# Troubleshooting real-time business metrics
## *with StarTree ThirdEye*

Madhumita Mantri
Product Lead
StarTree

Suvodeep Pyne
Founding Engineer
StarTree

startree

“Real-time business metrics are constantly changing, so it's important to be proactive in troubleshooting issues. Don't wait for a problem to become a crisis before you take action.”

# Agenda

- Challenges with Troubleshooting real-time business metrics with StarTree ThirdEye
- Introduction to Real-time Monitoring, Anomaly Detection and Troubleshooting
- Key Use-Case and Live Demo
- Benefits & Advantages
- Q&A Session

# Challenges with troubleshooting real-time business metrics

- **Data Volume and Velocity**: Handling in real-time requires efficient processing and storage solutions.
- **Data Quality and Consistency**: Ensuring accurate and consistent real-time data can be challenging, impacting troubleshooting decisions.
- **Data Integration**: Integrating data from diverse sources in real-time is complex, with discrepancies hindering troubleshooting.
- **Latency and Scalability**: Minimizing latency for real-time monitoring and scaling systems as data complexity and volume grow are ongoing challenges.

# What is the Solution? Anomaly Detection!


## Monitor Critical KPIs


### Unusual patterns or outliers!

- Unexpected **spike** or **dip**
- **Gradual** change over a period of time

### Caused by

- **User behavior changes**
- **Errors in data collection or processing**
- **Underlying system or process**
- **Seasonal fluctuations (ex: holidays)**
- **External events**

# What it takes to Automate Anomaly Detection?

**Monitor KPIs (Key performance indicators)**

**Identify anomalies (Proactively)**

What went wrong?

**Find impact (Insights)**

Why went wrong?

**Archive/Live/Streaming Data**

**Descriptive/Predictive/Diagnostic/Prescriptive Analysis**

**Actionable Insights**

# What are the different Anomaly Scenarios in business metrics?

## Stable Time Series

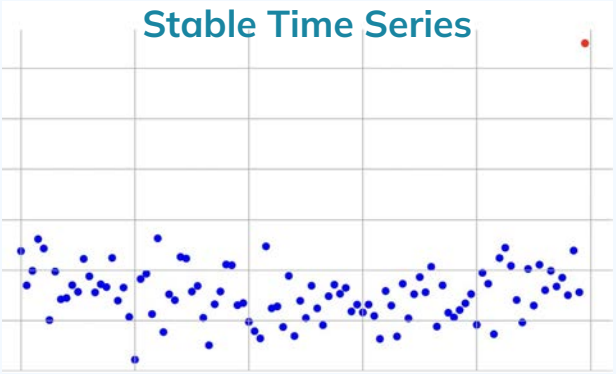

## Uptrending Time Series

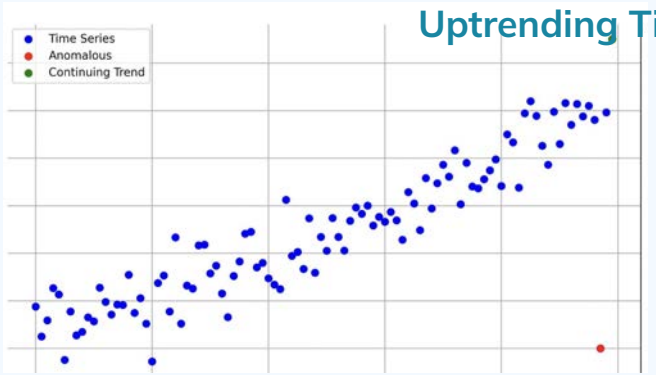


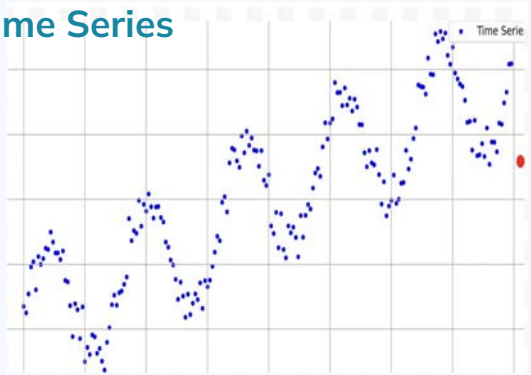


## Time Series with Seasonality

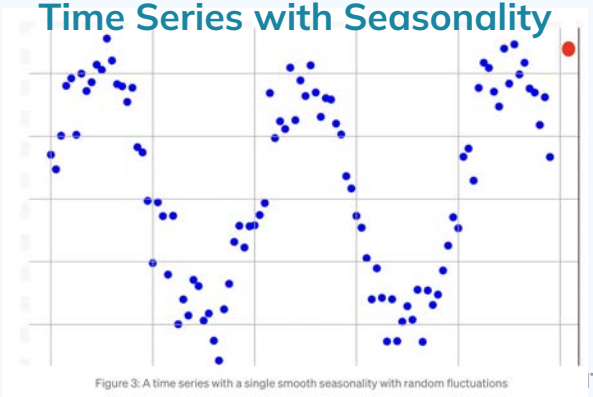

Figure 3: A time series with a single smooth seasonality with random fluctuations

TIAL.

## Time Series With Spikes At Regular Intervals

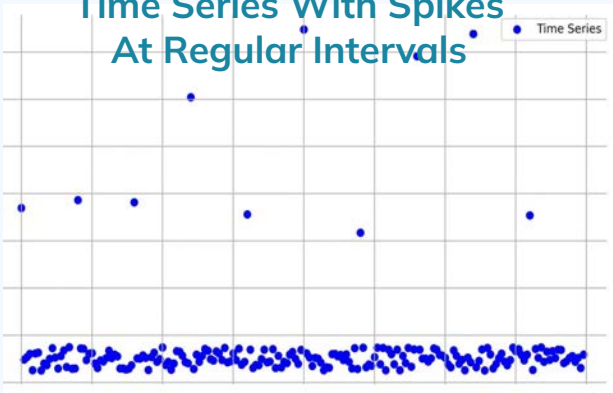

# StarTree ThirdEye

“Real-time monitoring and anomaly detection” of large and complex time series data and fast tracks problem solving by unlocking actionable insights.

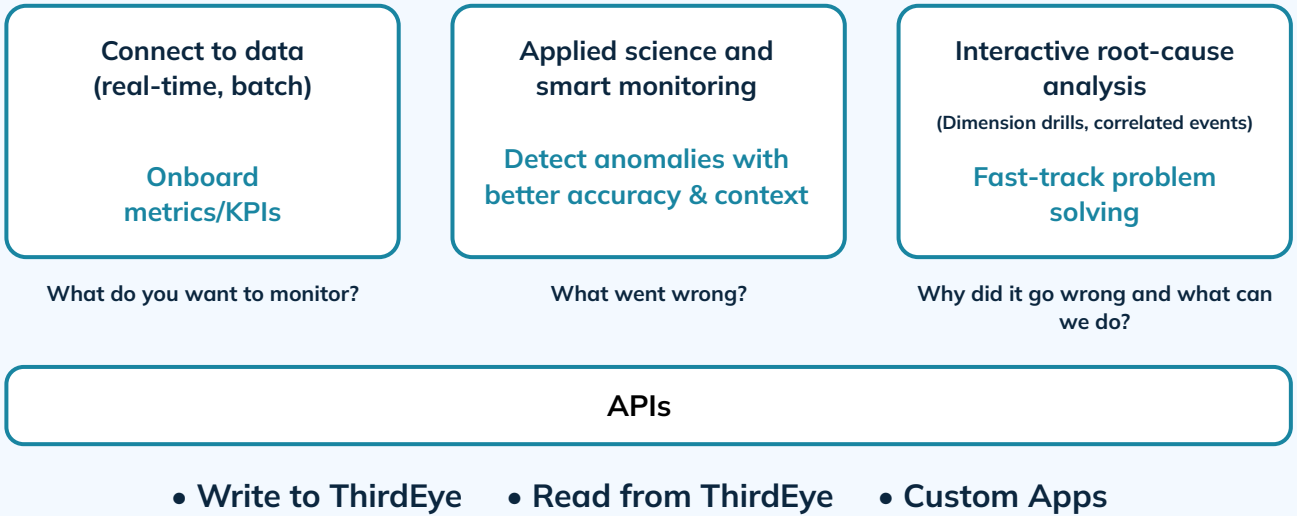

# Use-Case & Demo

# Use Cases

- Rideshare
- Sensordata

# Under the hood

# Data Modeling for real-time analysis

- Goal
  - Ensure metric retrieval is fast and efficient
  - Especially for lower time granularities like 1m, 5m
- Use a Time Series Database: **Apache Pinot**
- Partitioning: Managing your time column
  - **Datatyping**: Strings are harder to process than int/long
  - **Time Buckets**: Reducing expensive transforms at runtime
- Indexing
  - **StarTree Index**: precompute metrics across dimensions

# StarTree Advantage: Vertical Integration


- Real-time Time Series DB
- Optimized for aggregations
- Powerful Indexing Capabilities
- High Throughput


- High Query Volume
- Low Latency
- Aggregation Heavy
- Powerful slicing/dicing capabilities

# Flexible DAG Workflows

- Leveraged by detection pipeline
- Enabling
  - Derived Metrics
  - Post Processing Filters and Business Use Cases
- Multidimensional Workflows
- Dimension Recommender

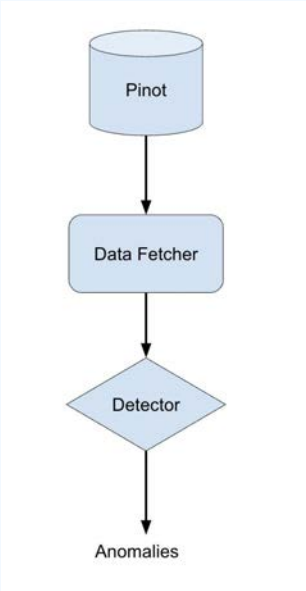


Simple Alert

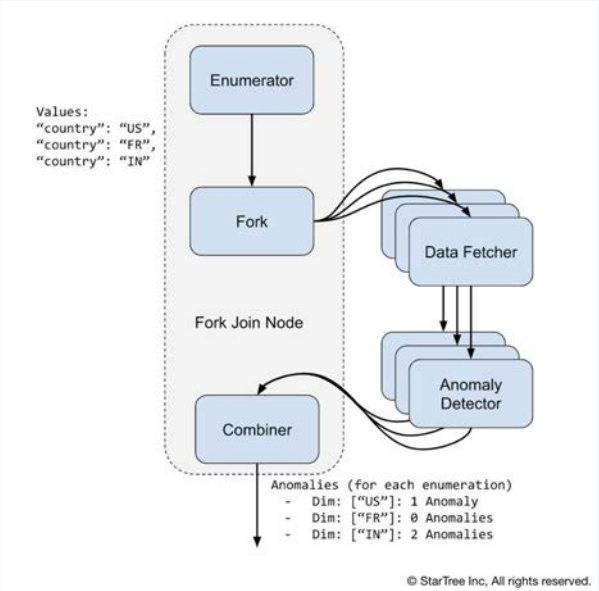


Multidimensional Alert

# Notifications

Email

Slack

Pagerduty
(2024)

also, build your custom integrations with **webhooks**

# API First

- Leveraged by react.js UI
- Easily perform batch operations
- Easily integrate with your internal systems
- Security

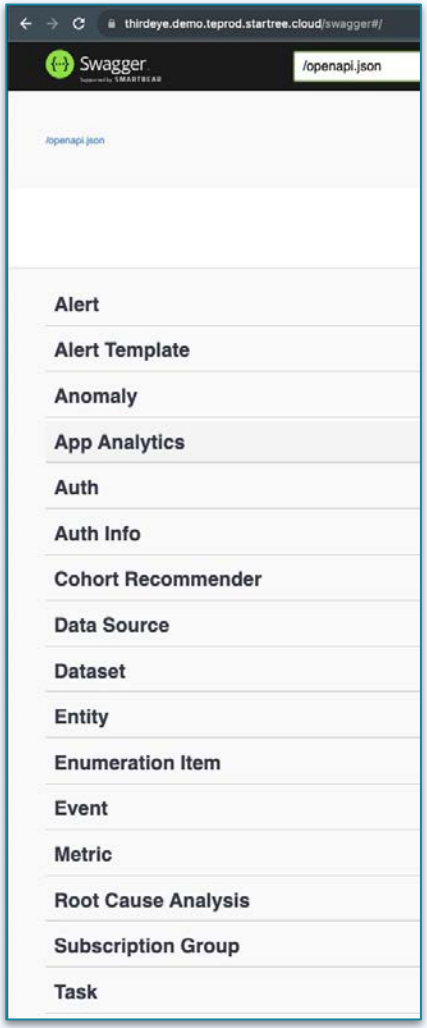

# Cloud Ready with K8S

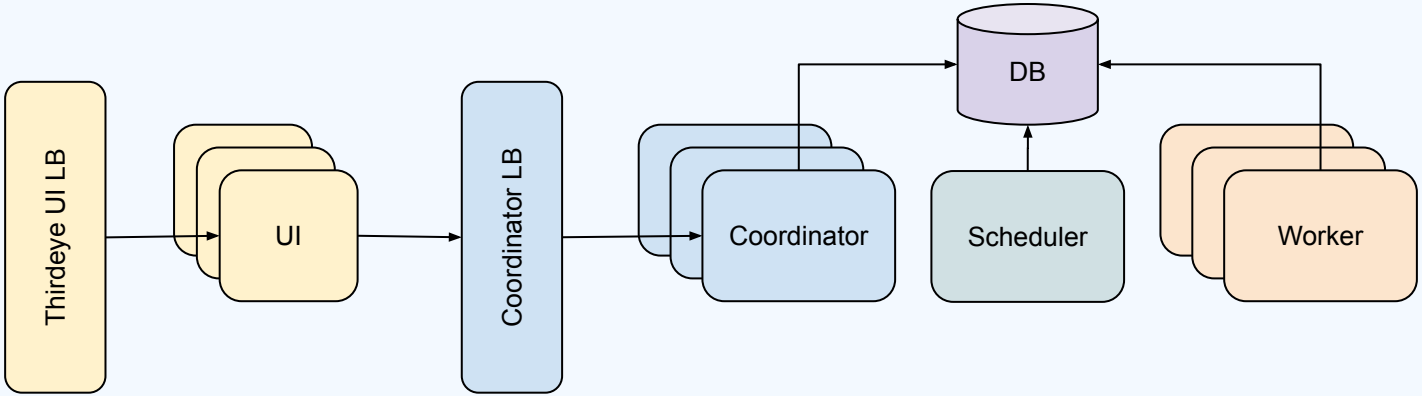

# StarTree ThirdEye Community Edition

github.com/startreedata/thirdeye

# Thank You!

Join StarTree Community

StarTree ThirdEye Community Edition

Try StarTree Enterprise Self-Serve