Microsoft

Magpie

Python at Speed and Scale using Cloud Backends

Alekh Jindal, K. Venkatesh Emani, Maureen Daum, Olga Poppe, Brandon Haynes, Anna Pavlenko, Ayushi Gupta, Karthik Ramachandra, Carlo Curino, Andreas Mueller, Wentao Wu, Hiren Patel

Gray Systems Lab, Microsoft
Azure Data, Microsoft
Microsoft Research
University of Washington

The Python and The Cloud

+

- De-facto for ad-hoc analysis
- Pandas dataframes highly popular
- Performance is a challenge

- Hyper-scale performance
- Several SQL processing backends
- Enterprise data already on cloud

The current landscape … is a fragmented jungle!

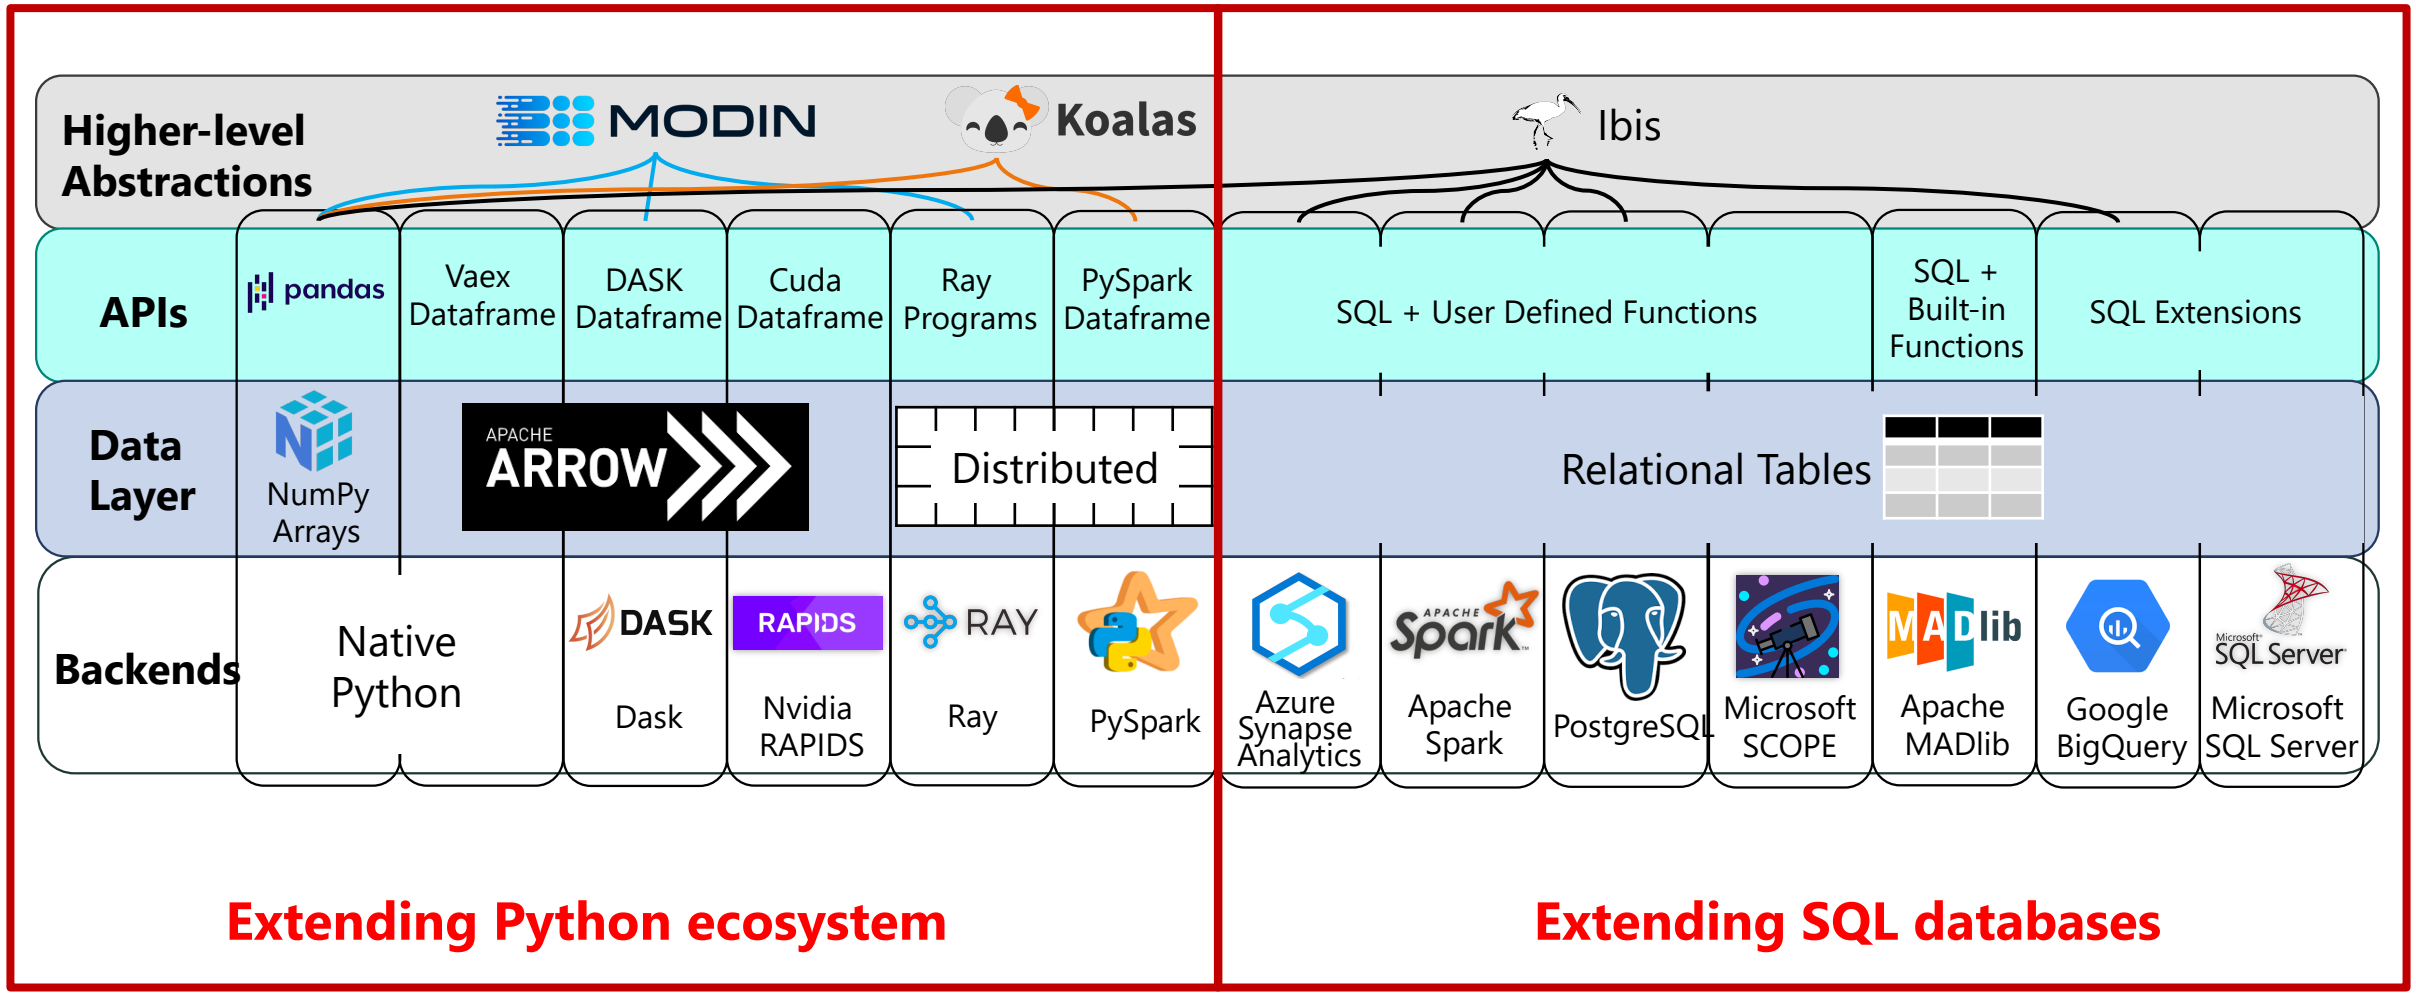

Magpie

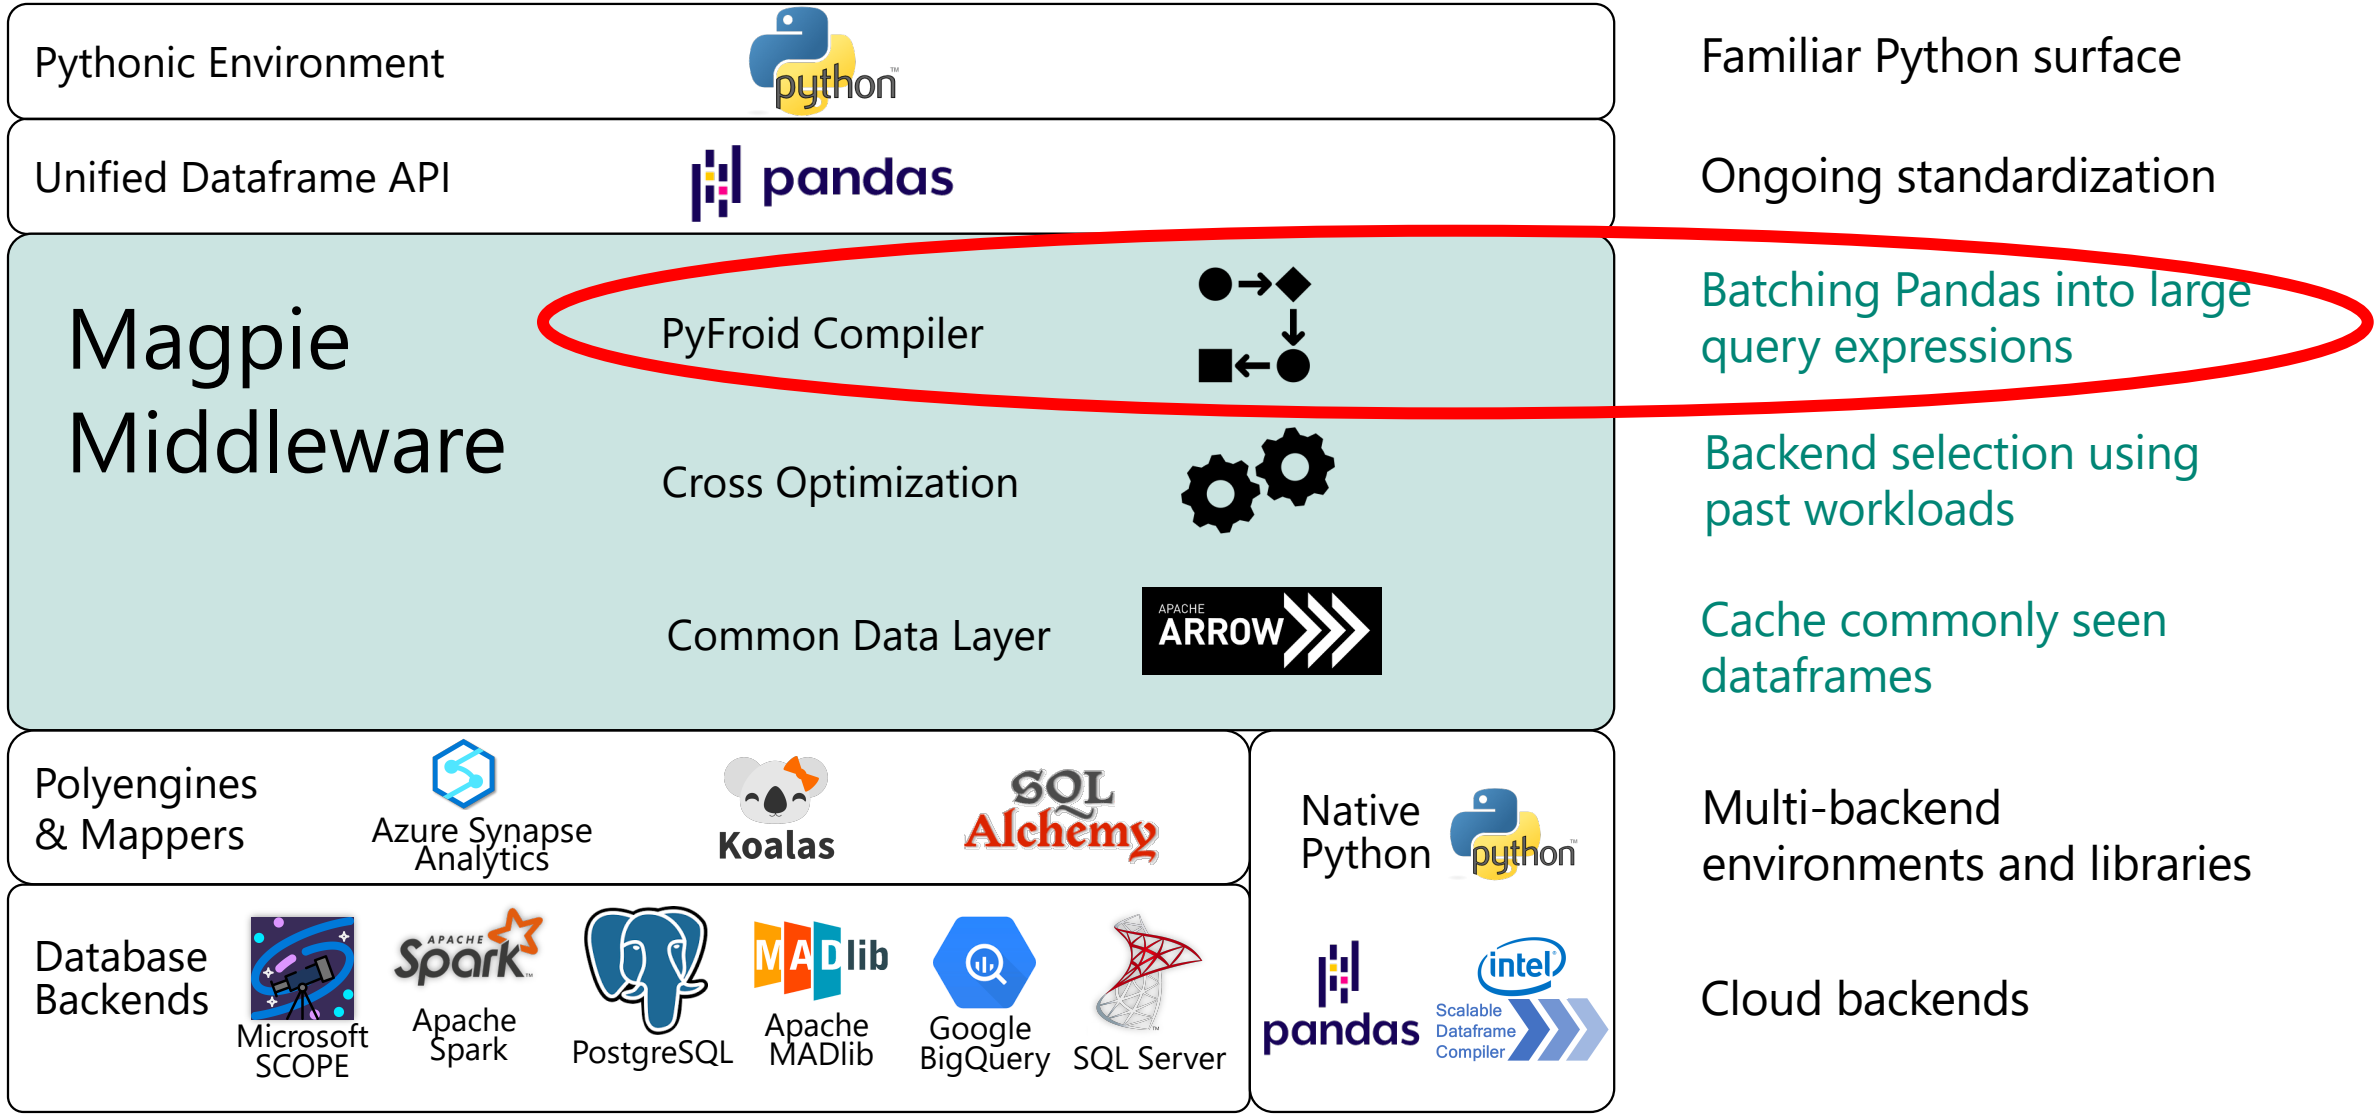
Pythonic Environment
python
Unified Dataframe API
pandas
Magpie Middleware
PyFroid Compiler
Cross Optimization
Common Data Layer
APACHE ARROW
Polyengines & Mappers
Azure Synapse Analytics
Koalas
SQL Alchemy
Database Backends
Microsoft SCOPE
Apache Spark
PostgreSQL
Apache MADlib
Google BigQuery
SQL Server
Native Python
python
pandas
intel
Scalable Dataframe Compiler
Familiar Python surface
Ongoing standardization
Batching Pandas into large query expressions
Backend selection using past workloads
Cache commonly seen dataframes
Multi-backend environments and libraries
Cloud backends

Batching Pandas

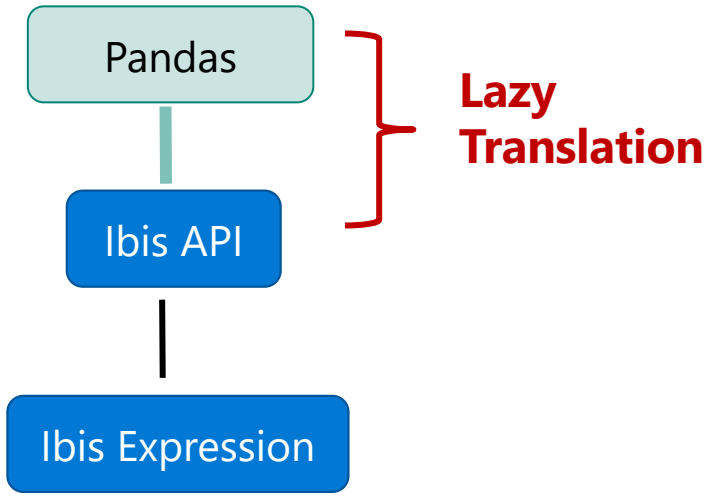

The number of taxi trips per weekday over the NYC Taxi dataset

```
1 import pyfroid.pandas as pd                          # vs import pandas as pd
2 df = pd.read_sql('nyctaxi', con)                     # fetch data
3 df = df[df.fare_amount > 0]                          # filter bad rows
4 df['day'] = df.pickup_datetime.dt.dayofweek          # add features
5 df = df.groupby(['day'])['passenger_count'].sum()    # aggregation
6 print(df)                                            # use dataframe
```

Pandas Dataframe Program

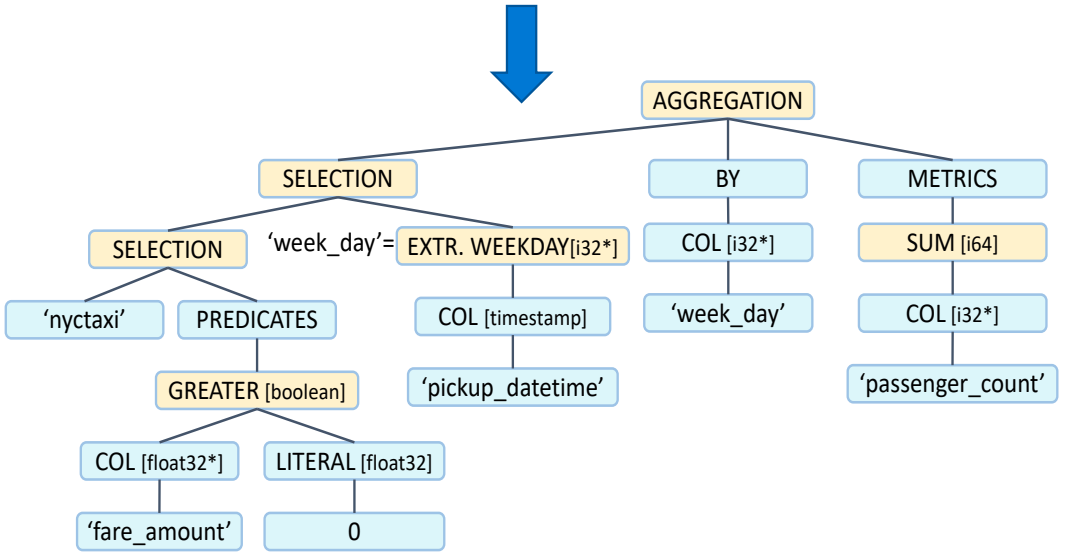

Intermediate Representation

Blue parts: already in IBIS, **Green** parts: our contributions

Pushing Data Science down

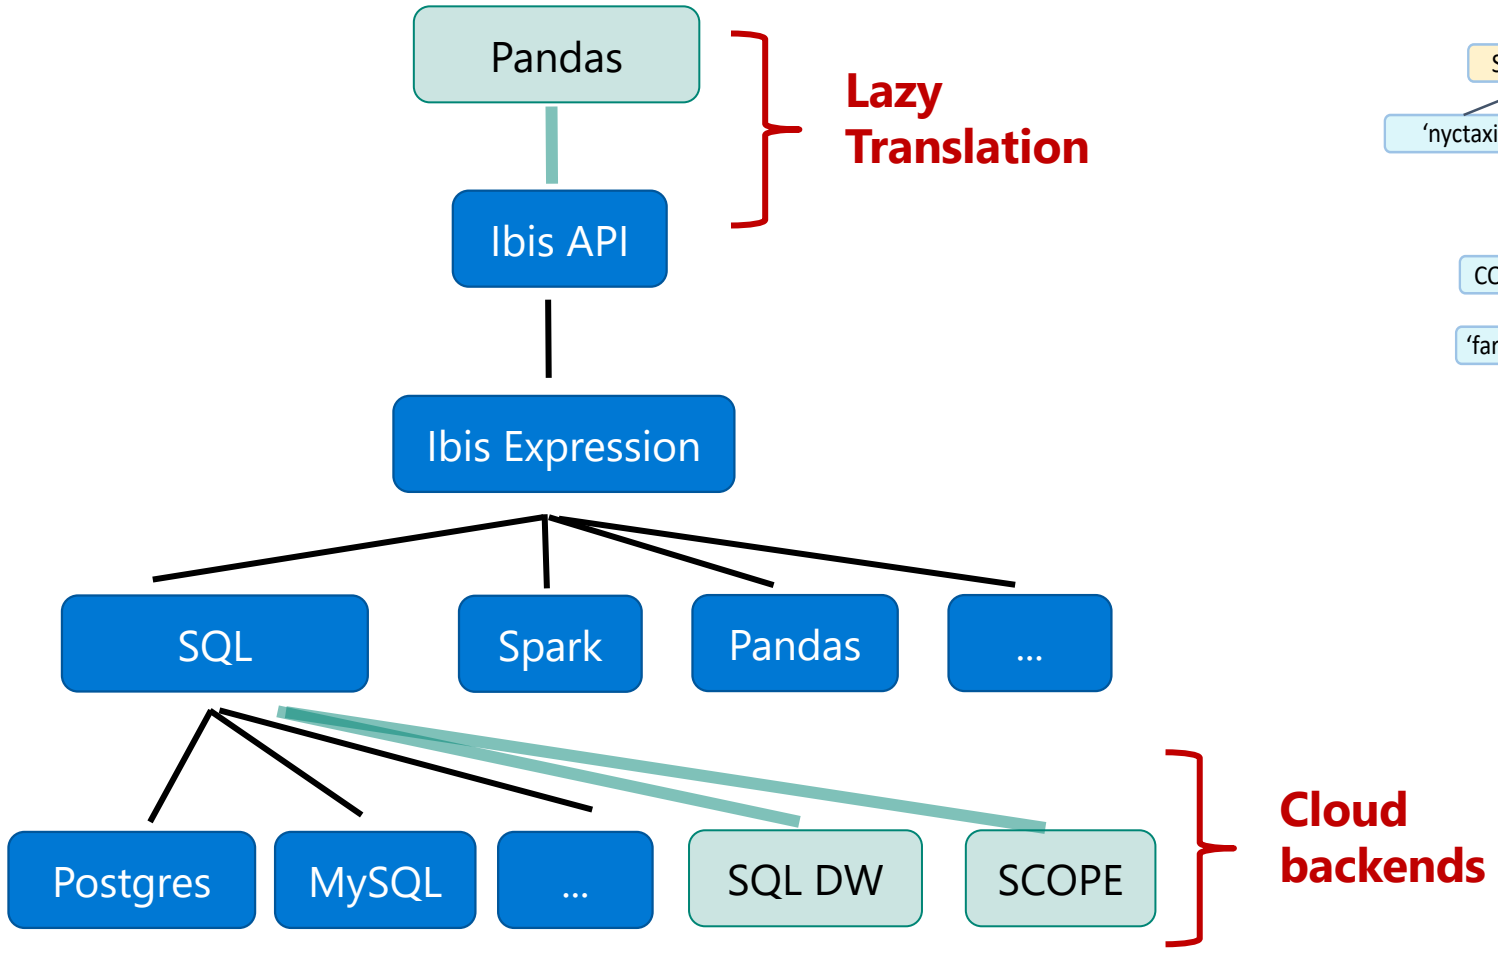

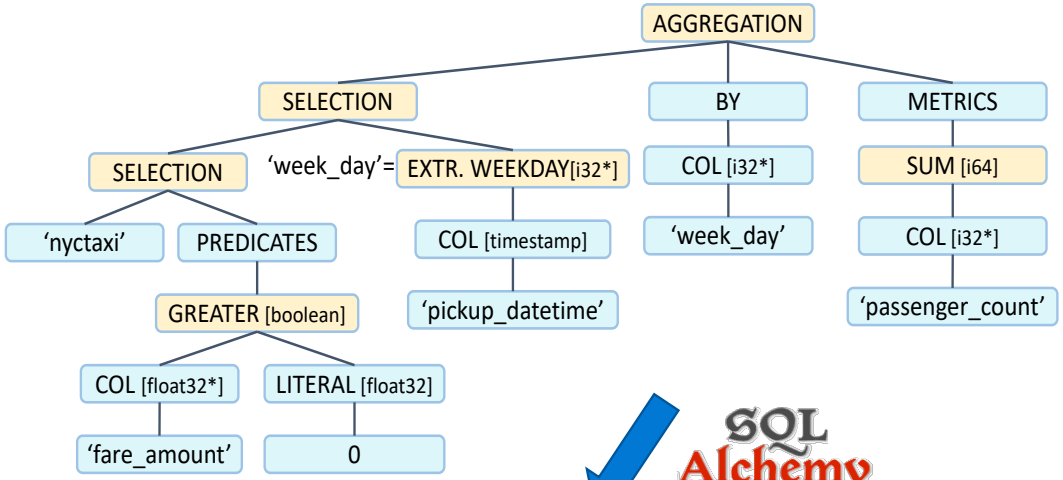

SQL Alchemy

```
SELECT DATEPART(WEEKDAY, pickup_datetime) AS day,
       SUM(passenger_count)
FROM nyctaxi WHERE fare_amount > 0
GROUP BY DATEPART(WEEKDAY, pickup_datetime)
```

T-SQL Statement

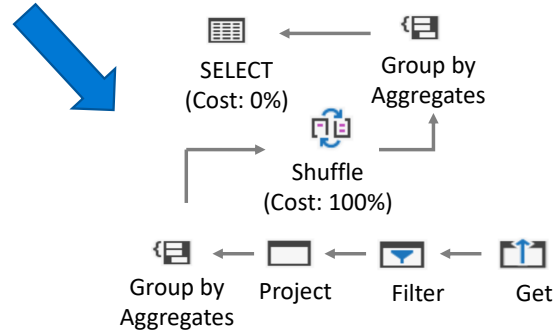

SQL DW Execution Plan

Blue parts: already in IBIS, **Green** parts: our contributions

Pushing Data Science down

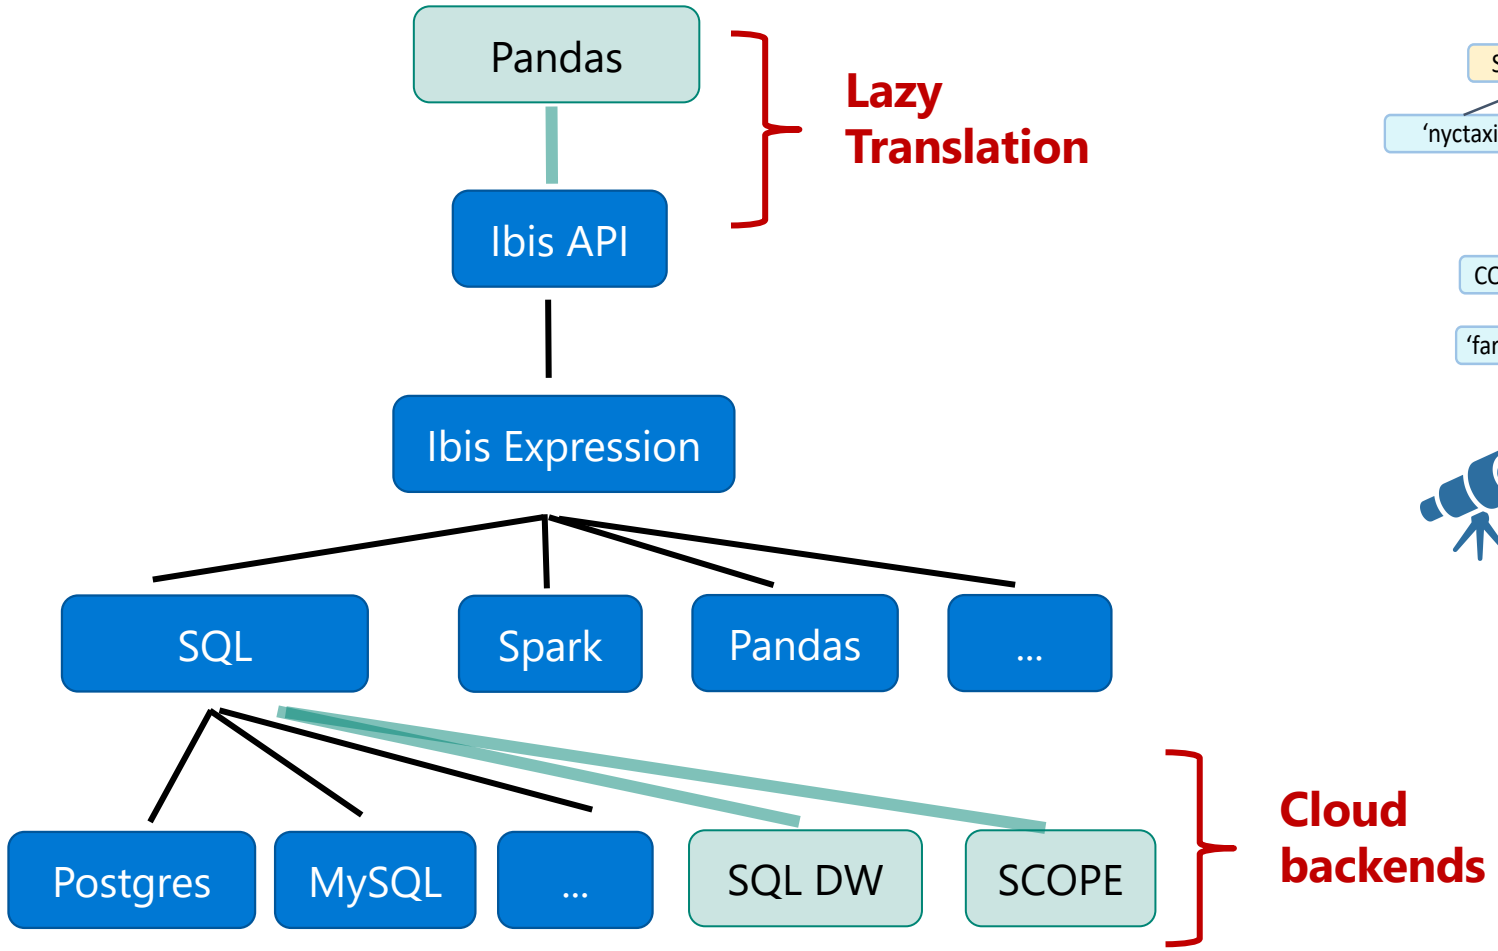

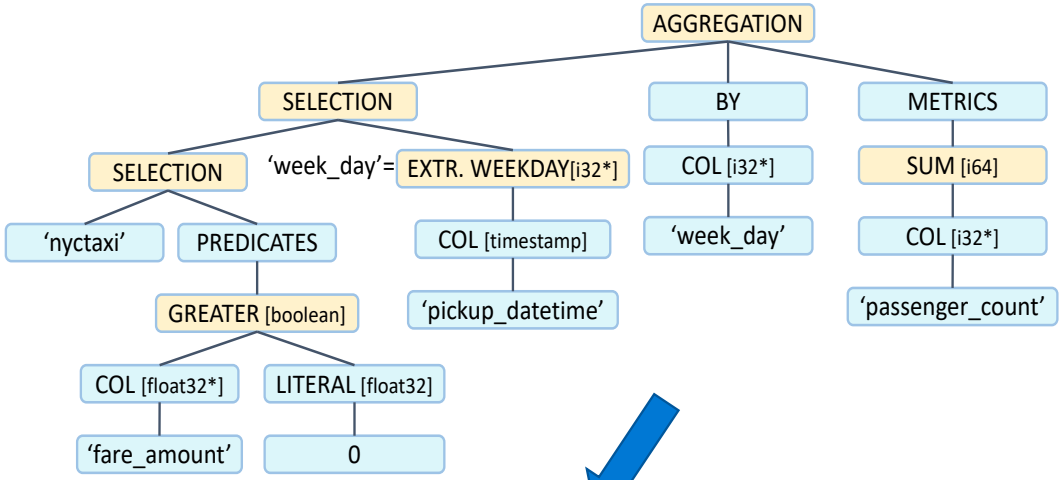

pyscope

```
df = script.extract(path, schema
      .select("fare_amount > 0")
      .groupby("day")
      .project("pickup_datetime
.DayOfWeek.ToString() AS day",
"passenger_count")
```

SCOPE Script

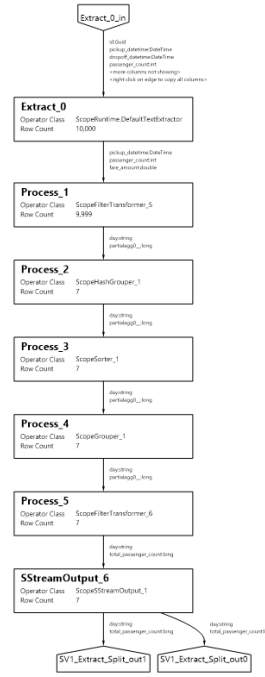

Blue parts: already in IBIS, **Green** parts: our contributions

SCOPE Execution Plan

Impact: speed-up using SQL DW

Growing input size

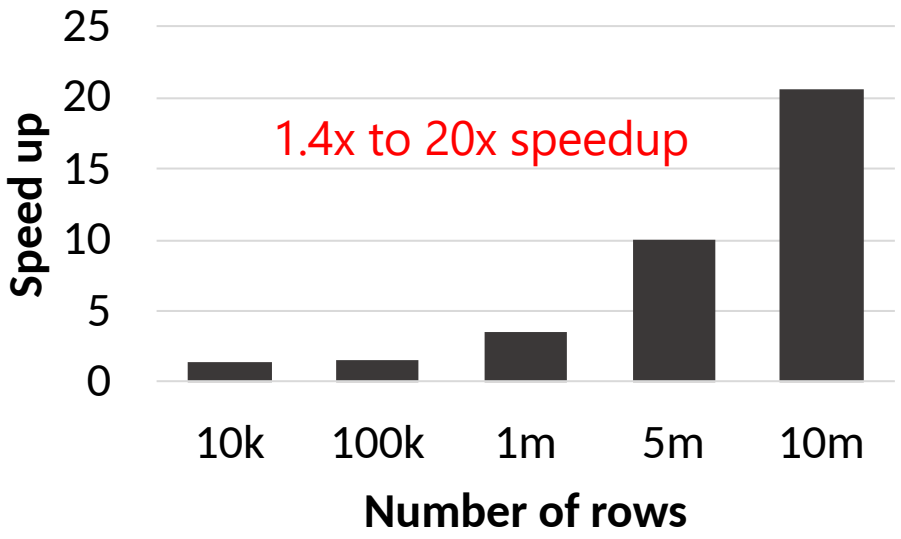

Growing query complexity

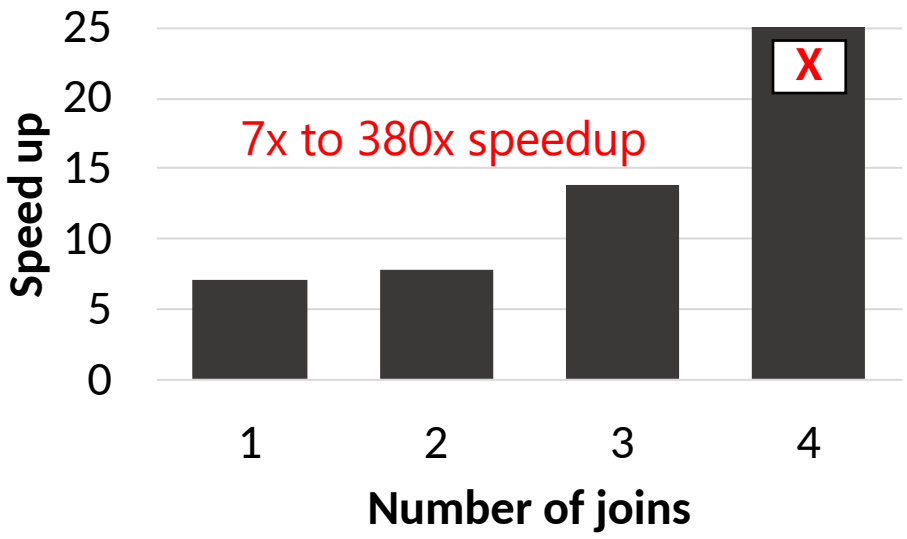

Impact: scale-out using SCOPE

Scale data science to big data!

SCOPE vs Spark

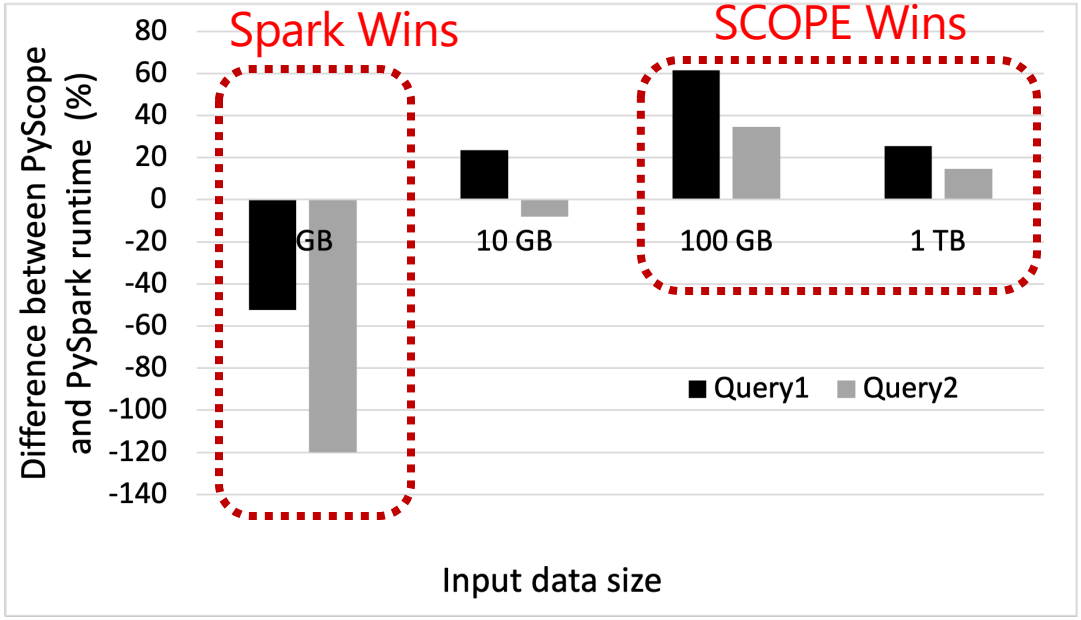

Magpie

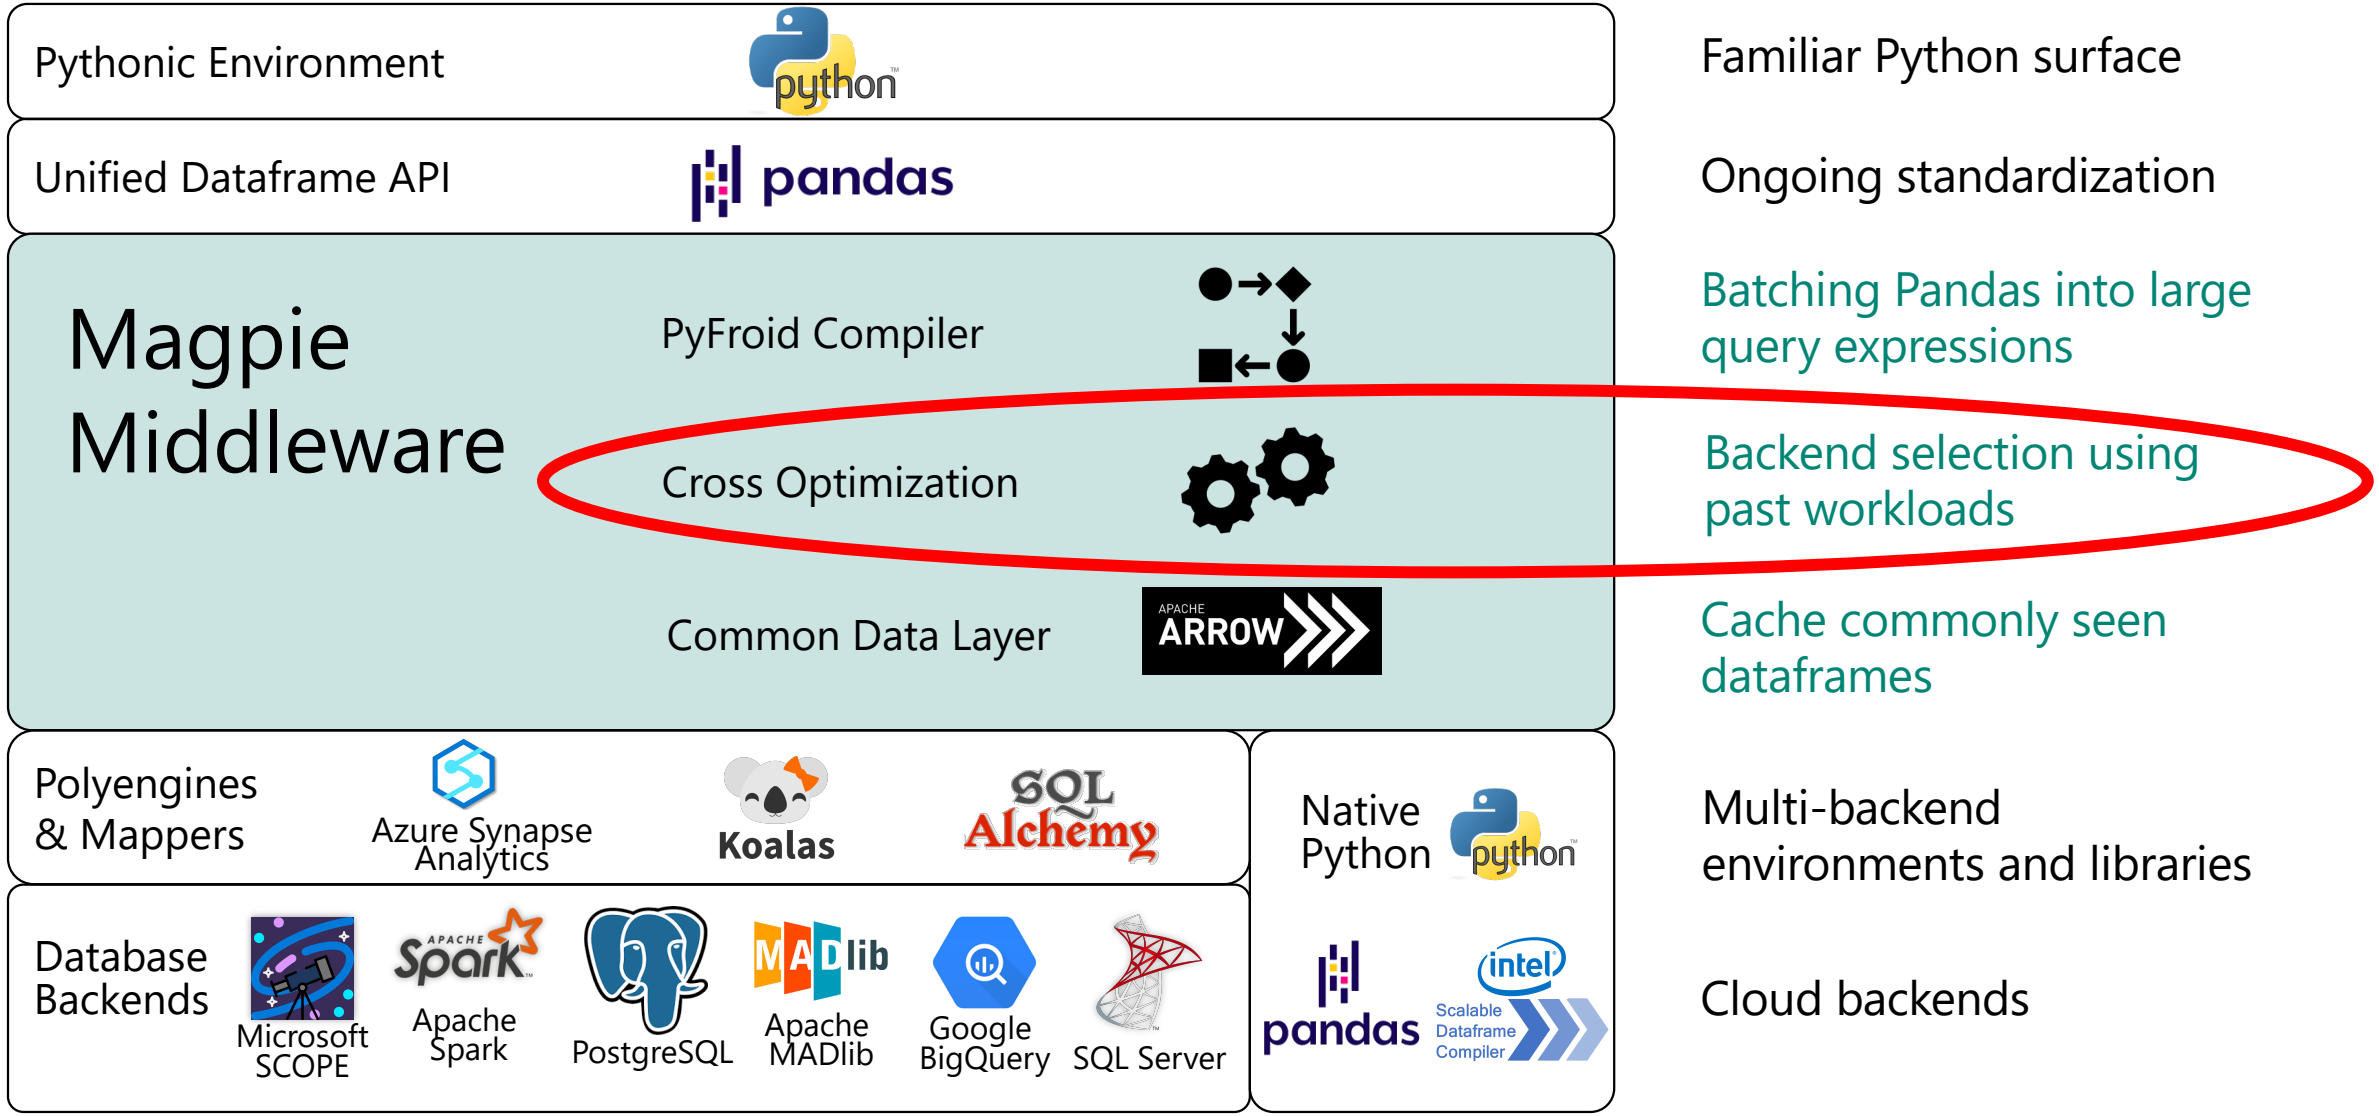
Pythonic Environment
python
Unified Dataframe API
pandas
Magpie Middleware
PyFroid Compiler
Cross Optimization
Common Data Layer
APACHE ARROW
Polyengines & Mappers
Azure Synapse Analytics
Koalas
SQL Alchemy
Database Backends
Microsoft SCOPE
Apache Spark
PostgreSQL
Apache MADlib
Google BigQuery
SQL Server
Native Python
python
pandas
intel
Scalable Dataframe Compiler
Familiar Python surface
Ongoing standardization
Batching Pandas into large query expressions
Backend selection using past workloads
Cache commonly seen dataframes
Multi-backend environments and libraries
Cloud backends

Backend Selection

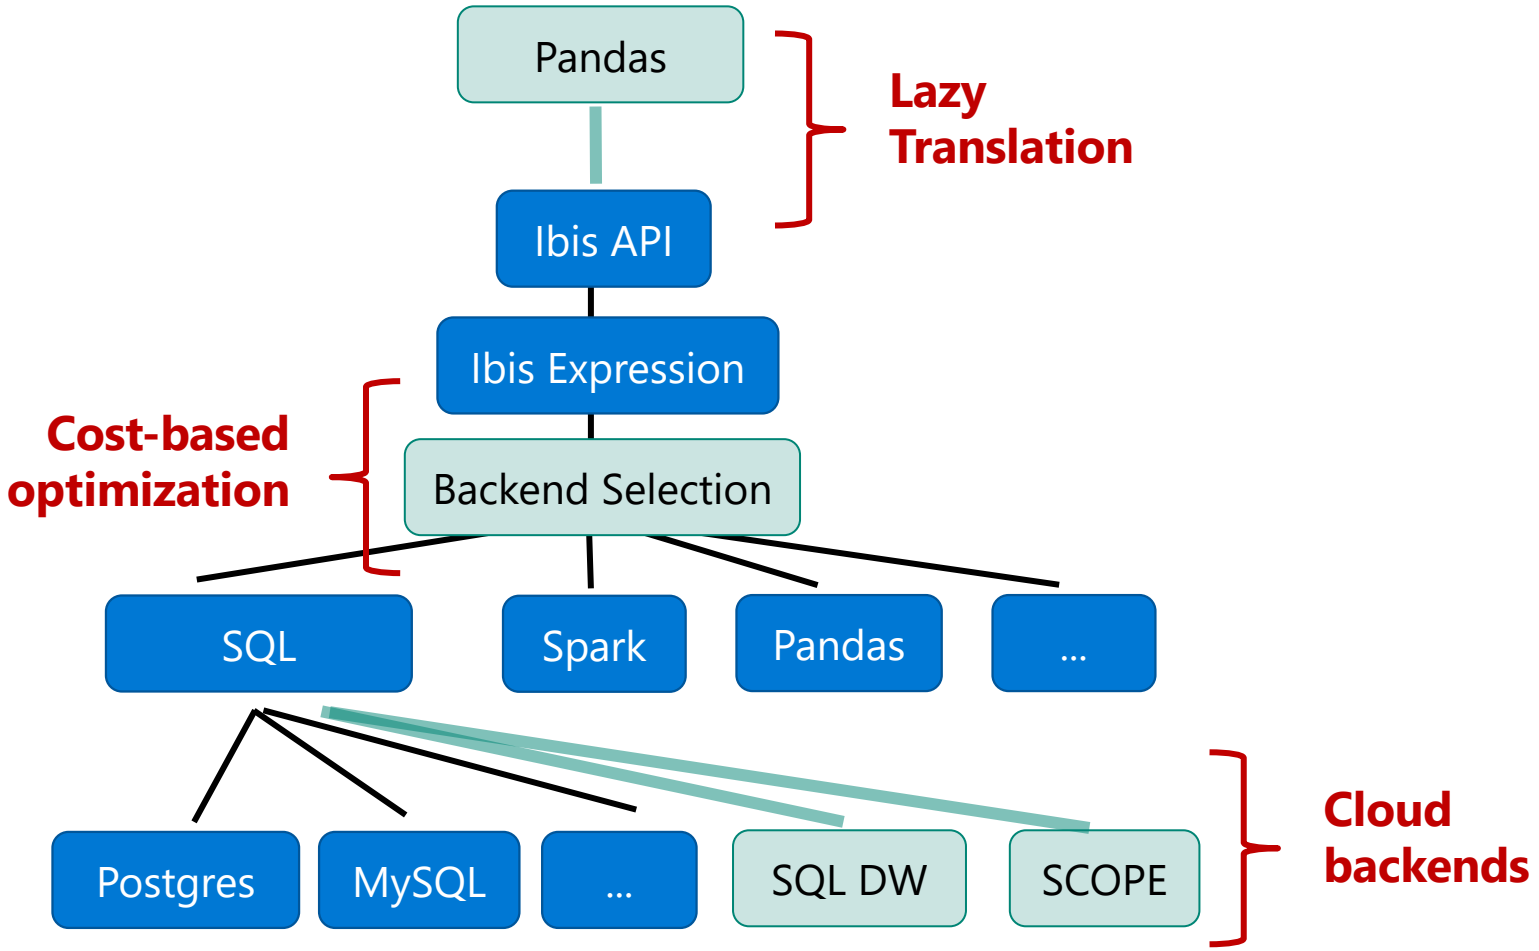

- Leverage past workloads from cloud backends to learn a decision tree
- At compile time:
  - User provides the list of available backends
  - Compile the plan into a common representation
  - Infer best backend using the decision tree

Blue parts: already in IBIS, **Green** parts: our contributions

Scenario 1: Pandas vs PySpark

- Question:
  - When to switch to a cluster?
  - Or to local execution?
- Decision tree:
  - 84% accuracy on test set
  - On Pandas:
    - 84% median improvement
    - Up to 99% improvement

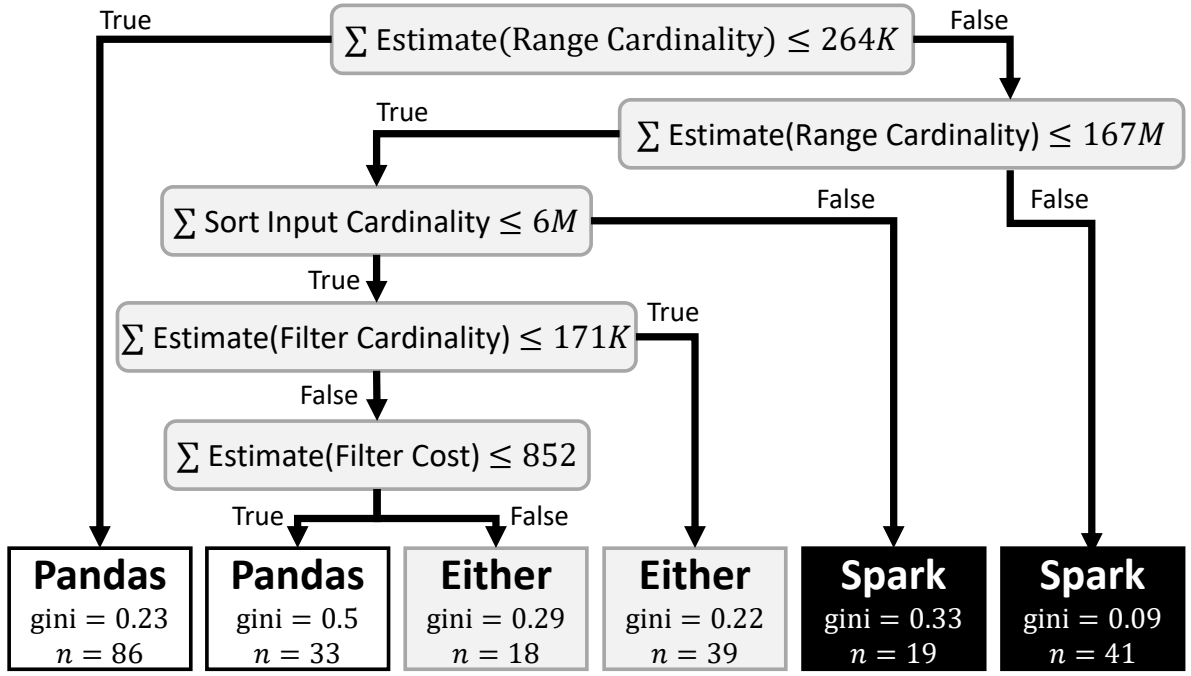

Scenario 2: PyScope vs PySpark

- Cosmos has both SCOPE and Spark engines now
- Question: which one to use for data science?
- Decision tree
  - 87% accuracy on test set
  - On Spark:
    - Median improvement 85%
    - Up to 98% improvement

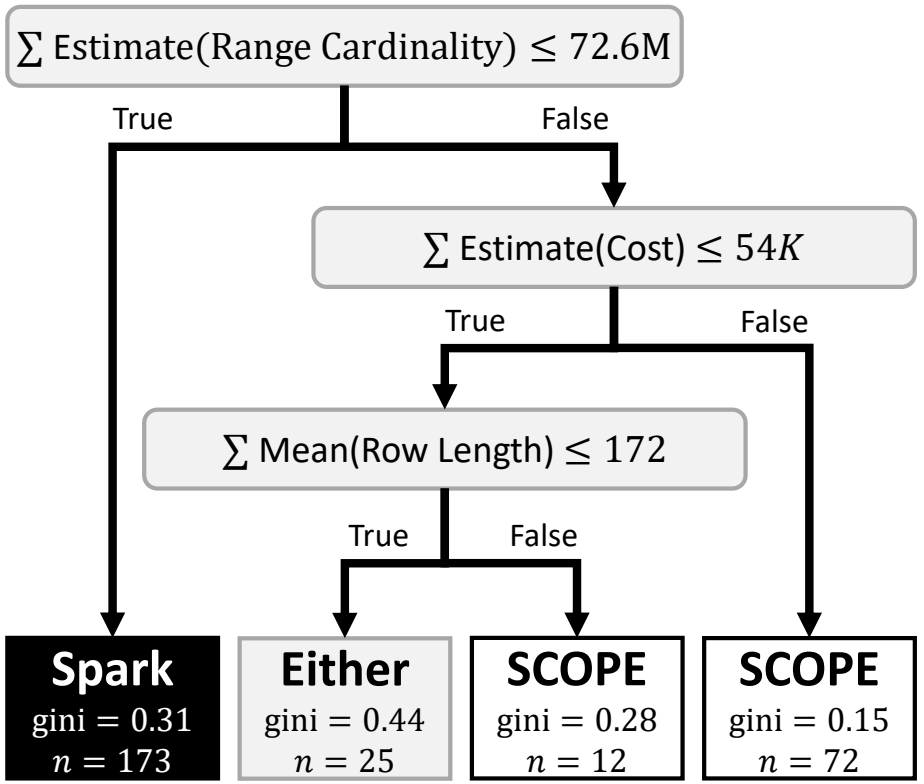

Magpie

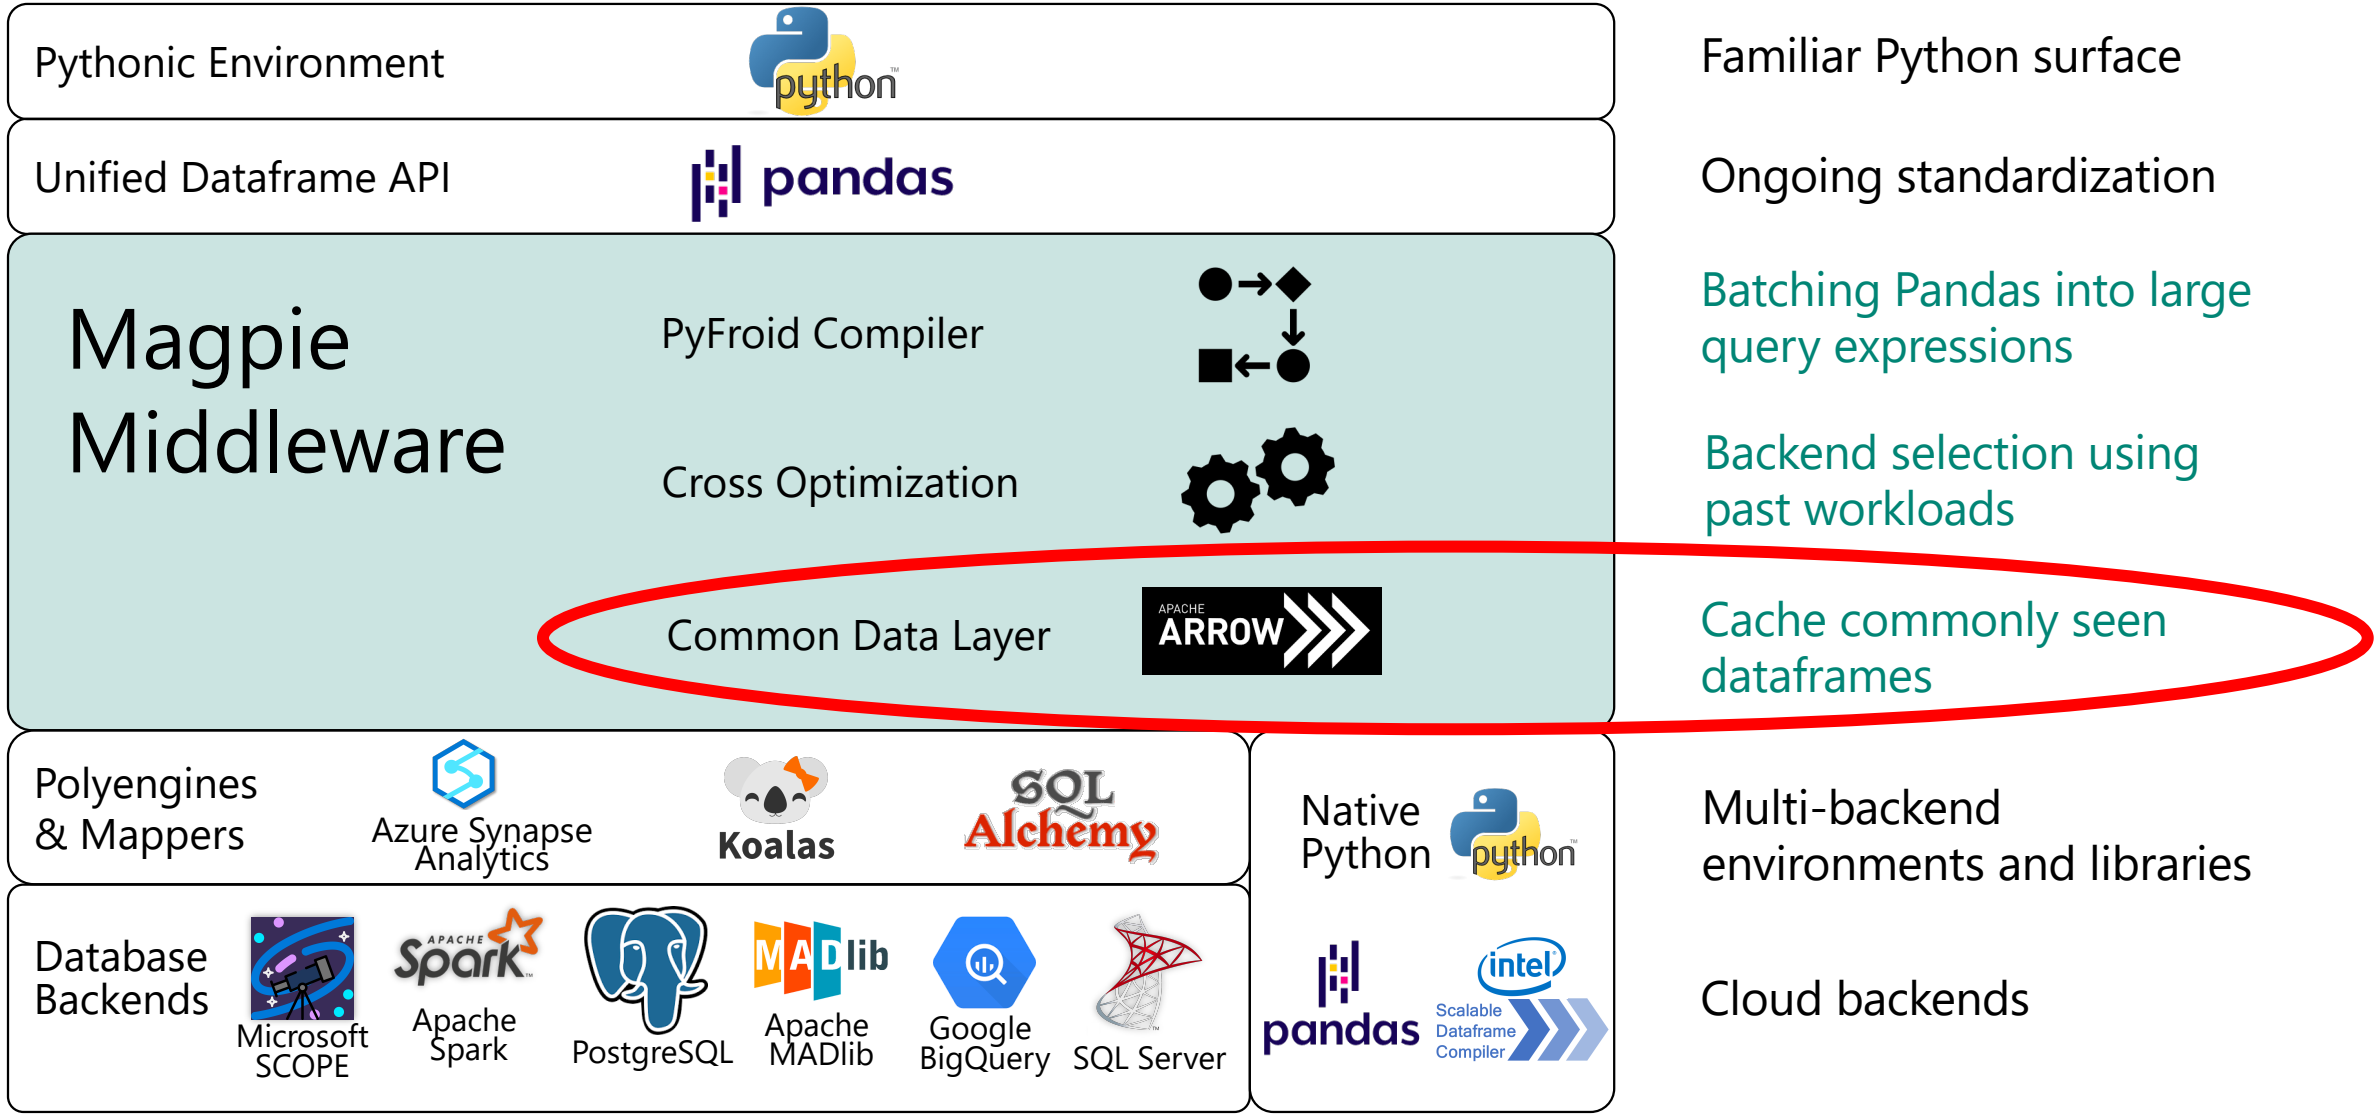
Pythonic Environment
python
Unified Dataframe API
pandas
Magpie Middleware
PyFroid Compiler
Cross Optimization
Common Data Layer
APACHE ARROW
Polyengines & Mappers
Azure Synapse Analytics
Koalas
SQL Alchemy
Database Backends
Microsoft SCOPE
Apache Spark
PostgreSQL
Apache MADlib
Google BigQuery
SQL Server
Native Python
python
pandas
intel
Scalable Dataframe Compiler
Familiar Python surface
Ongoing standardization
Batching Pandas into large query expressions
Backend selection using past workloads
Cache commonly seen dataframes
Multi-backend environments and libraries
Cloud backends

Common Data Layer

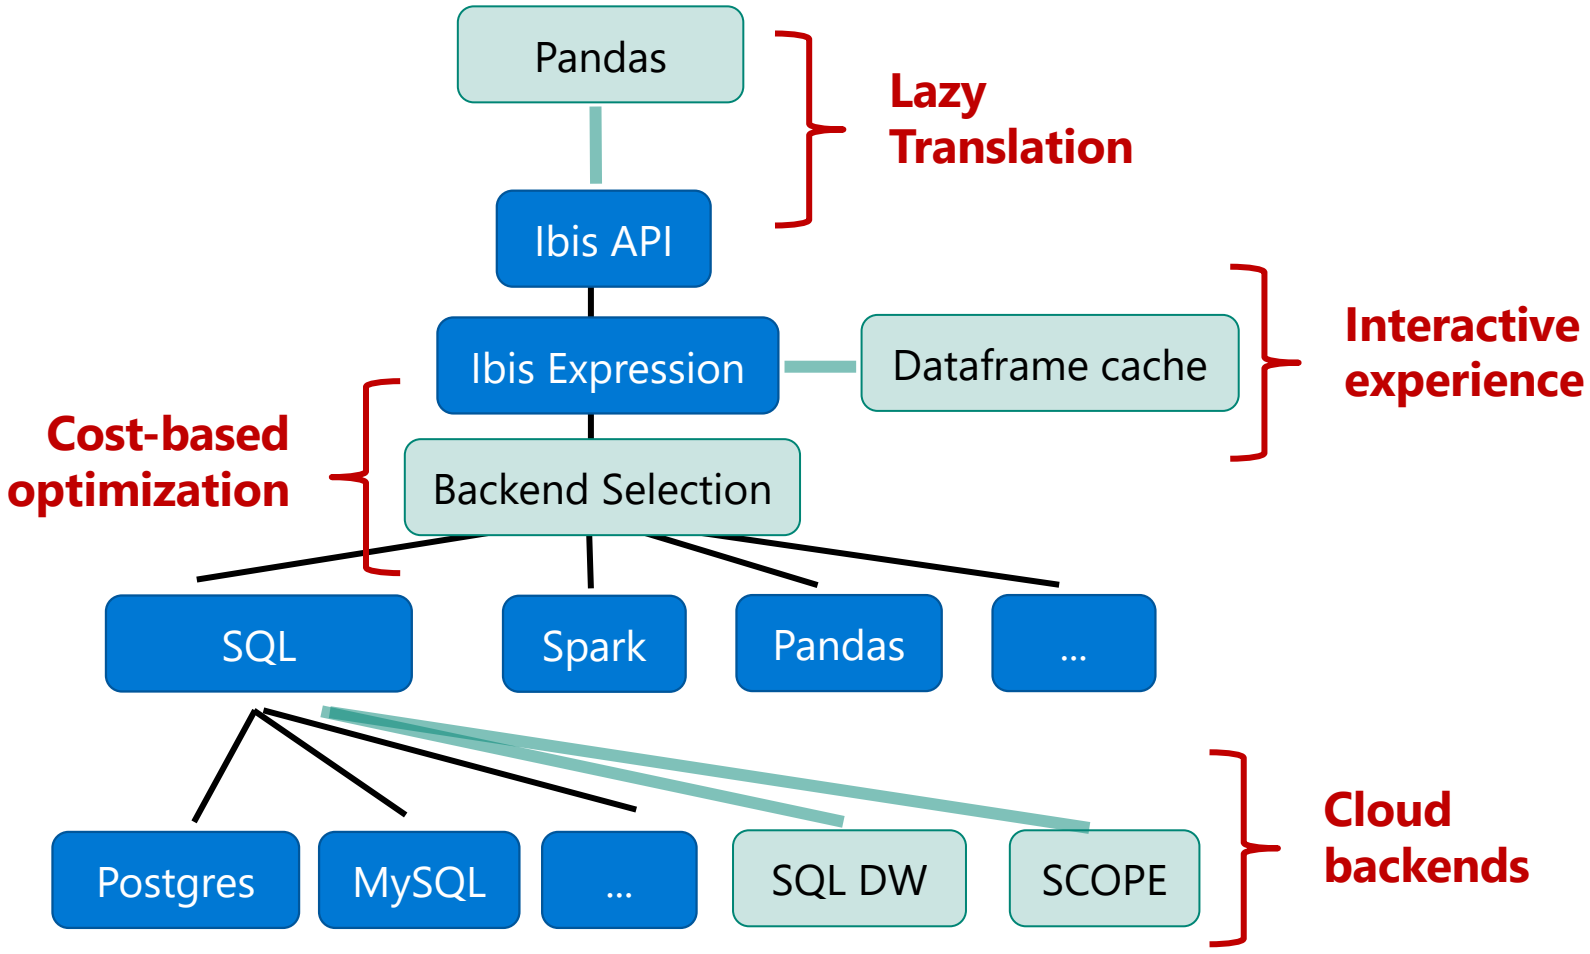

Blue parts: already in IBIS, **Green** parts: our contributions

- Dataframe cache
  - Generate unique signatures
  - Store repeated dataframes in ArrowFlight server
  - Skip accessing the backend in case of cache hit

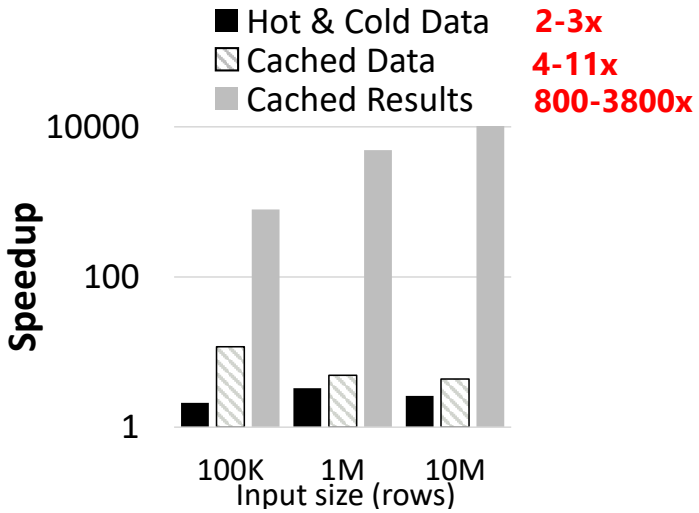

Summary

Pythonic Environment	Lingua franca for many analyses
Unified Dataframe API	Increasingly getting standardized
Magpie Middleware	**Pandas Without Regret!** **Write once, execute anywhere** **Abstracting Data Processing Complexity**
Polyengines & Mappers	From polystores to polyengines
Database Backends	Hyperscale performance Data already in the cloud
Native Python	

Gray Systems Lab

https://azuredata.microsoft.com/

Hiring Summer Interns!

Microsoft