




Polyglot & Hybrid Persistence Architectures for Big Data Analytics

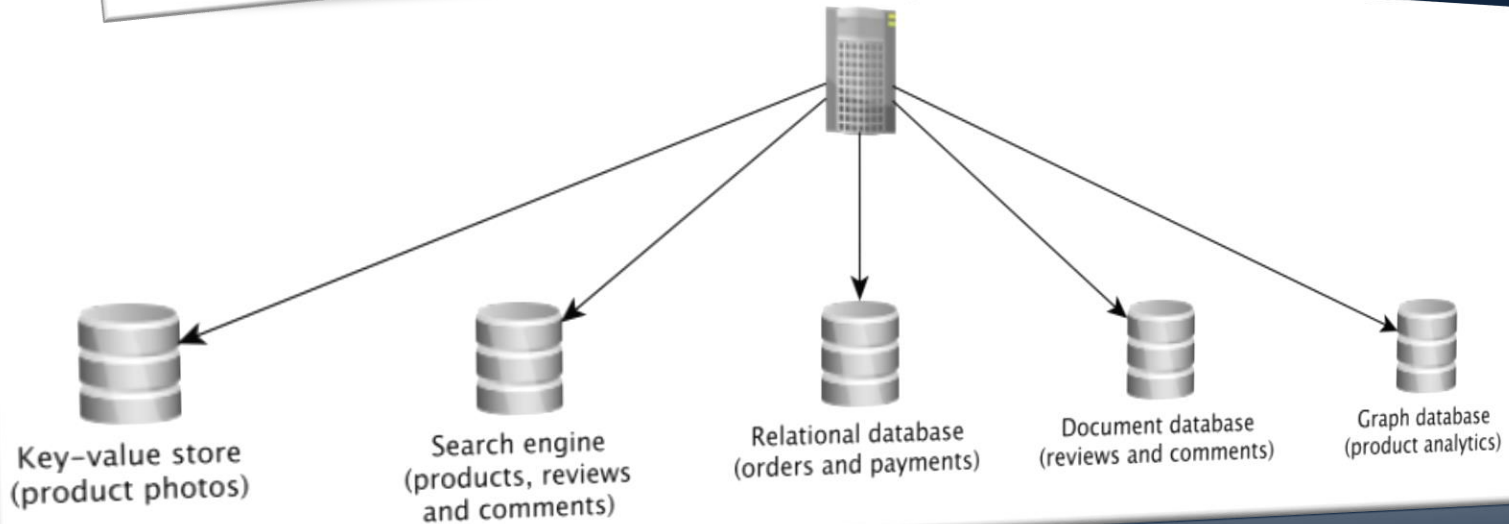
Infocom World Conference, Athens
26 November 2019

Hybrid Data Persistence Architectures


Photos	
Key	Value
p1	
p1	
p2	
...	



e-Commerce system



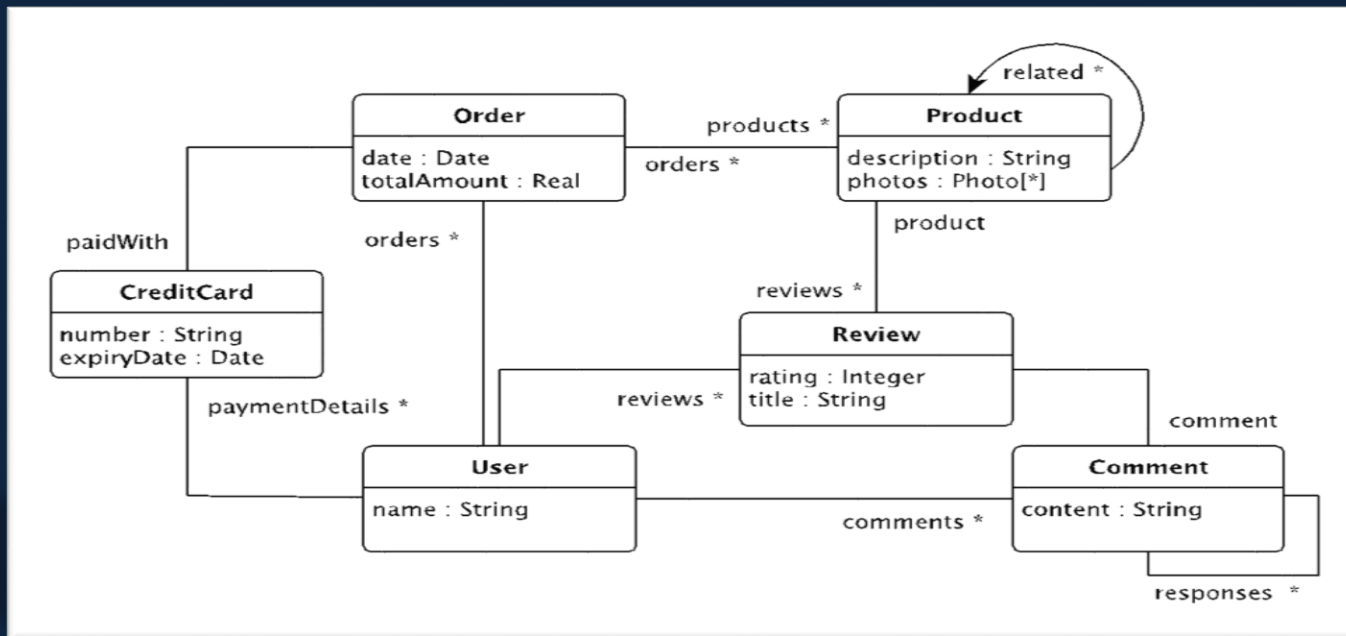
Challenges



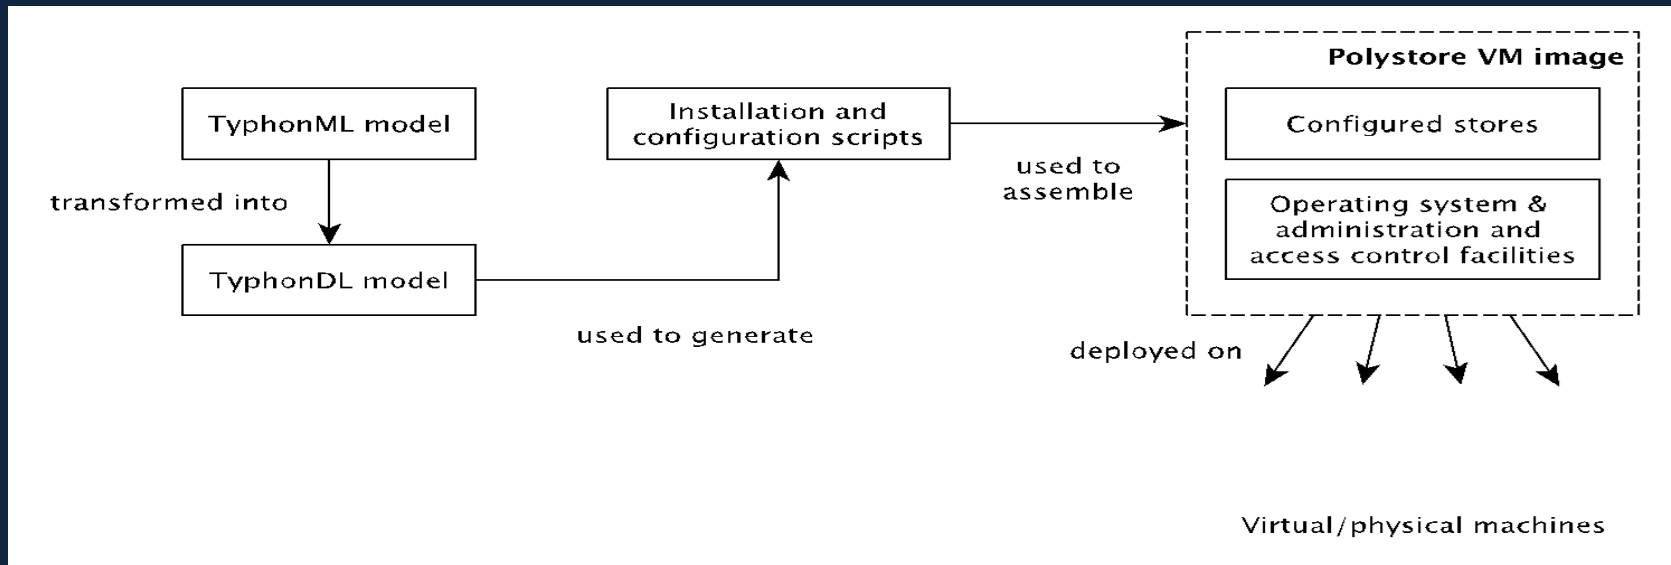
- Design
- Deployment
- Querying
- Monitoring & Analytics
- Evolution

Polystore Design

- Develop a tailored modelling language (TyphonML)
- Support for capturing availability, consistency & partitioning requirements
- Supporting modelling tools
- Extensibility to accommodate future data stores



Polystore Deployment

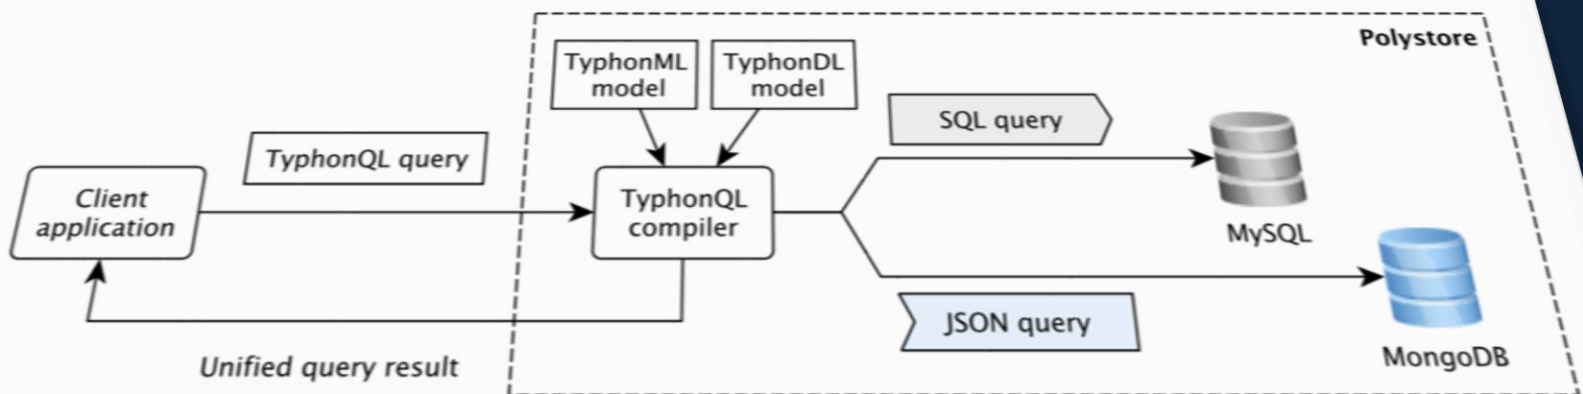


- Develop a language for modelling hybrid polystore deployments (TyphonDL)
- Develop automated TyphonML→TyphonDL transformations
- Develop automated transformations of TyphonDL models to installation & configuration scripts

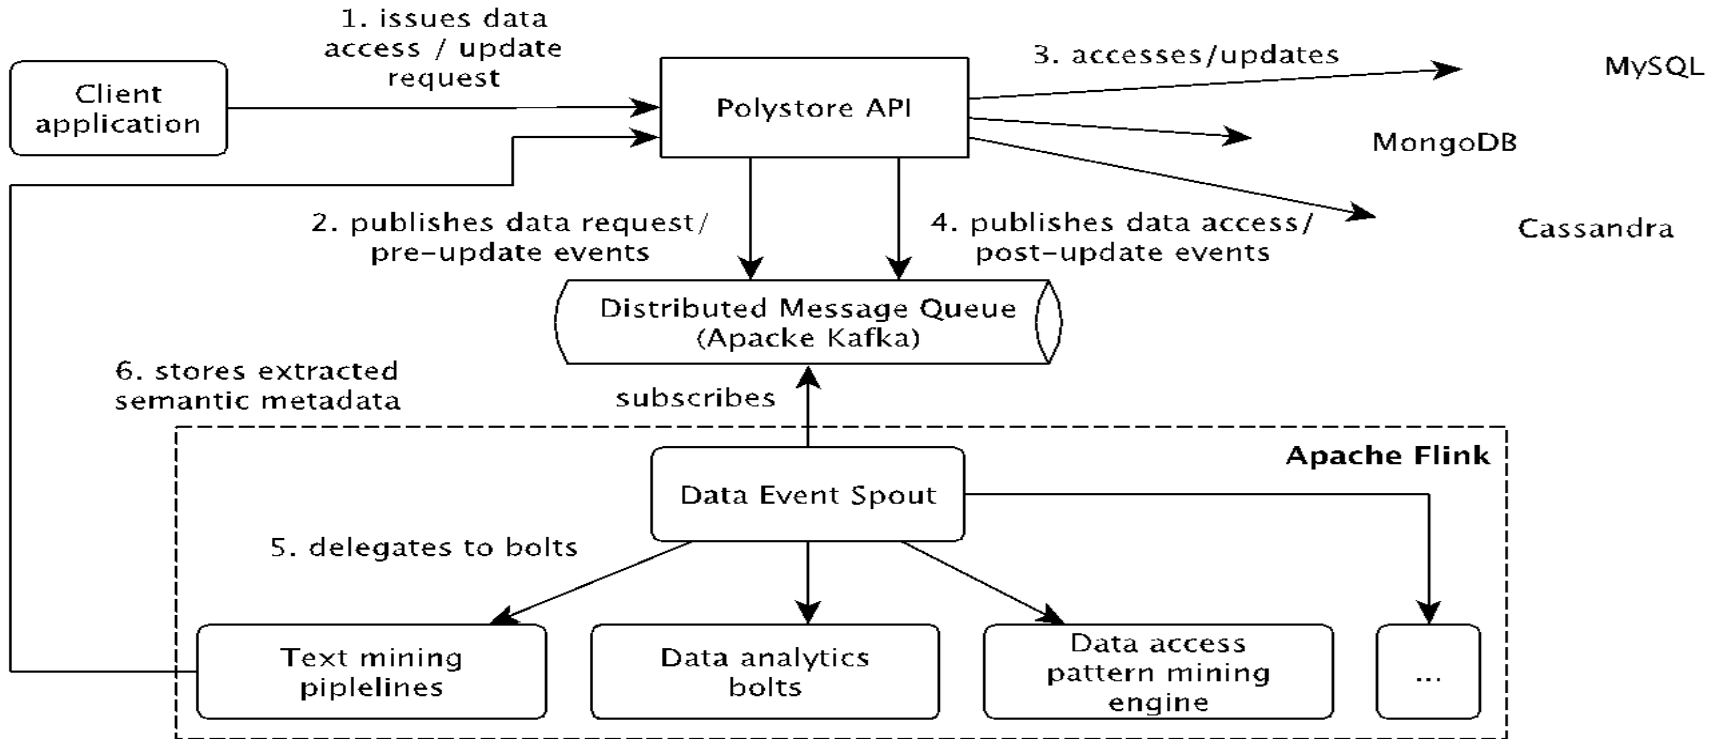
Polystore Querying

- Unified language for querying data in polystores
- Compilation to native querying facilities
- Static analysis, lazy loading and prefetching

```
find c:Customer where  
c.orders.totalAmount.sum() > 1000 and  
c.reviews.comment.responses.size() > 20
```



Polystore Monitoring & Analytics



- Record data access and update events
- Publish data events to distributed ledger
- Facilitate subscription-based data analytics & text mining pipelines

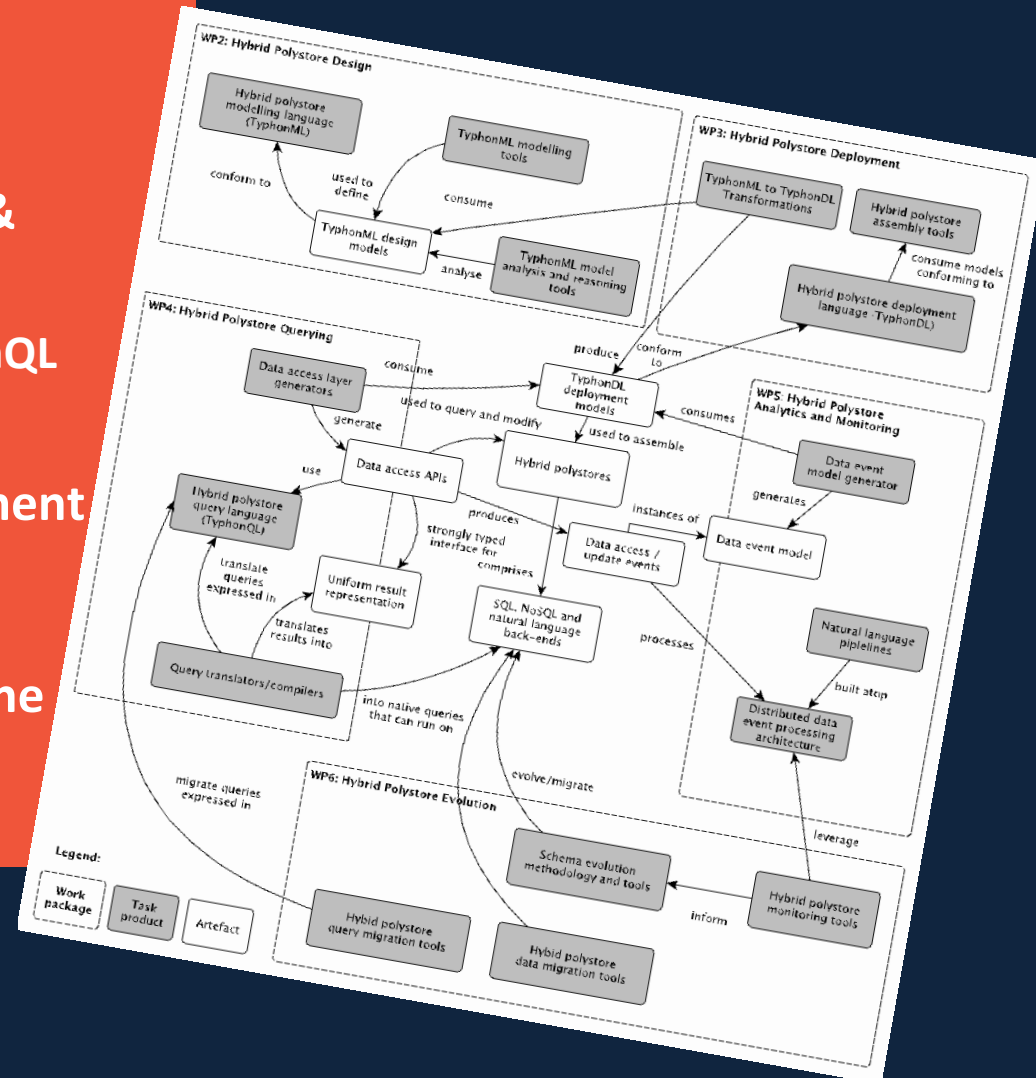
Polystore Evolution



- Support for schema evolution
- Support for internal data migration
- Query migration tools
- Data usage monitoring tools to recommend evolutions

Technology Stack

- Java as the main programming language
- Eclipse as the base modelling & development platform
 - ❖ TyphonML, TyphonDL, TyphonQL editors & tools
- Docker for container management
- OpenAPI for REST APIs
- Open-source software under the Eclipse Public License 2.0





OTE's Contribution in Typhon

❖ Data for Predictive Maintenance and QoS

- Technology Requirements Designation
- Big Data Provision
- Typhon's Technology Deployment/Evaluation

❖ Dissemination and Communication Activities

- Publication of project news in internal networks



The Use Case ...

Why Predictive Maintenance and QoS?

AI-based Operations'
Ecosystem

- Minimize Maintenance Costs
- High Availability
- Service Reliability

The System

(Use Case Description – CRM System Functionality)



Use Case Objectives

Breakdown

**Descriptive
Reporting for
Prediction
Accuracy**



**Predictive
Maintenance for
improved QoS**

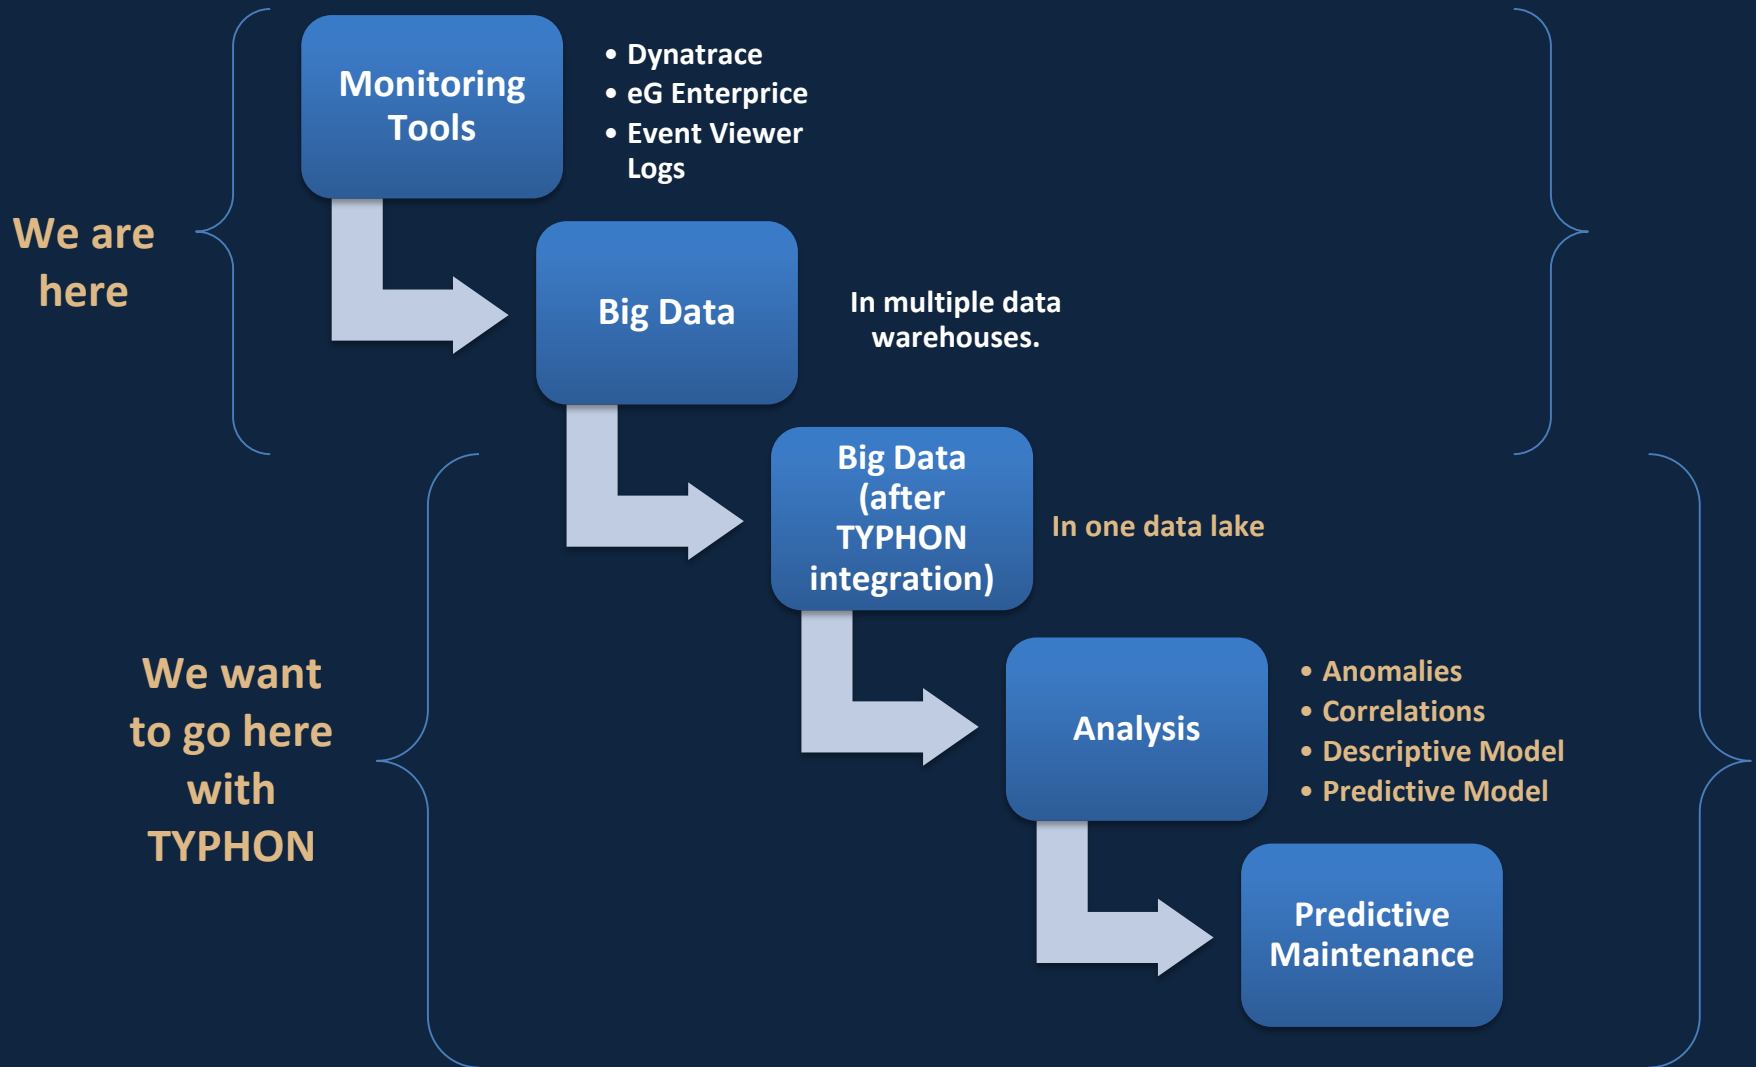
Identification of:

- Quantitative Variables
- Categorical Variables
- Correlations
- Anomalies

Implementation of:

- Prediction Models
- Training/Testing
- Evaluation

Our Approach ... in *Steps*



Information Flow

Data Sources

Human-generated Data 

Synthetic Data 

Agent Data 

Machine Data 

Wire Data 

Collect










Kafka



- TopicName
- TopicName
- TopicName
- TopicName

Preprocess

HDFS 

Store/
Process/
Analyze



OTE @ TYPHON EU Project

Q & A

An aerial night view of a modern, multi-story building with a grid-like facade. The building is illuminated from within, with warm yellow lights visible through the windows. The building's outline is highlighted with glowing green lines at the top and bottom, and glowing blue lines on the sides. The text "Q & A" is overlaid in large, white, sans-serif font in the center of the image. The background shows a cityscape at dusk, with other buildings and mountains in the distance under a dark sky.

Thank you!

