PDUG

Gdansk, 2-3 June 2026

BEYOND STATIC THRESHOLDS

AI Anomaly Detection for Db2 for z/OS

Jørn Thyssen

jthyssen@rocketsoftware.com

Senior Technical Staff Member, Rocket Software

POLISH DB2 USERS GROUP

Rocket Software Presentation Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON ROCKET'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY ROCKET WITHOUT NOTICE. ROCKET SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR NAY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM ROCKET (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF ROCKET PRODUCTS OR SOFTWARE.

Rocket's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at Rocket's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at out sole discretion.

Rocket and Rocket logo are trademarks of Rocket Software Inc. ("Rocket Software"), registered in many jurisdictions worldwide. Other product and service names might be trademarks of Rocket Software or other companies. A current list of other Rocket Software-owned trademarks and guidelines of the permitted use are available at https://www.rocketsoftware.com/company/legal/trademarks

Other company, product, or service names may be trademarks or service marks of others.

Context - Predictive AI alerts - **Generative AI**

GENERATIVE AI

SYNTHETIC CREATION

VS

PREDICTIVE AI

Using Machine Learning technics to predict based on historical data to provide forecasting.

Context - Predictive AI alerts - **Predictive AI**

Aspect	Generative AI	Predictive AI
Primary Goal	Create new content based on learned patterns	Estimate outcomes or probabilities
Typical Tasks	Text, image, code, and audio generation	Classification, regression, forecasting
Inputs	Prompts, context, documents, examples	Structured features and historical labels
Outputs	Drafts, answers, summaries, creative assets	Scores, probabilities, numeric predictions
Evaluation Metrics	Quality, helpfulness, safety, human feedback	Accuracy, AUC, RMSE, MAE
Common Models	Transformers, large language models, diffusion	Linear models, tree-based models, time-series

Rule of thumb: Predictive AI says "something is wrong / will be wrong"; Generative AI says "why it matters and what to do next."

When to Use Each Approach and How They Work Together

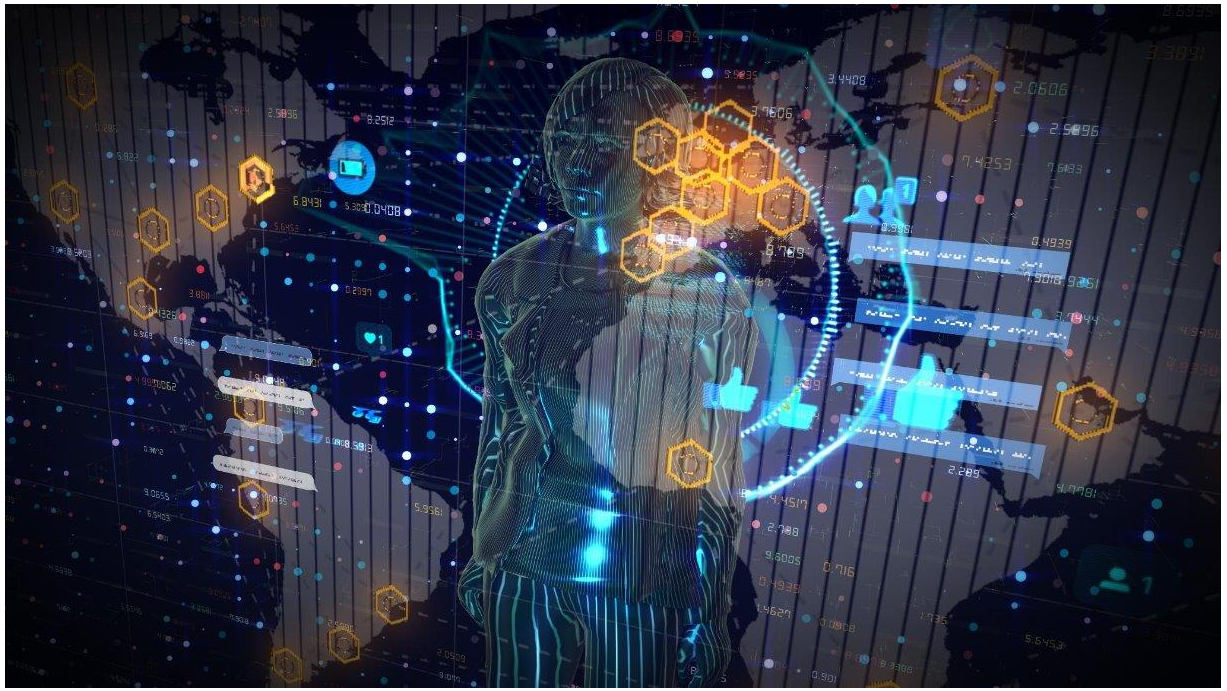

Predictive AI Use Cases

Predictive AI supports numeric decisions like risk scoring, forecasting, and **anomaly detection with consistent and explainable metrics.**

Generative AI Applications

Generative AI excels at communication, synthesis, and creativity tasks like drafting reports and conversational interfaces.

Combined AI Workflow

Combining predictive and generative AI enables actionable insights by analyzing data and producing explanatory or personalized outputs.

Context – Traditional alerts versus predictive AI alerts

Traditional alerts and predictive AI alerts are complementary…

Traditional alerts

Static thresholds
Domain rule based

Any metrics

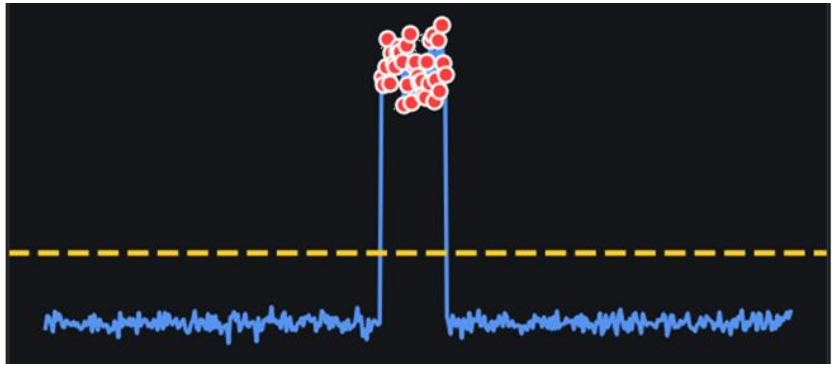

Predictive AI alerts

Dynamic baselines
Complex problems

Curated metrics
(CPU Time, Network Segments, CICS Response Time, …)

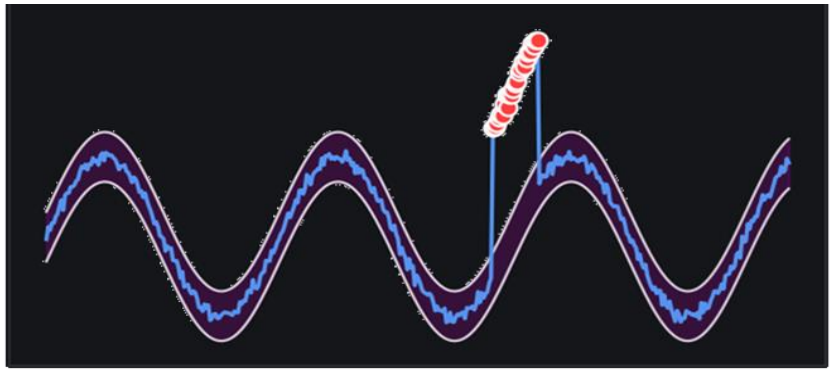

Context - Traditional alerts Versus Predictive AI alerts

Context - Traditional alerts Versus Predictive AI alerts

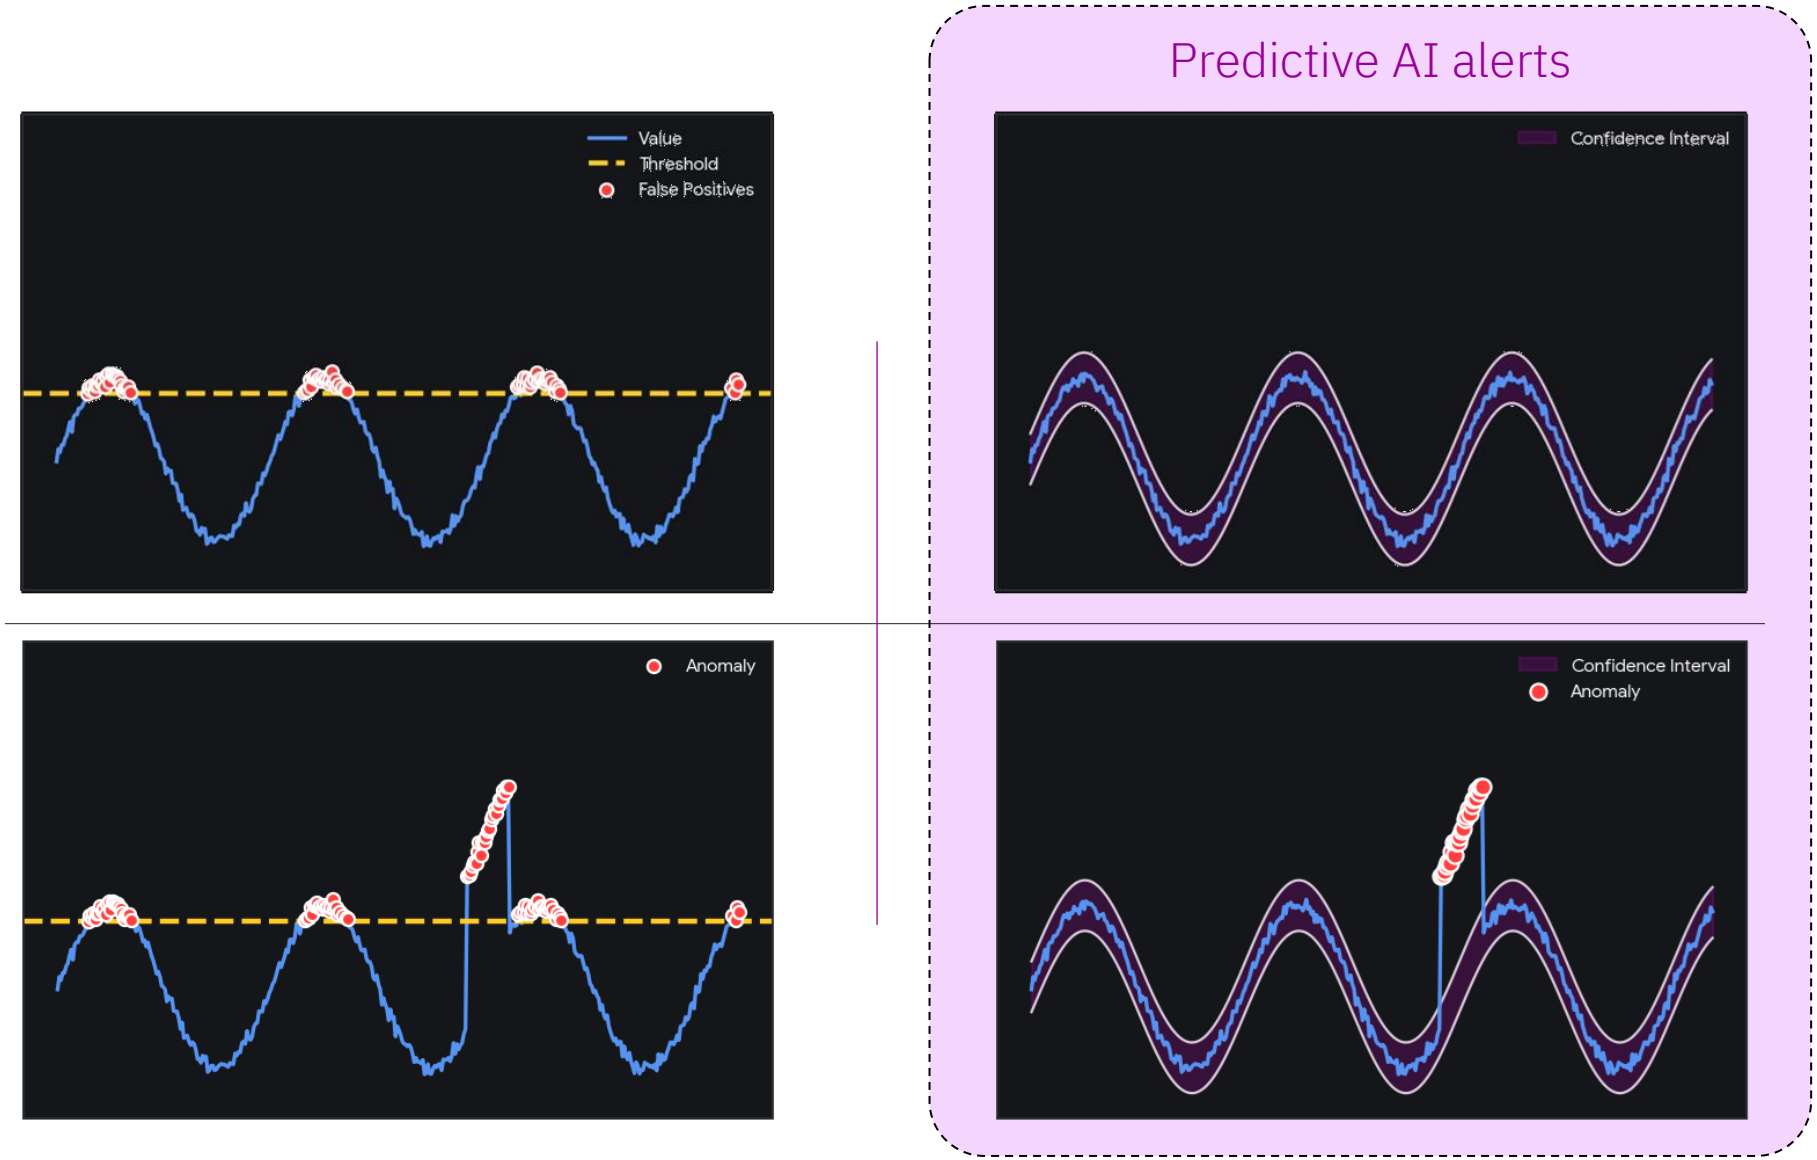

Context – high level architecture

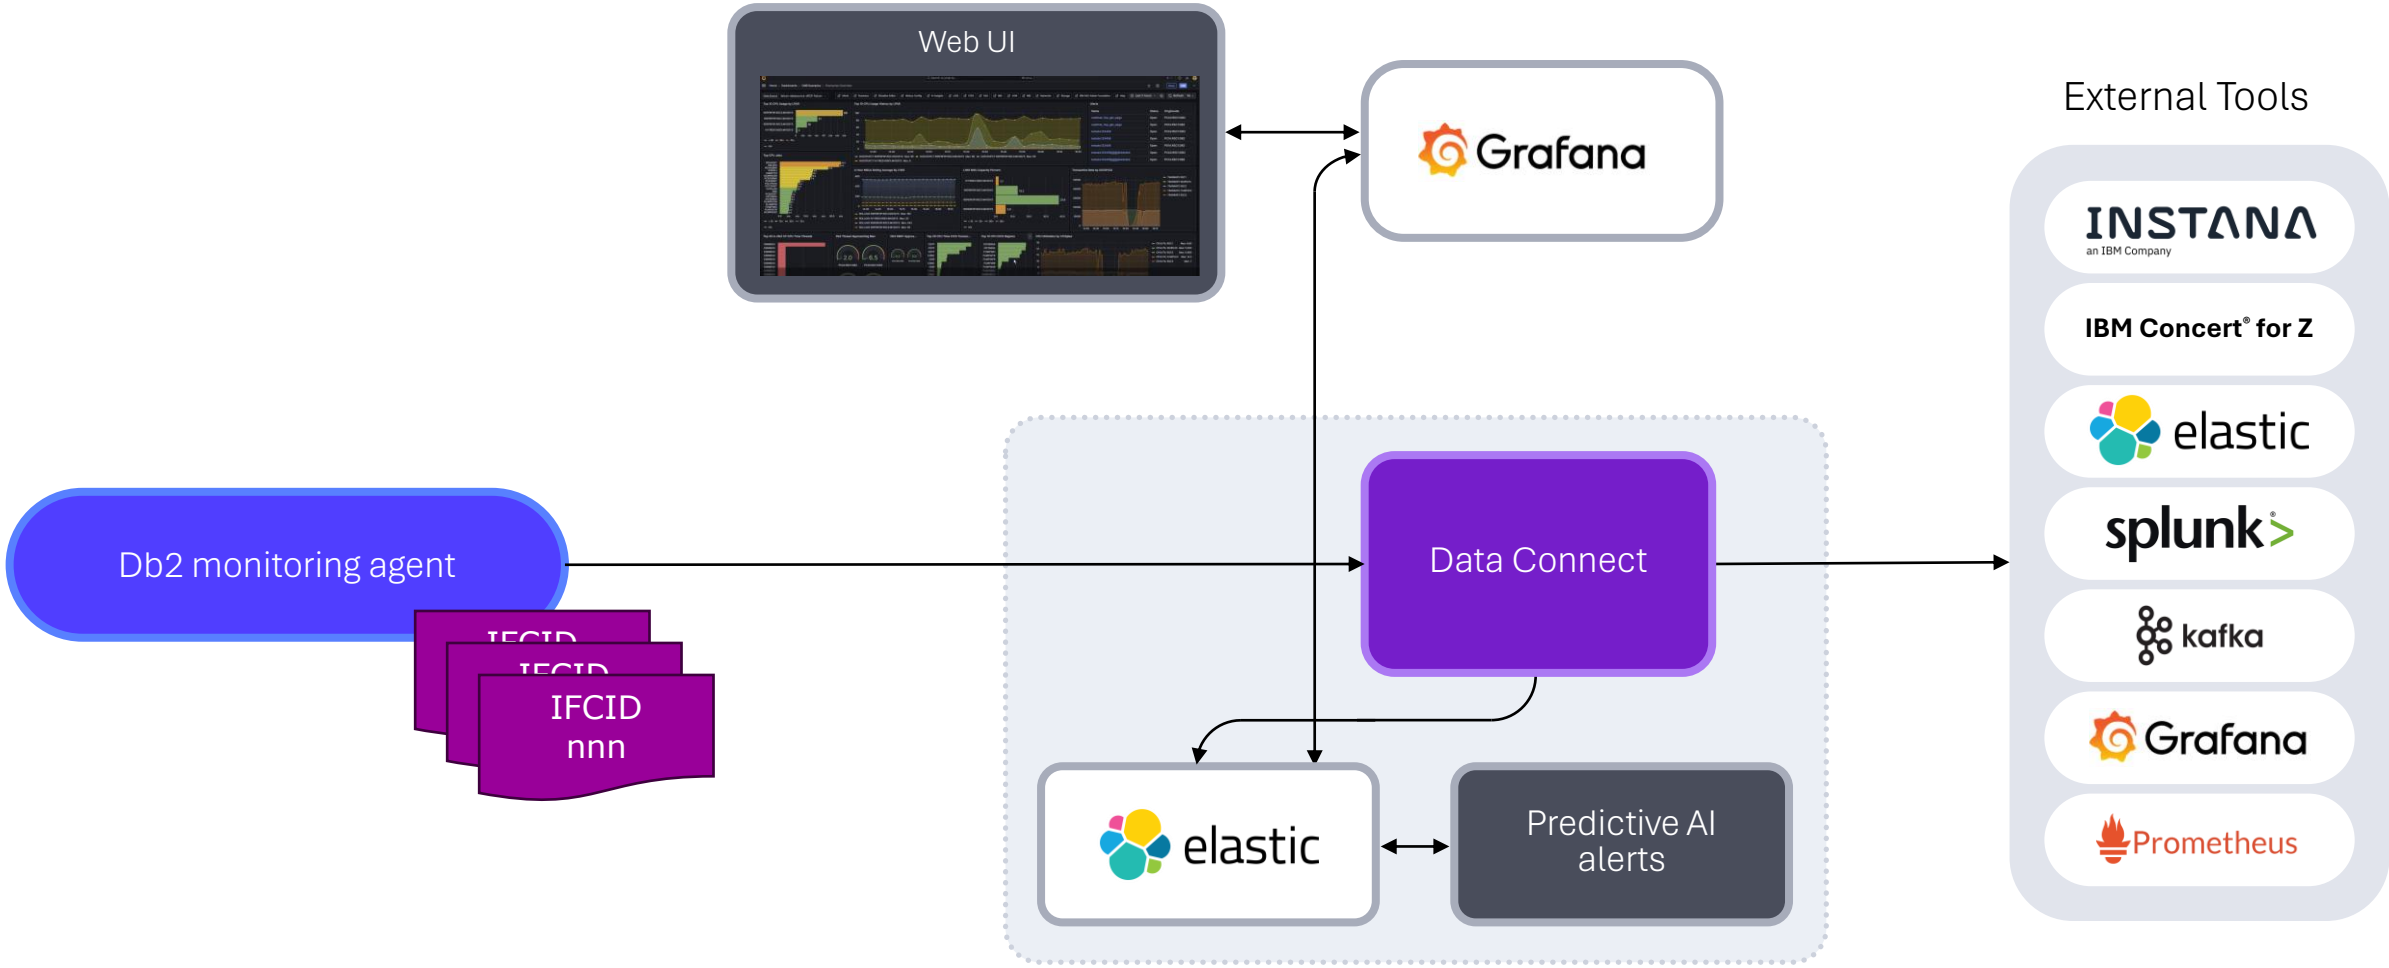

Db2 Use Case #1 – Connection Types Analysis at Db2 Subsystem Level

Alerting on significant divergence of CPU consumption per Connection Type to detect potential resource constraints or unexpected workload shifts ensuring performance and protection of transactional workload.

Context: Db2 workloads are categorized by Connection Type (CICS, IMS, REST, Utility, Batch, RRSAF, DDF) and the aggregated accounting information like Elapsed, CPU or suspension times is exposed in SMF 100 records (IFCID 369) at the Db2 Subsystem level.

👉*See it in action!*

Db2 Use Case #1 – Connection Types Analysis at Db2 Subsystem Level

Data: **3+ months of SMF 100 data that includes IFCID 369 for all production Db2 subsystems (all members of a Data Sharing Group)**

Primary KPI: In-Db2 CPU time

Secondary KPIs: Application Elapsed time, Application CPU time, In-Db2 Elapsed Time, In-Db2 Wait Time, Not Accounted times (Elapsed / CPU)

*Terminology:

- Application times = “Class 1” times
- In-Db2 times = “Class 2” times
- Wait/Suspension times = “Class 3” times

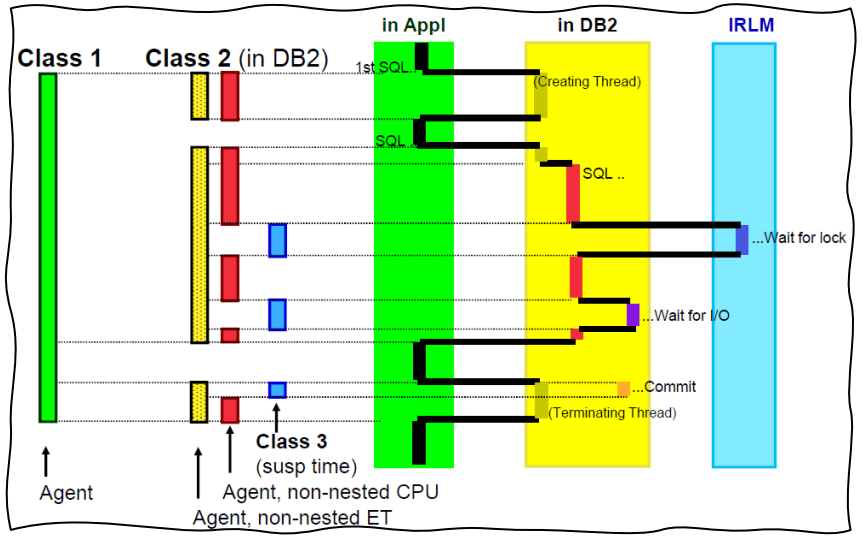

Db2 Use Case #1 – Connection Types Analysis at Db2 Subsystem Level

IF369 Collection - STATS Trace class:

- Db2 12: Restriction lifted in https://www.ibm.com/support/pages/apar/PH43916
- Db2 13: Always started with STATS CLASS(1)

Output – Connection Types Analysis at Db2 Subsystem Level

Proactive detection of poor performance, cost saving, workload optimizations

Client impact:

*"Predictive AI alerts" discovers anomalies in Db2 on the fly –giving us the possibility to **solve problems before they impact the end users experience.**"*

- Customer

Without AI Insight

Application Change
Wed. Aug. 28

Real life notification
Friday Aug. 30
Event Start

Real life resolution
Tuesday Sept. 04
Event Fixed

With AI Insight

Automated AI detection
Wed. Aug. 28
Event Detected

3 days of proactive window
to fix the issue before end of month business peak

46k CPU seconds and application outage avoided!

Output – Connection Types Analysis at Db2 Subsystem Level

Examples: Outliers on CICS workload

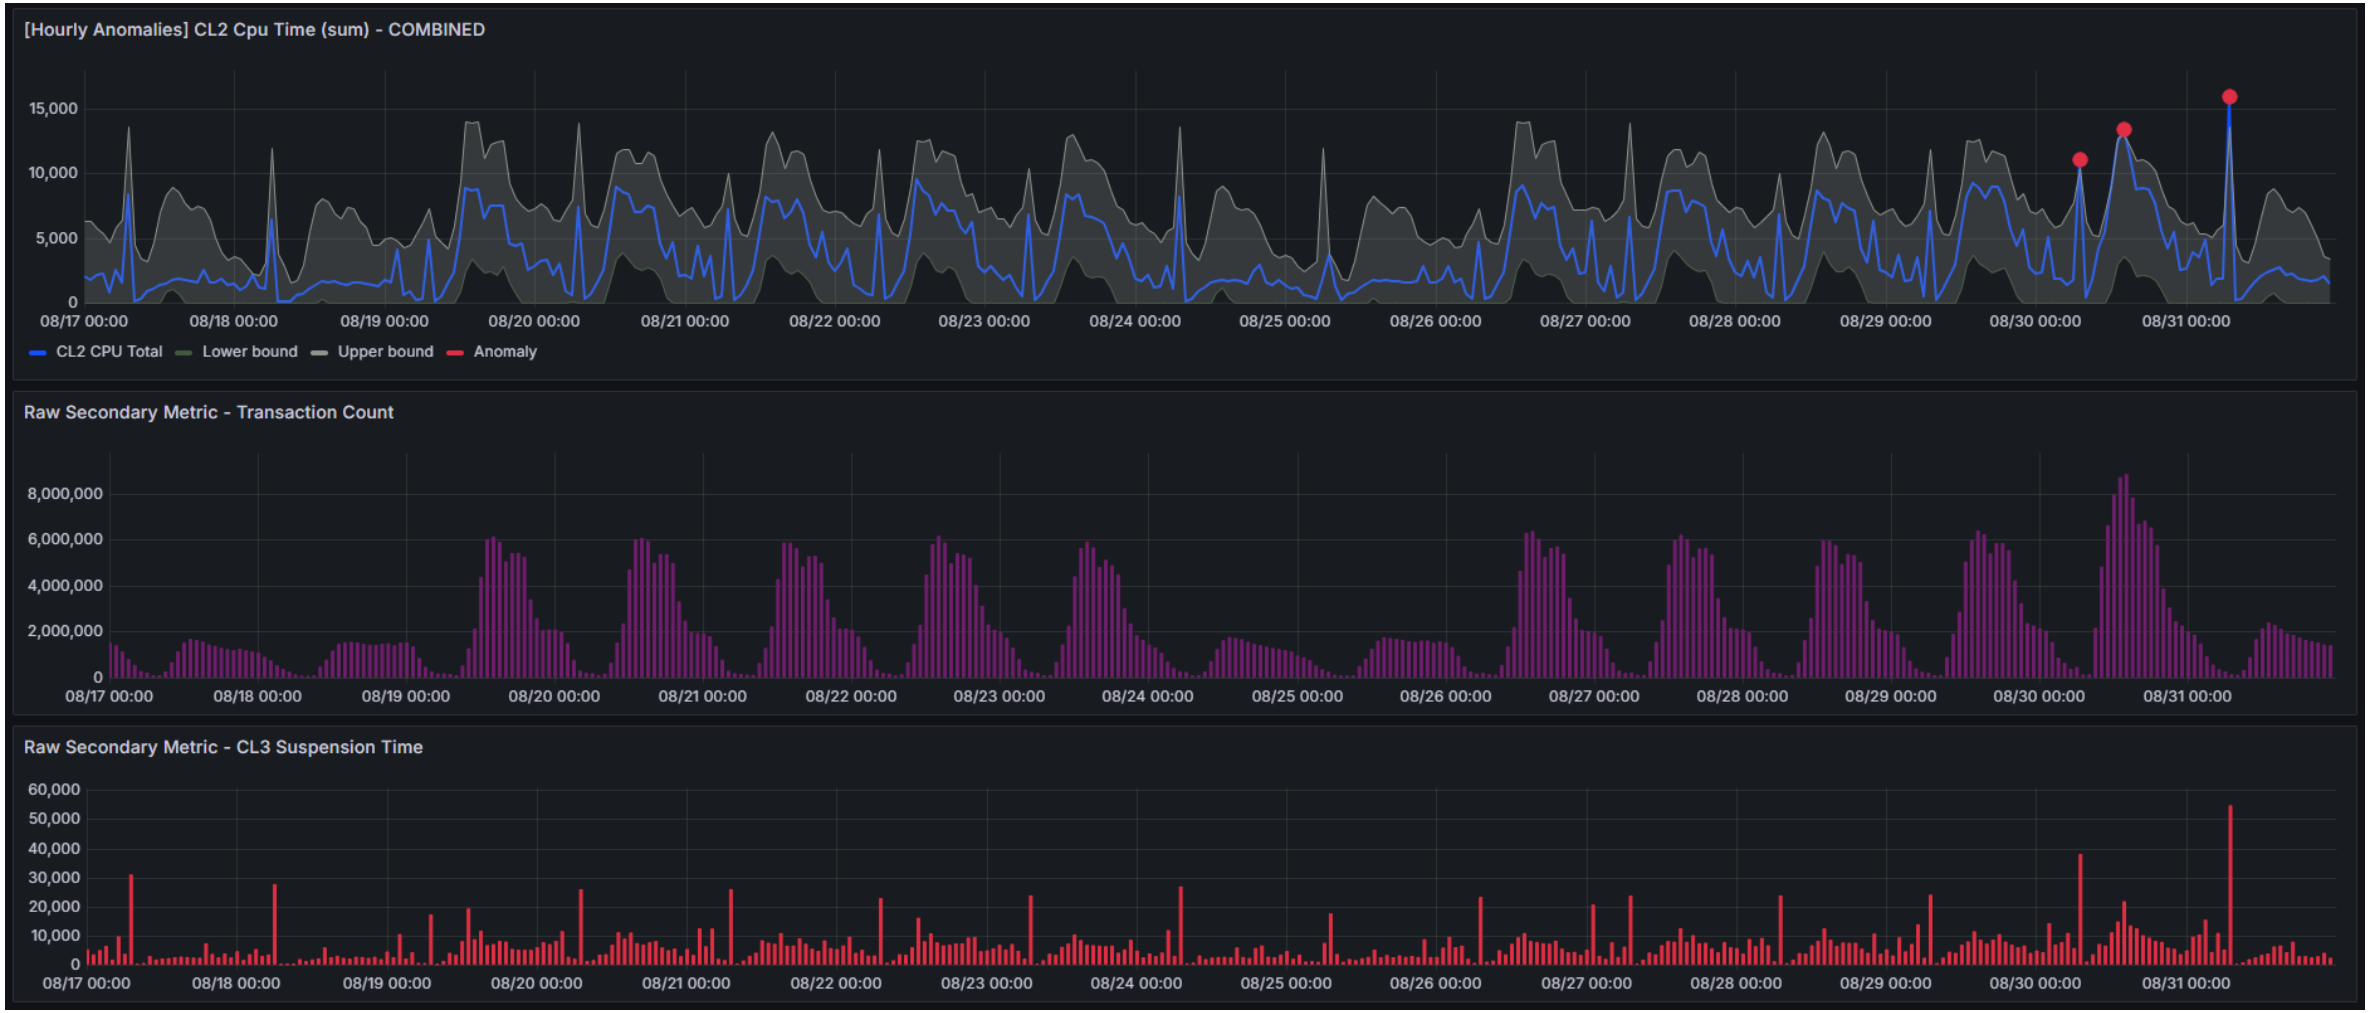

Output – Connection Types Analysis at Db2 Subsystem Level

Examples: Broken pattern on RRSAF workload

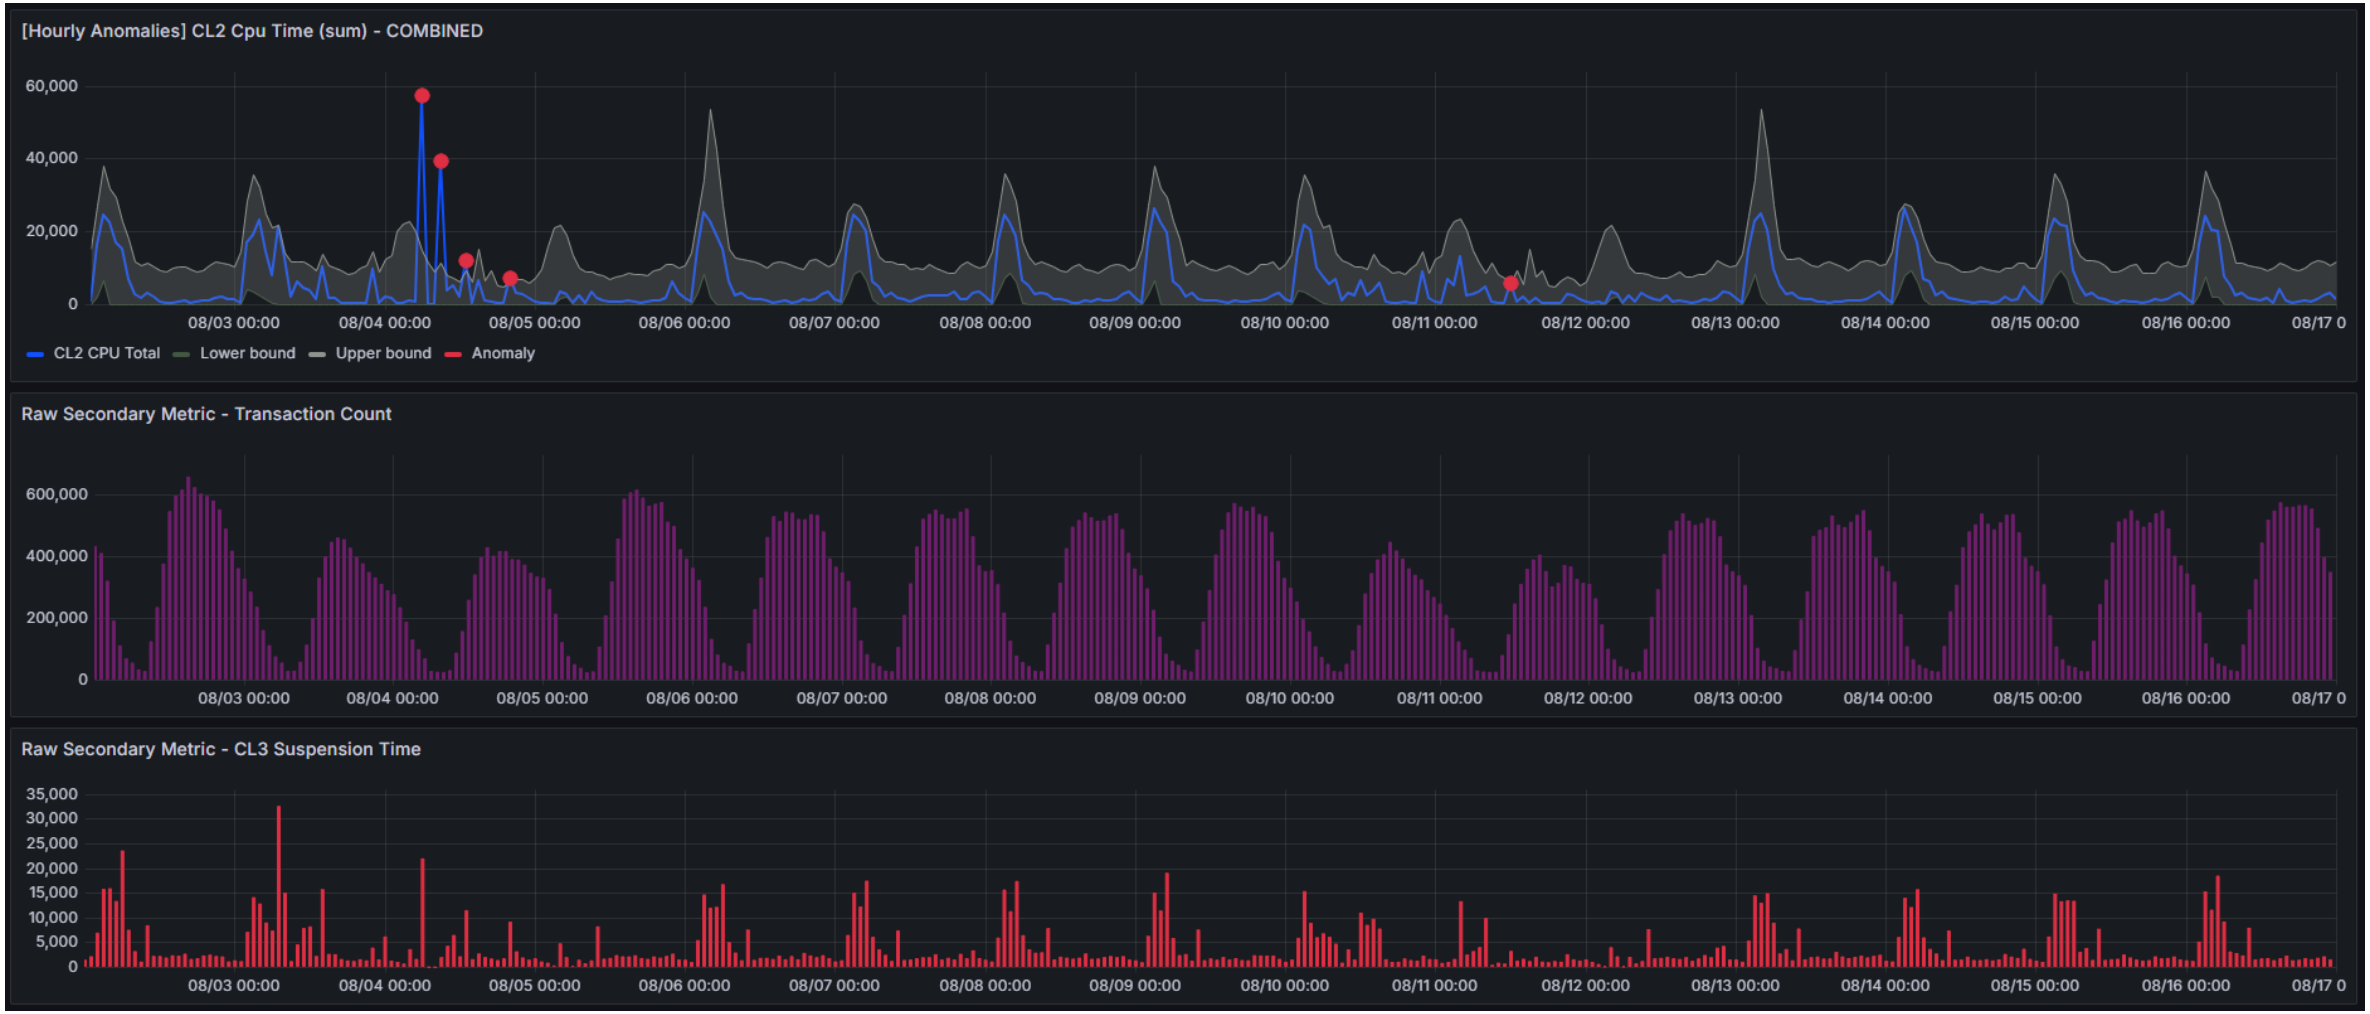

Predictive AI alerts - Explainability

Not every spike is an anomaly!

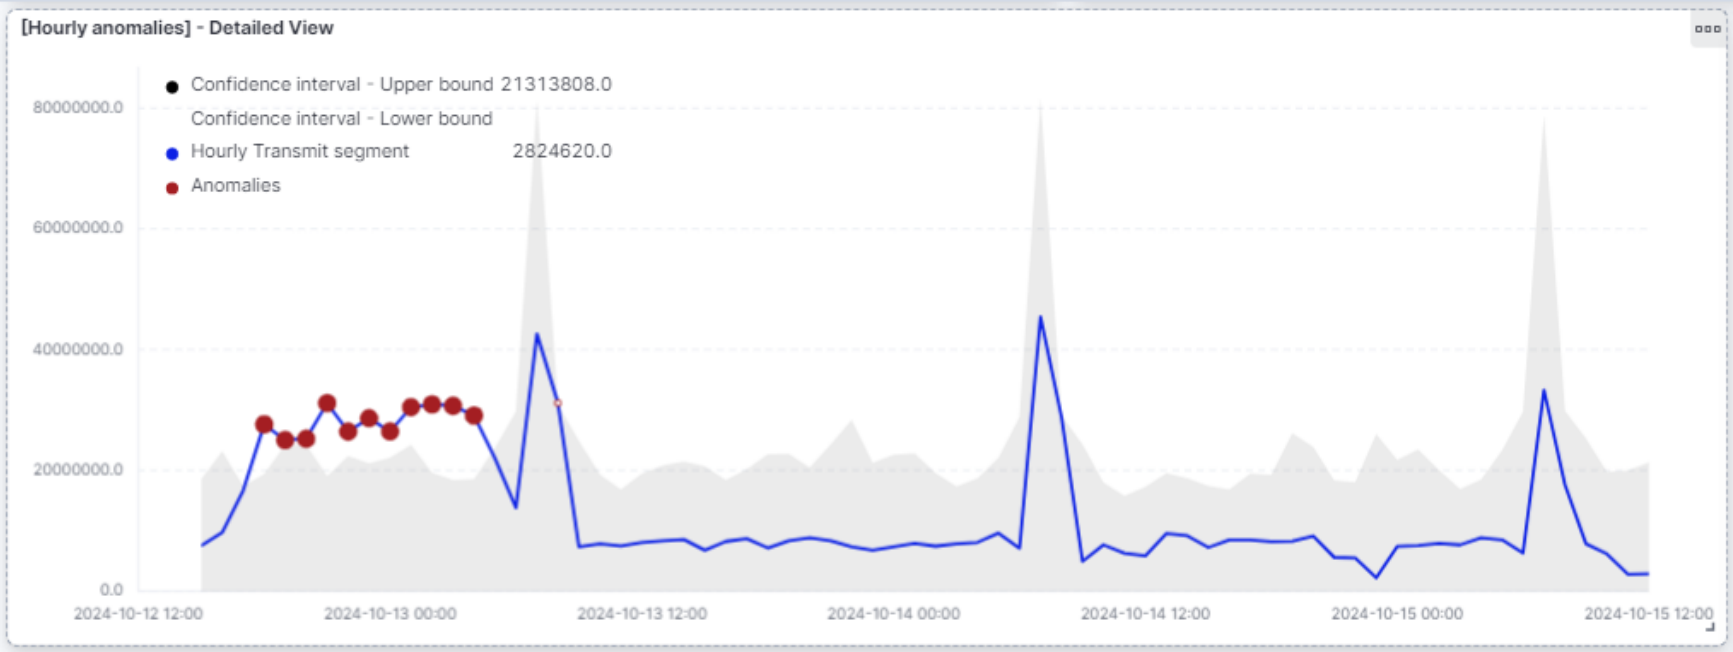

Captured timeseries for the metric subject to Machine Learning

Confidence bounds computed by the AI Models based on history

Anomaly, out of bounds behavior compared to normal

Predictive AI Models learn ranges, patterns, seasonality from a single metric and for thousands of Service Classes, hundreds of CICS Regions, Lpars, ...

Creating a dynamic threshold for each of them!

Predictive AI alerts - Explainability

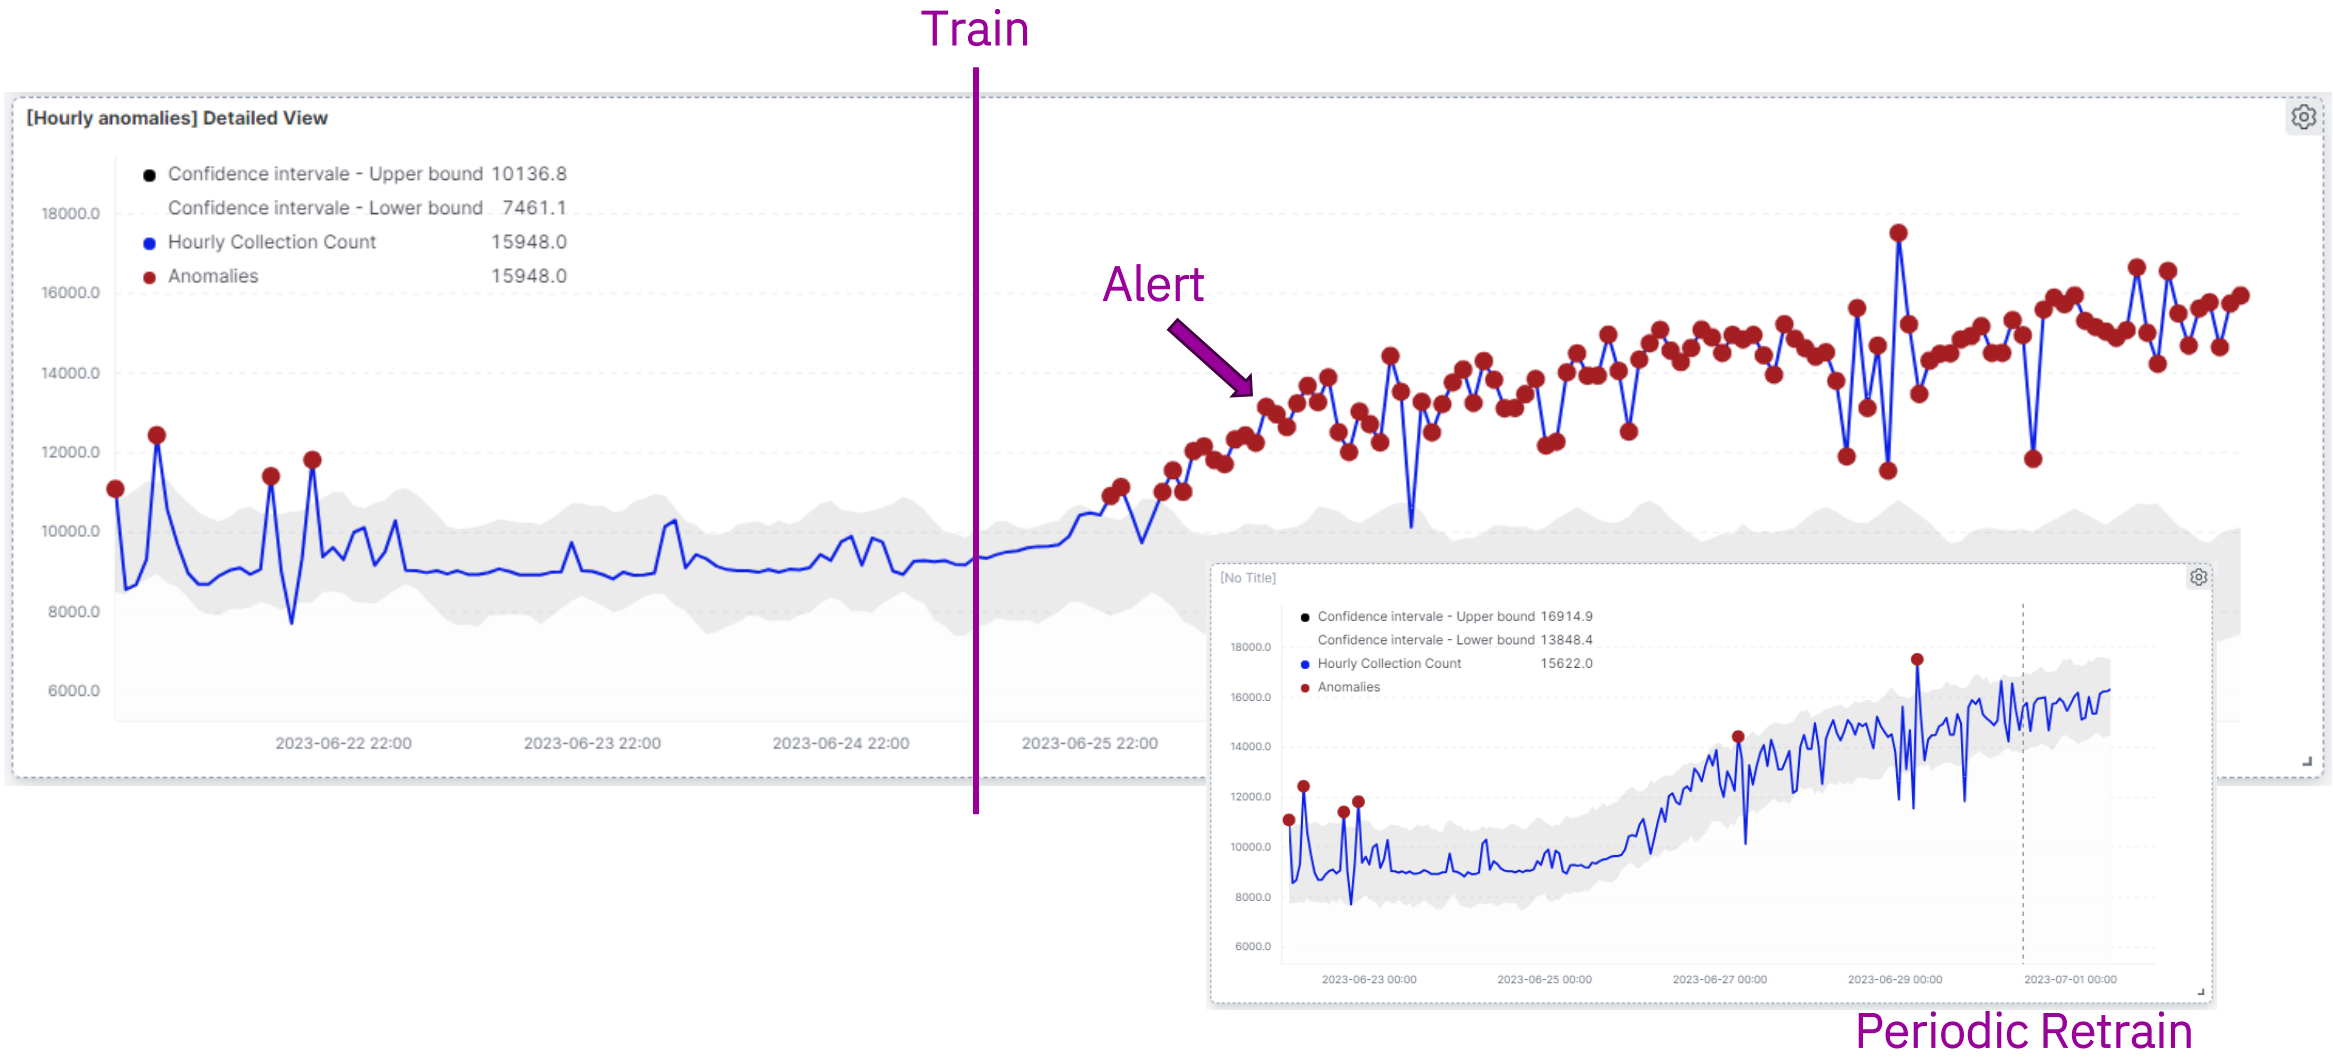

Thank you.