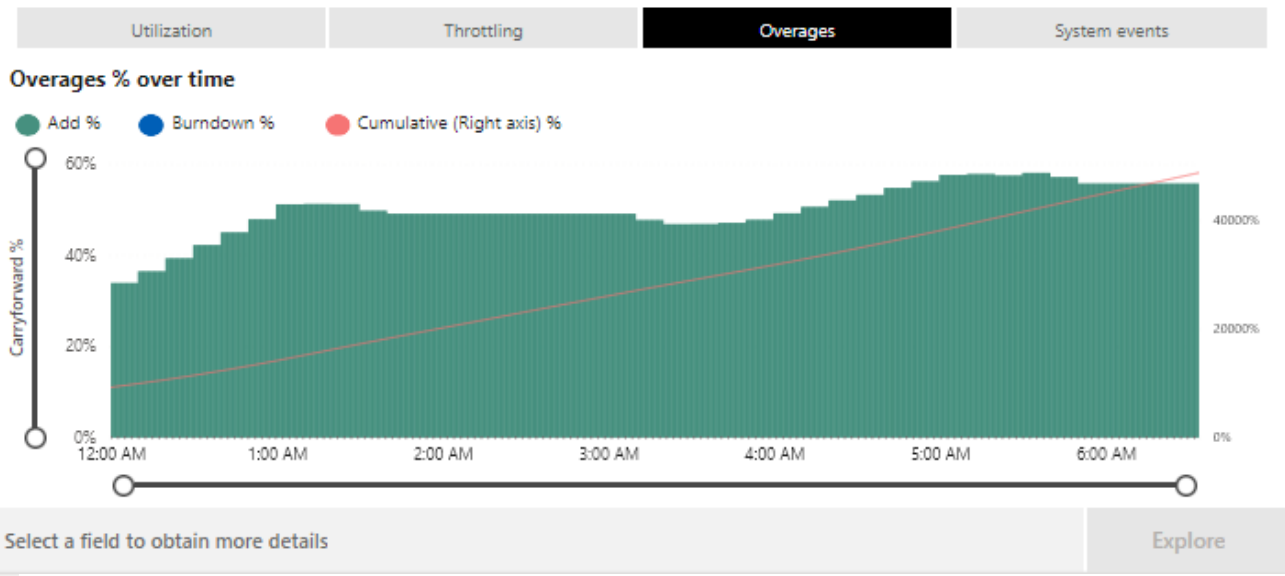




Can =/= Should

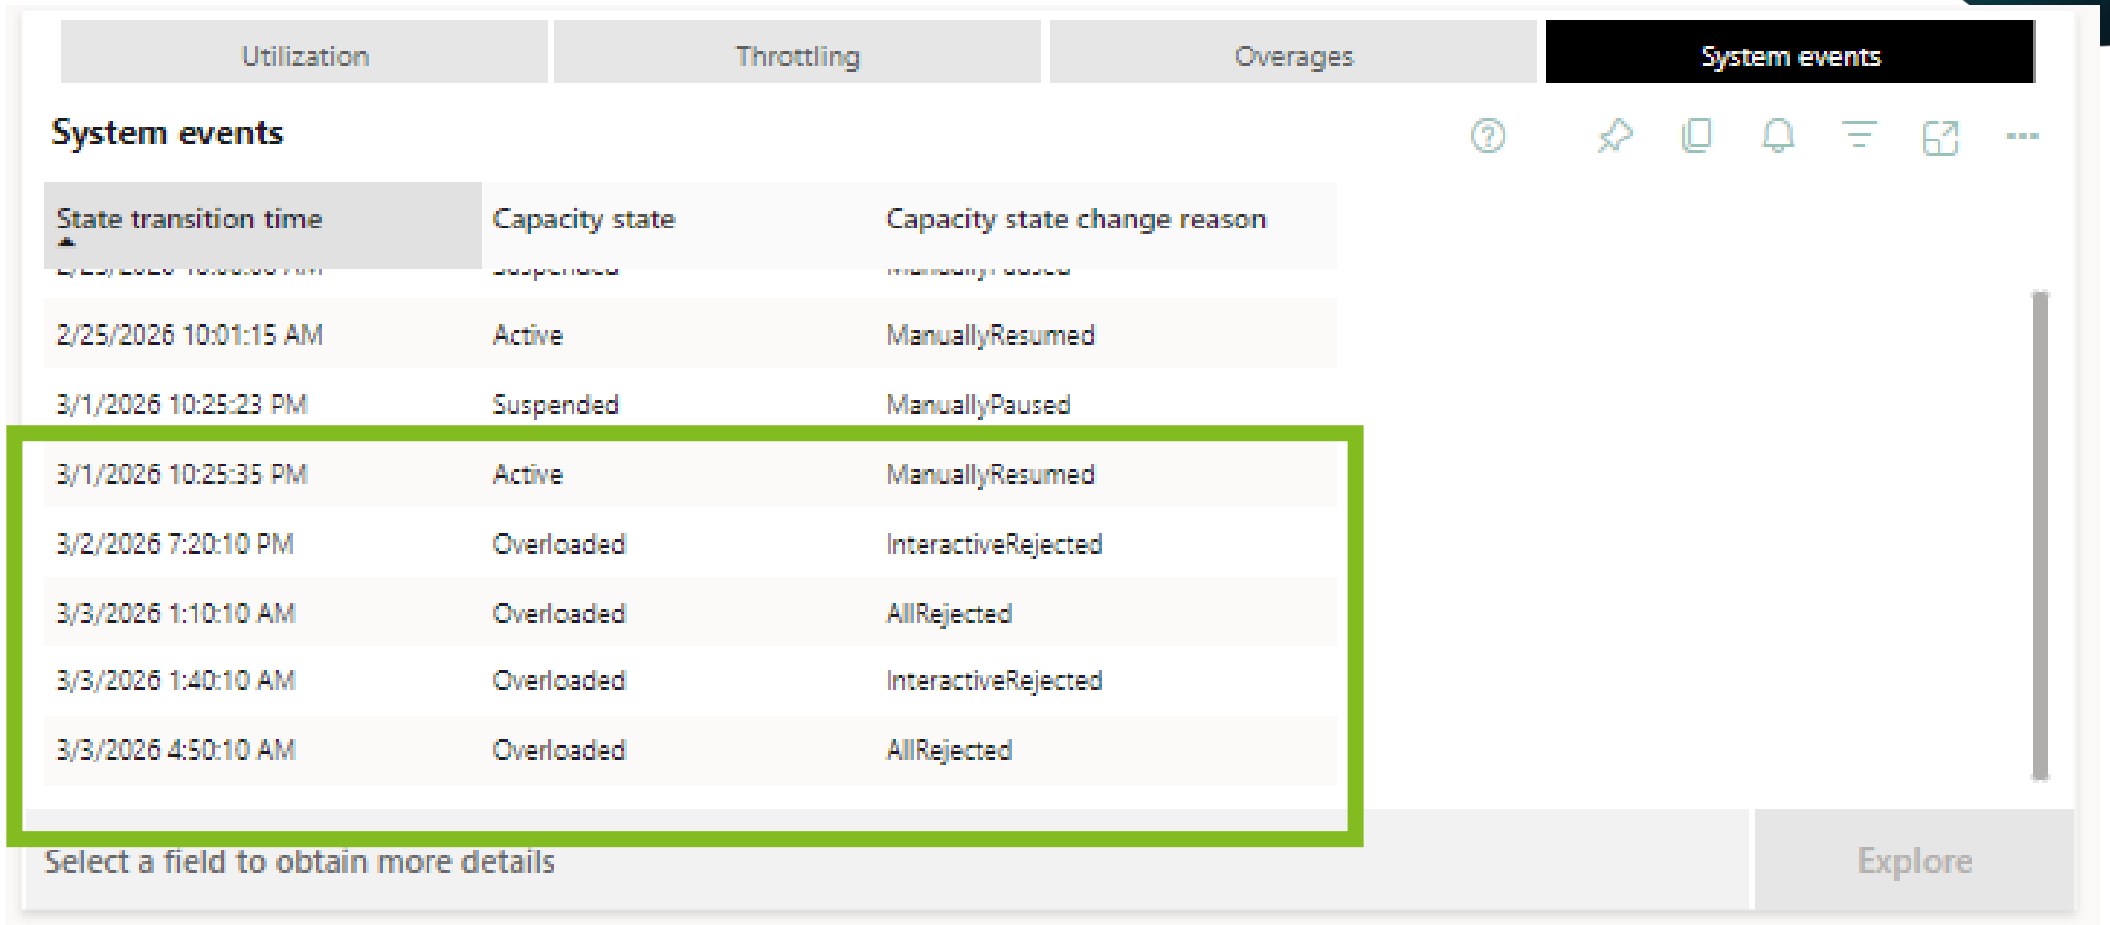


Day 1



Day 2

Can =/= Should


Utilization
Throttling
Overages
System events
System events
State transition time
Capacity state
Capacity state change reason
2/25/2026 10:01:15 AM
Active
ManuallyResumed
3/1/2026 10:25:23 PM
Suspended
ManuallyPaused
3/1/2026 10:25:35 PM
Active
ManuallyResumed
3/2/2026 7:20:10 PM
Overloaded
InteractiveRejected
3/3/2026 1:10:10 AM
Overloaded
AllRejected
3/3/2026 1:40:10 AM
Overloaded
InteractiveRejected
3/3/2026 4:50:10 AM
Overloaded
AllRejected
Select a field to obtain more details
Explore

SQL
KONFERENZ
2026



🛠️ Building on F2

Development reality and team pitfalls

SQL KONFERENZ 2026

DEV vs PROD capacities

Usual recommendation:



F2 DEV Capacity

<

F4 PROD Capacity

SQL KONFERENZ 2026

```python
# 🚀 Spark + Required Libraries
from pyspark.sql import SparkSession
from pyspark.sql.functions import col, current_timestamp, max as spark_max, lit, lower, when, to_date, to_timestamp, coalesce, concat_ws
from delta.tables import DeltaTable
import uuid, re, datetime
from pyspark.sql.types import TimestampType, DateType


#config for schema modification (adding column to existing table)
spark.conf.set("spark.databricks.delta.schema.autoMerge.enabled", "true")
```

PySpark (Python)

```
InvalidHttpRequestToLivy: [TooManyRequestsForCapacity] This spark job can't be run because you have hit a spark compute or API rate limit. To run this spark job, cancel an active Spark job through the Monitoring hub, choose a larger capacity SKU, or try again later. HTTP status code: 430 {Learn more} HTTP status code: 430.
```



Collaboration session joined

Session error

[TooManyRequestsForCapacity] This spark job can't be run because you have hit a spark compute or API rate limit. To run this spark job, cancel an active Spark job through the Monitoring hub, choose a larger capacity SKU, or try again later. HTTP status code: 430 {Learn more} HTTP status code: 430.

Diagnostics

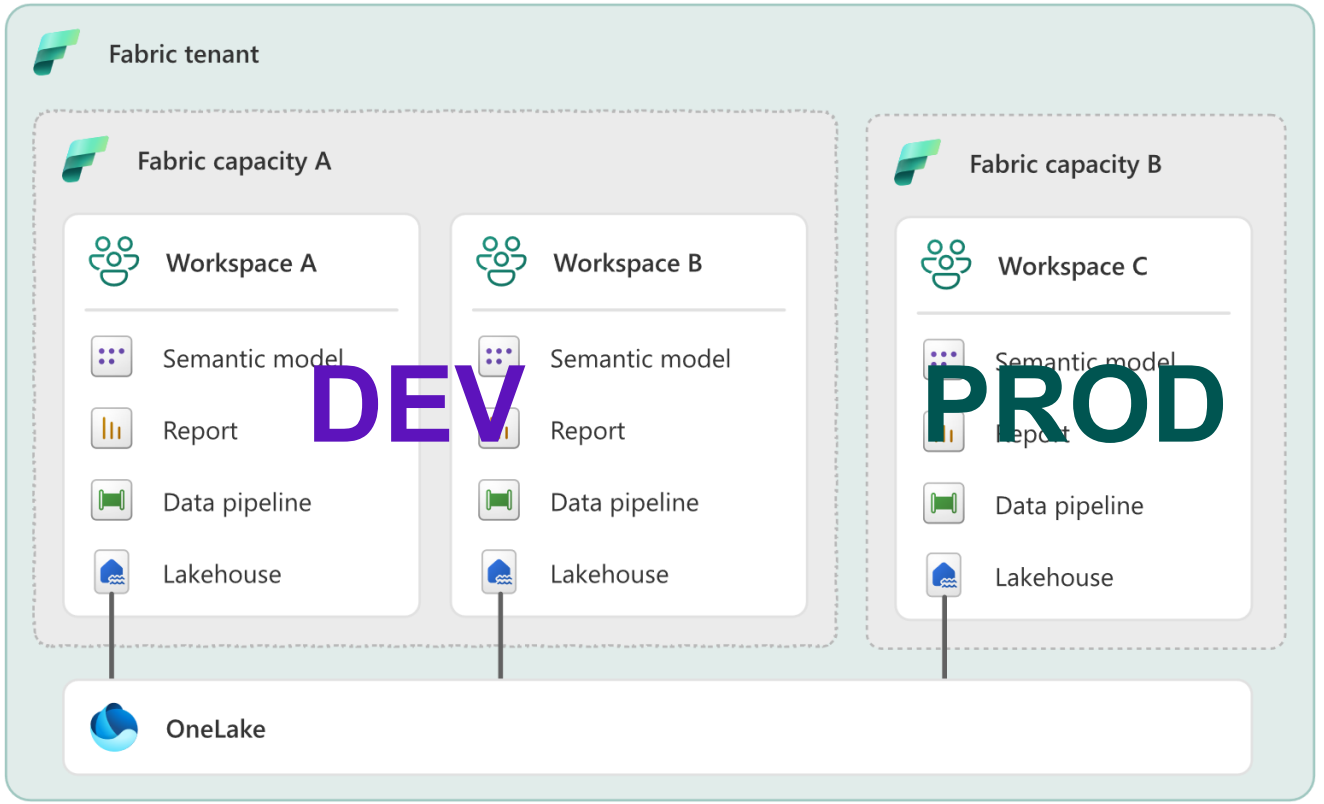
Copy and provide the following information if you are seeking technical support.

Expand Copy to clipboard

SQL KONFERENZ 2026

DEV vs PROD „paradox"

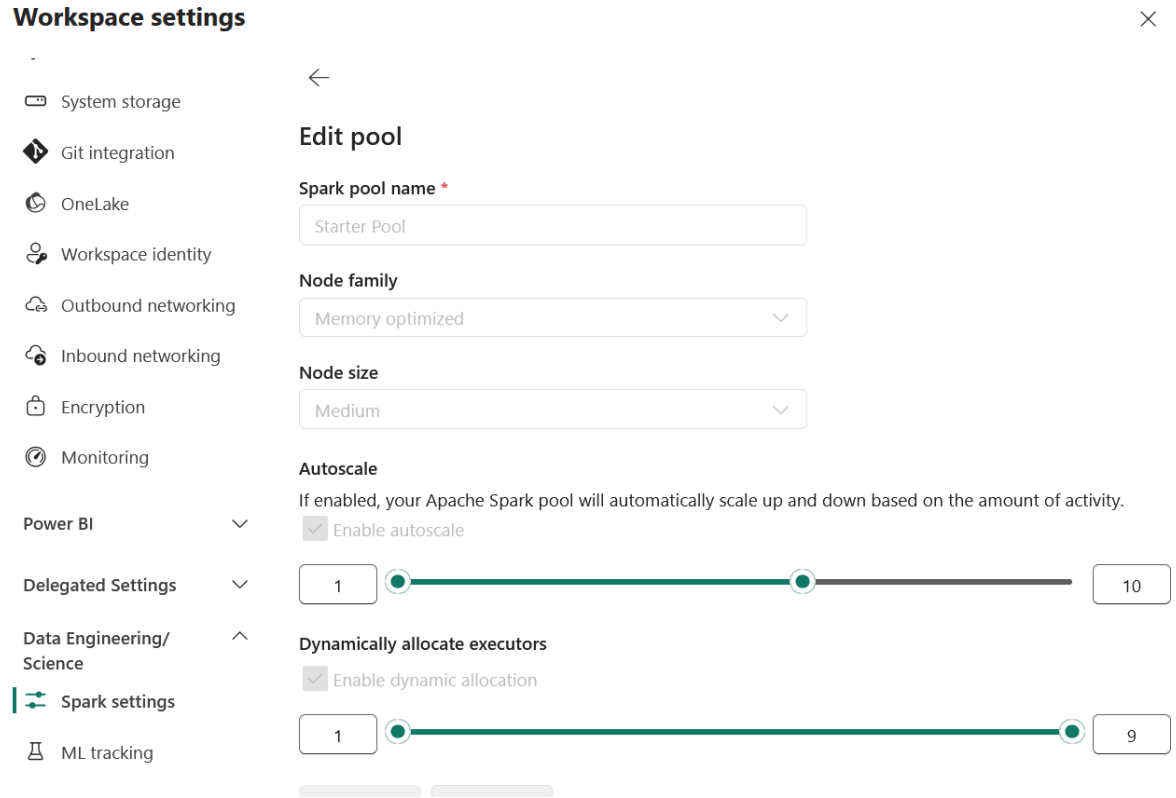




SQL KONFERENZ 2026

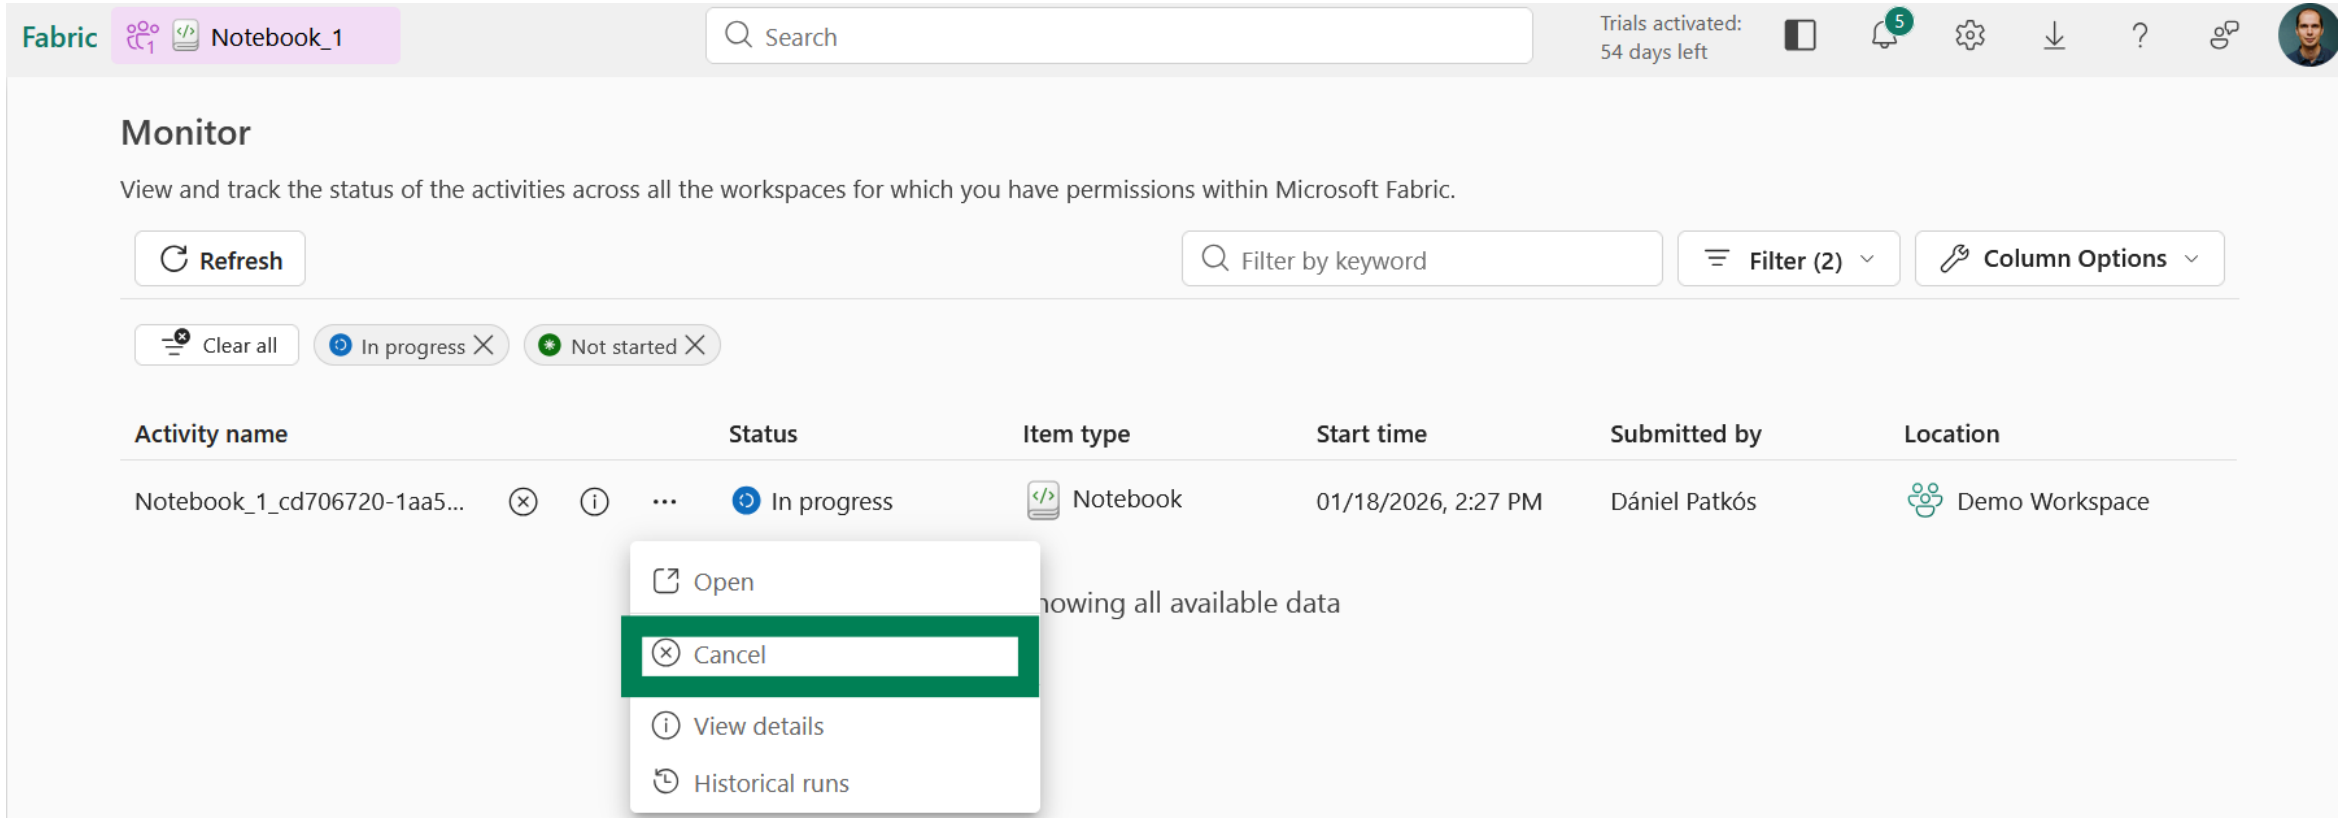
Default settings are not F2 friendly





```
Py4JJavaError: An error occurred while calling z:notebookutils.udf.run.
: java.lang.Exception: Request to https://10cb77630ad947ba9f864cc9160fda33.pbidedicated.windows.net/webapi/
capacities/10CB7763-0AD9-47BA-9F86-4CC9160FDA33/workloads/FuncSet/FuncSetService/automatic/functionsets/cd1e60e0-a549-4e43-a1d1-d00859c36510/
functions/get_envs/invoke?workspaceId=572dc55f-df4a-4e27-82cb-5d5423305390 failed with status code: 429, response:
{"code":"CapacityLimitExceeded","subCode":0,"message":"Unable to complete the action because your organization's Fabric compute capacity has
exceeded its limits. Try again later.","timeStamp":"2026-02-12T15:43:07.4609302Z","httpStatusCode":429,"hresult":-2147467259,"details":
[{"code":"RootActivityId","message":"e0d99dce-fa0f-4c8d-b9d2-acce0b1b05b8"}],"exceptionCategory":1}, response headers: Array(Content-Length: 377,
Content-Type: application/json, Strict-Transport-Security: max-age=31536000; includeSubDomains, x-ms-routing-hint: host004_funcset-016, x-ms-
workload-resource-moniker: bc293ea2-ac73-4eb8-95cd-054cbcfff6a8, x-ms-root-activity-id: e0d99dce-fa0f-4c8d-b9d2-acce0b1b05b8, x-ms-current-utc-
date: 2/12/2026 3:43:07 PM, Date: Thu, 12 Feb 2026 15:43:07 GMT)
    at com.microsoft.spark.notebook.workflow.client.BaseRestClient.request(BaseRestClient.scala:95)
    at com.microsoft.spark.notebook.workflow.client.FabricMwcClient.request(FabricMwcClient.scala:29)
    at com.microsoft.spark.notebook.msutils.impl.fabric.UdfUtilsImpl$.run(UdfUtilsImpl.scala:216)
    at notebookutils.udf$.$anonfun$run$1(udf.scala:47)
    at com.microsoft.spark.notebook.common.trident.CertifiedTelemetryUtils$.withTelemetry(CertifiedTelemetryUtils.scala:82)
    at notebookutils.udf$.run(udf.scala:46)
    at notebookutils.udf.run(udf.scala)
    at java.base/jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
    at java.base/jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
    at java.base/jdk.internal.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
    at java.base/java.lang.reflect.Method.invoke(Method.java:566)
    at py4j.reflection.MethodInvoker.invoke(MethodInvoker.java:244)
    at py4j.reflection.ReflectionEngine.invoke(ReflectionEngine.java:374)
    at py4j.Gateway.invoke(Gateway.java:282)
    at py4j.commands.AbstractCommand.invokeMethod(AbstractCommand.java:132)
    at py4j.commands.CallCommand.execute(CallCommand.java:79)
    at py4j.GatewayConnection.run(GatewayConnection.java:238)
    at java.base/java.lang.Thread.run(Thread.java:829)
```

Details to pay attention to 2/2



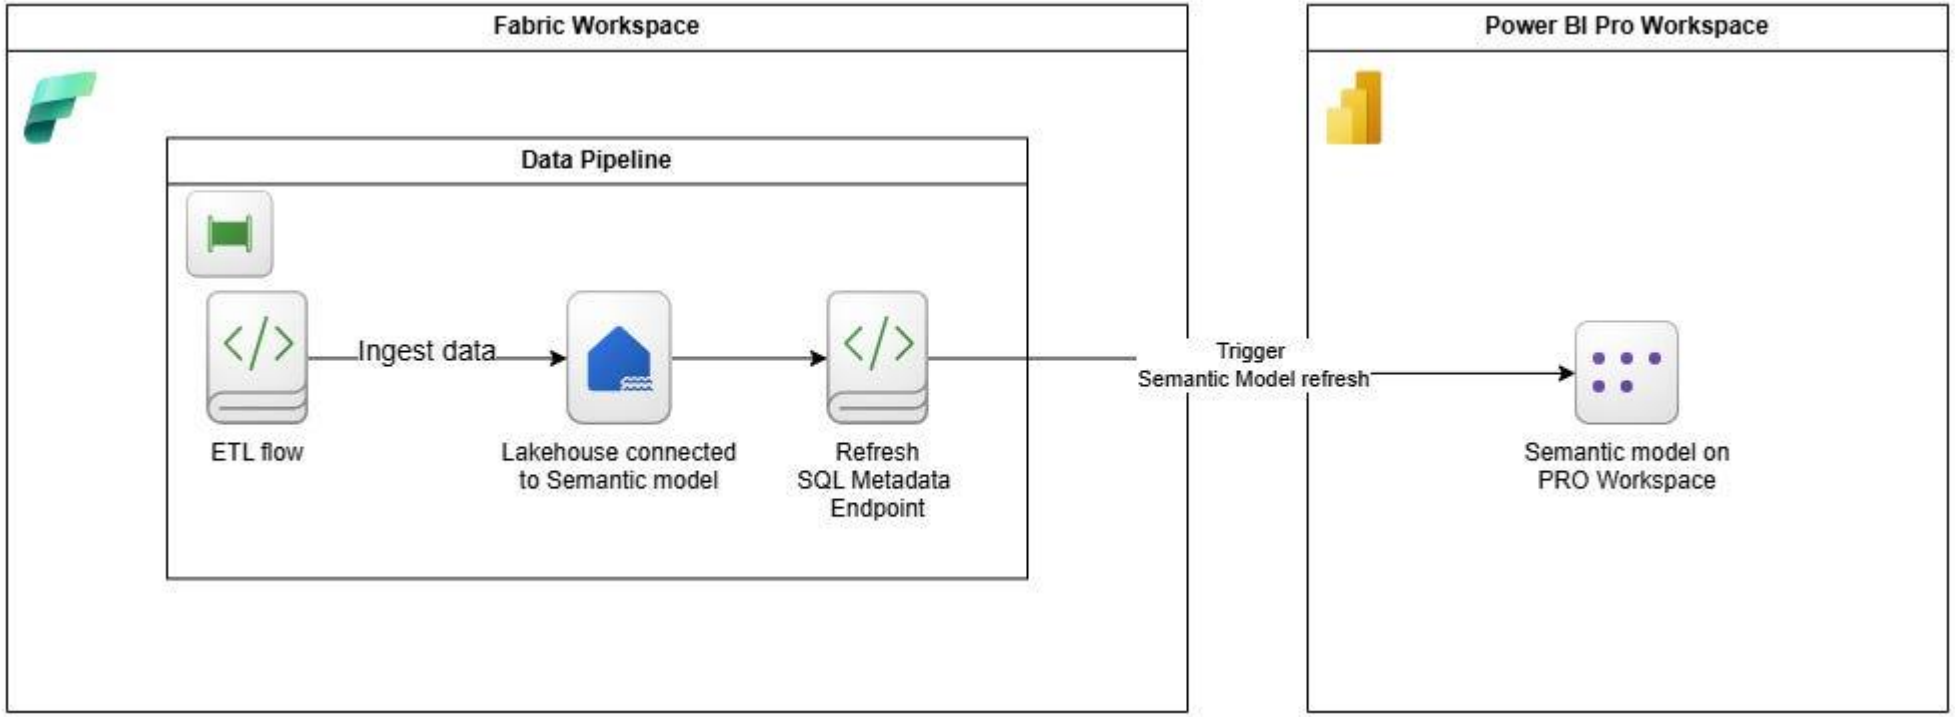
SQL KONFERENZ 2026

Why is this report sooo slow? Yesterday was everything faster!



VS

Split Fabric and Power BI items



SQL
KONFERENZ
2026



🚧 Where F2 hits its limits

SQL KONFERENZ 2026

F2 vs F4+ (vs Power BI Pro)



SKU	CU / sec	Max memory (GB)	Live connection / sec	Max memory per query (GB)	Model refresh parallelism	Direct Lake rows per table (million)
F2	2	**3**	2	1	1	300
F4	4	**3**	2	1	2	300
F8	8	**3**	3.75	1	5	300
F16	64	25	30	10	40	1500
...	...	...	...	...	...	...
F2048	2048	400	960	40	1,280	24,000

SQL
KONFERENZ
2026

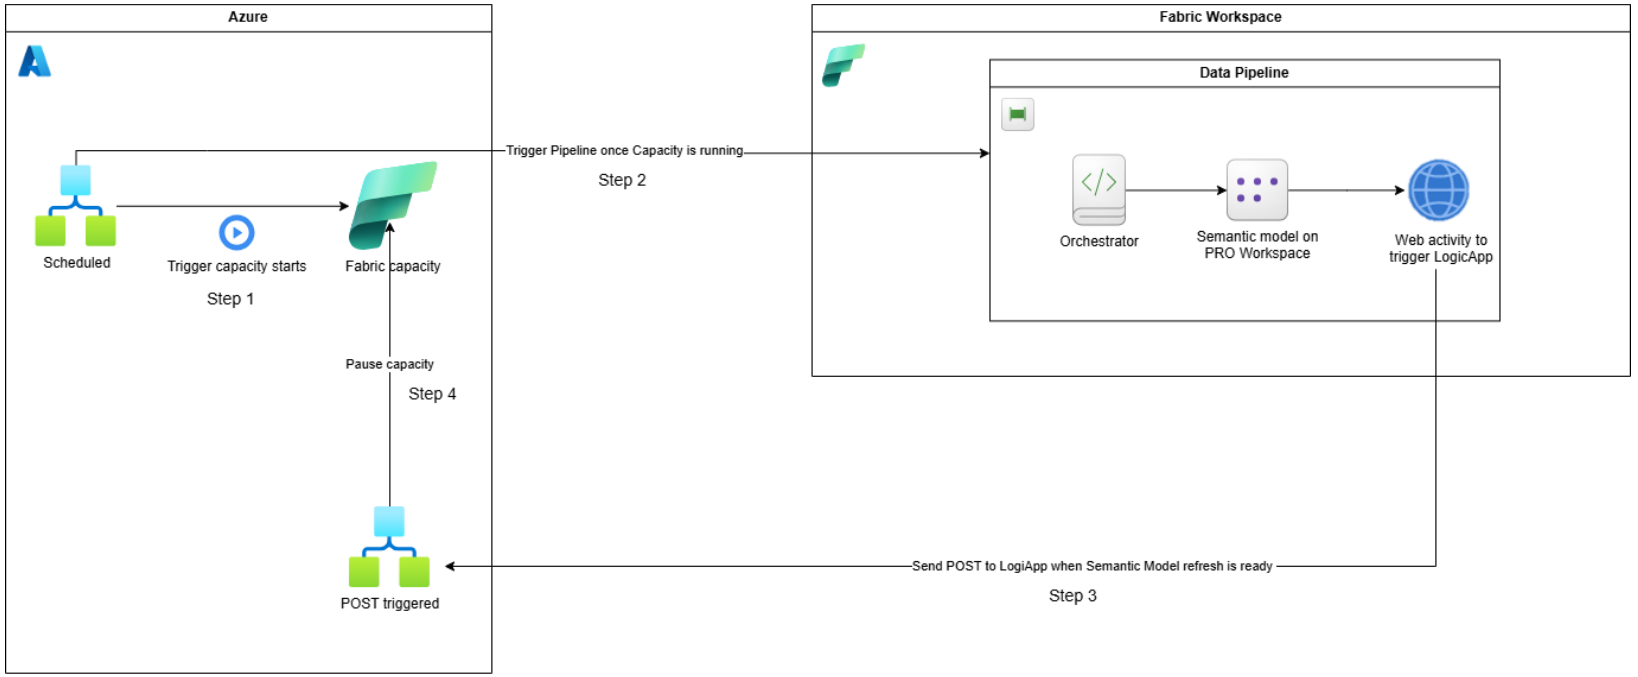


🧱 Architecture recommendations for F2

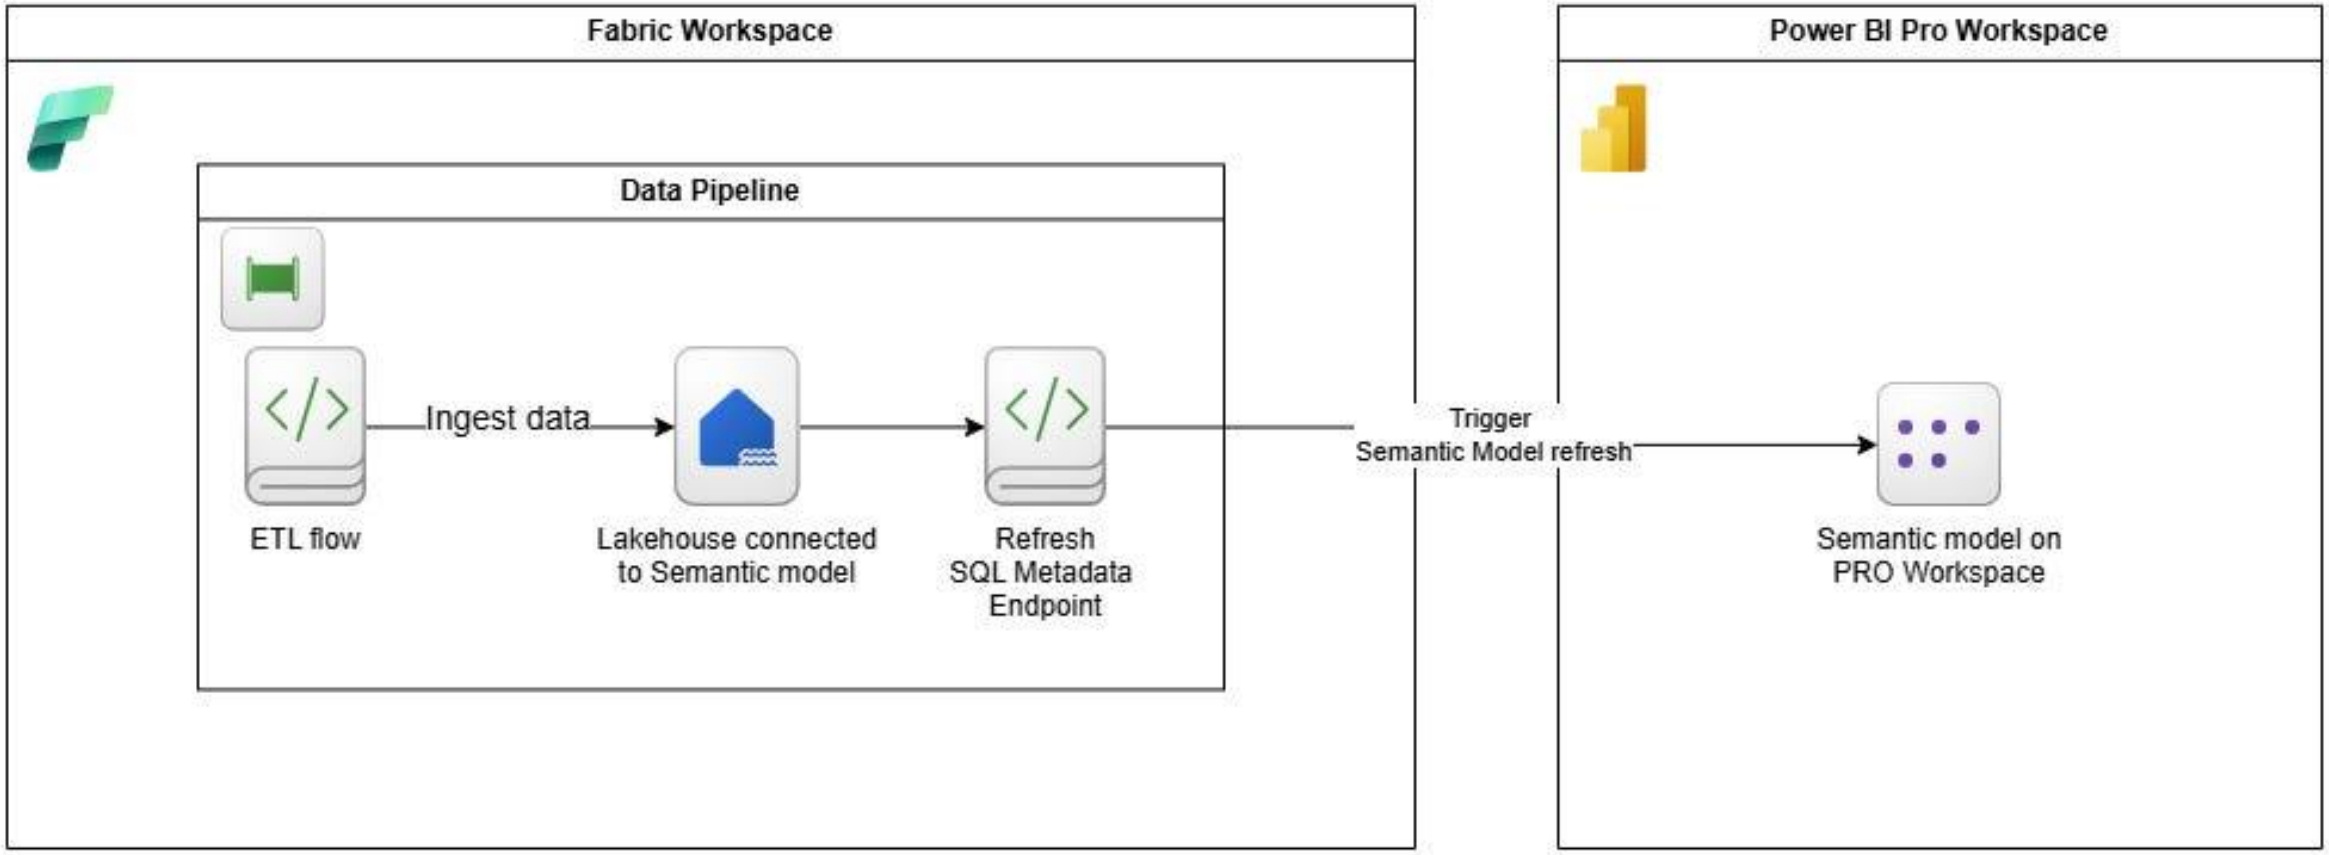
SQL KONFERENZ 2026

Pay-as-You-Go model

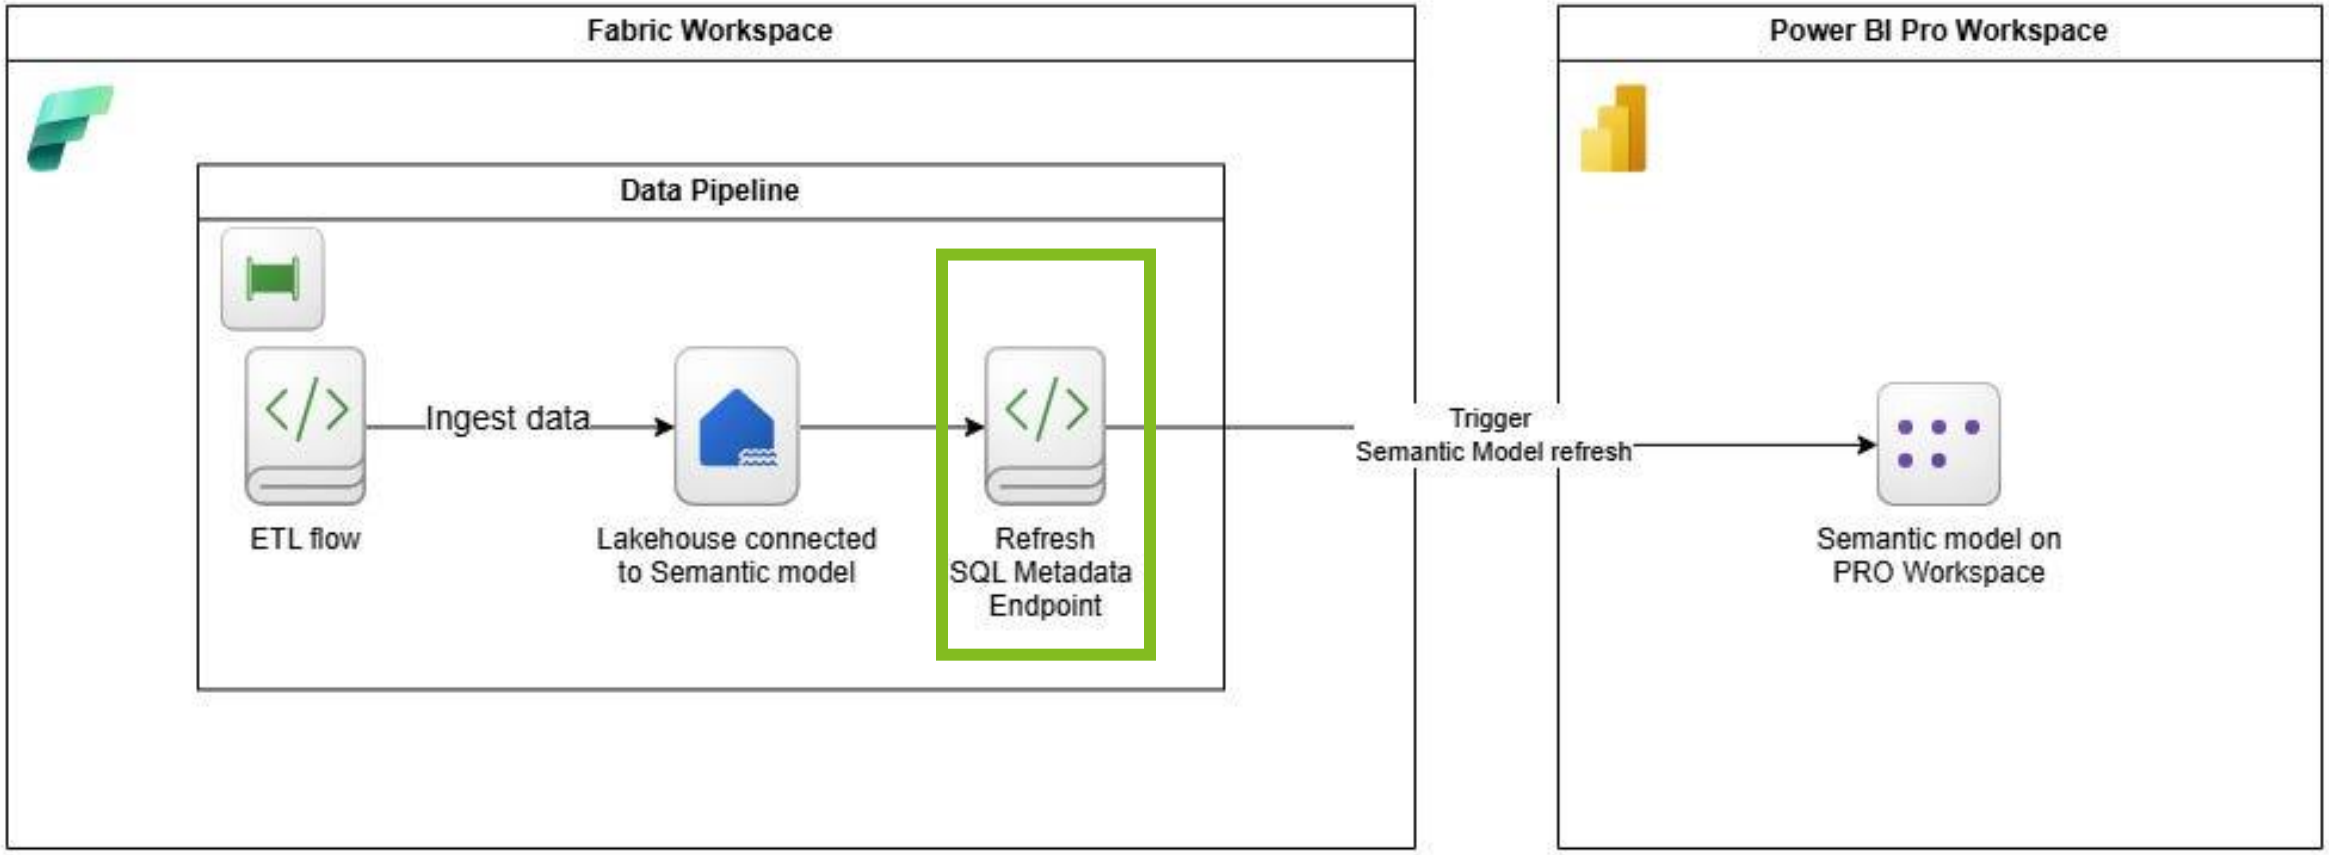




Split Power BI and Fabric items



Don’t forget to refresh the SQL metadata endpoint



SQL KONFERENZ 2026

📜 Summary



- **F2 is disciplined Fabric**
- Keep Power BI consumption separate from Fabric capacity
- F2 can be a valid PROD option for small and mid-sized data volumes
- F2 works best with predictable workloads and small, coordinated dev teams
- Advanced networking is still missing (not F2 specific)
- Scaling capacity helps, but it won’t fix design issues

SQL
KONFERENZ
2026

Help me improve this session!



SQL KONFERENZ 2026

Useful resources

- Explaining throttling Part1, Part2, Part3
- Understand bursting, smoothing
- Fabric Billing series



Time for:

thoughts, questions, or reactions

SQL KONFERENZ 2026

