




# Using AI/ML to Enhance Db2 Subsystem and Application Performance

**Antonio Couto**

*Broadcom*


# Agenda

- **Machine Learning Concepts**
- Applied Machine Learning to **enhance Db2 Performance**
- **Customer Use Cases**
  - What makes a Business Application?
  - Db2 Application Elapsed Time on CICS
  - High CPU usage on Db2 DIST Address Space
  - Db2 Package and Plan performance investigation



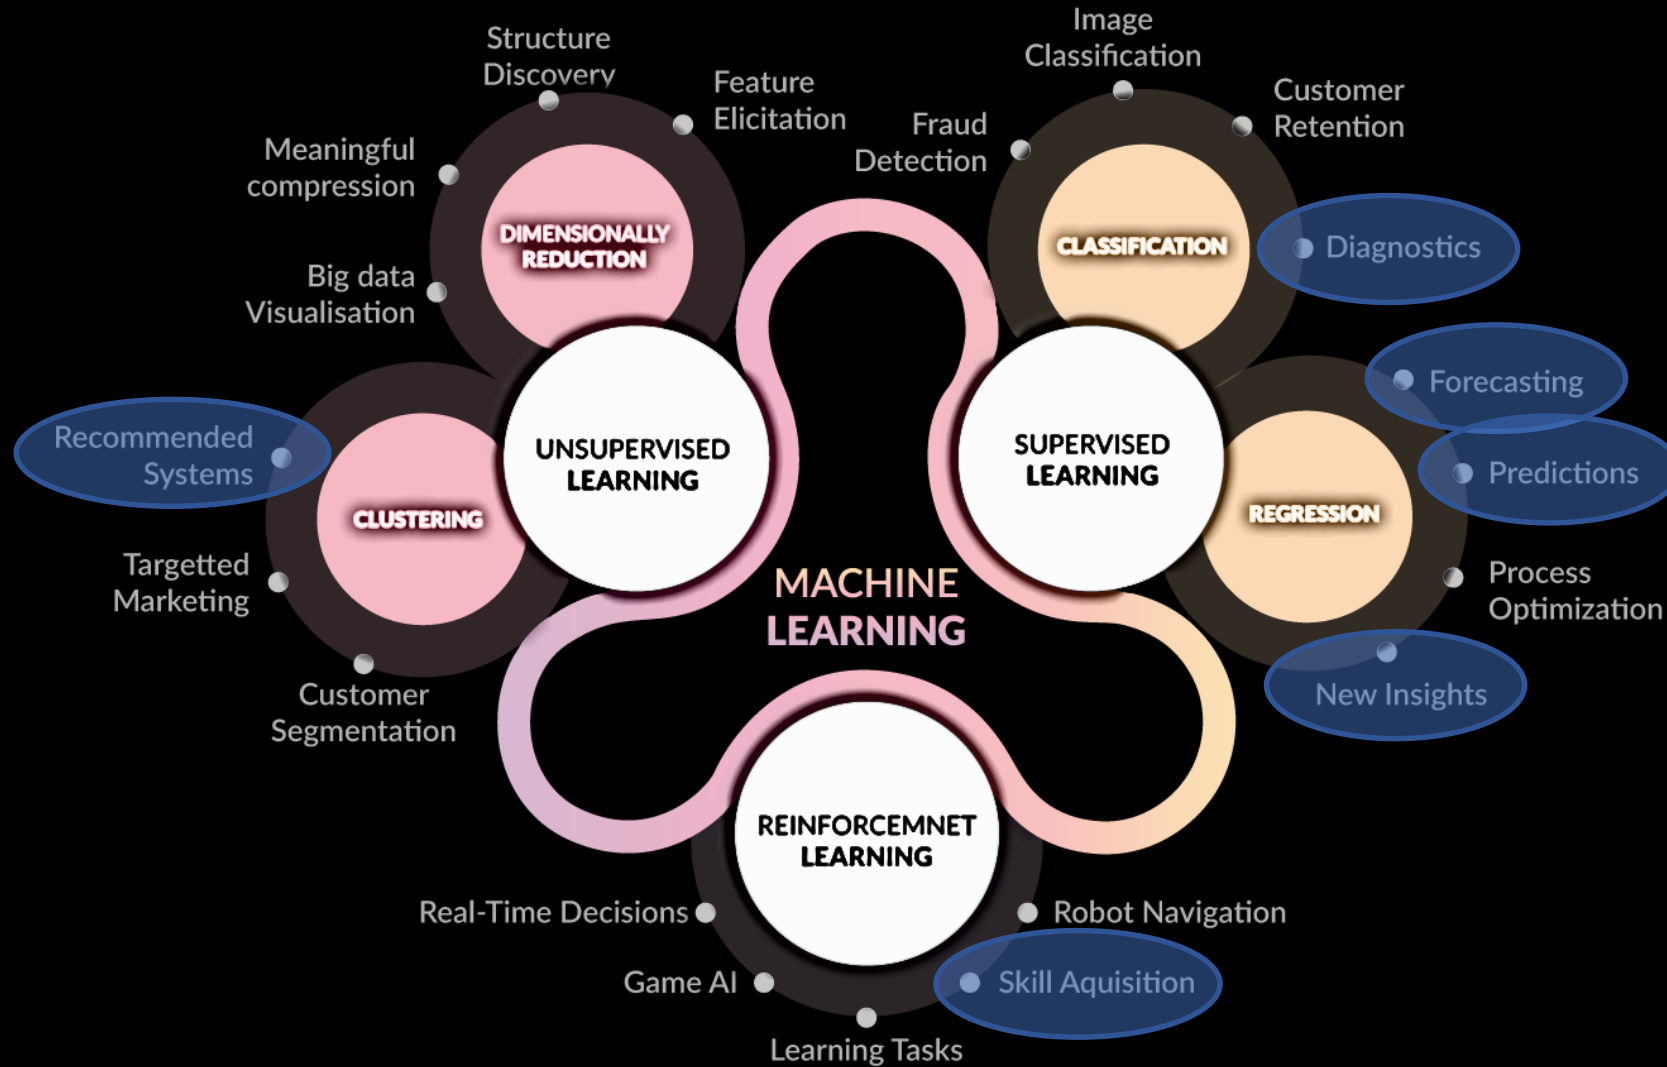
# MACHINE LEARNING CONCEPTS

# Machine Learning



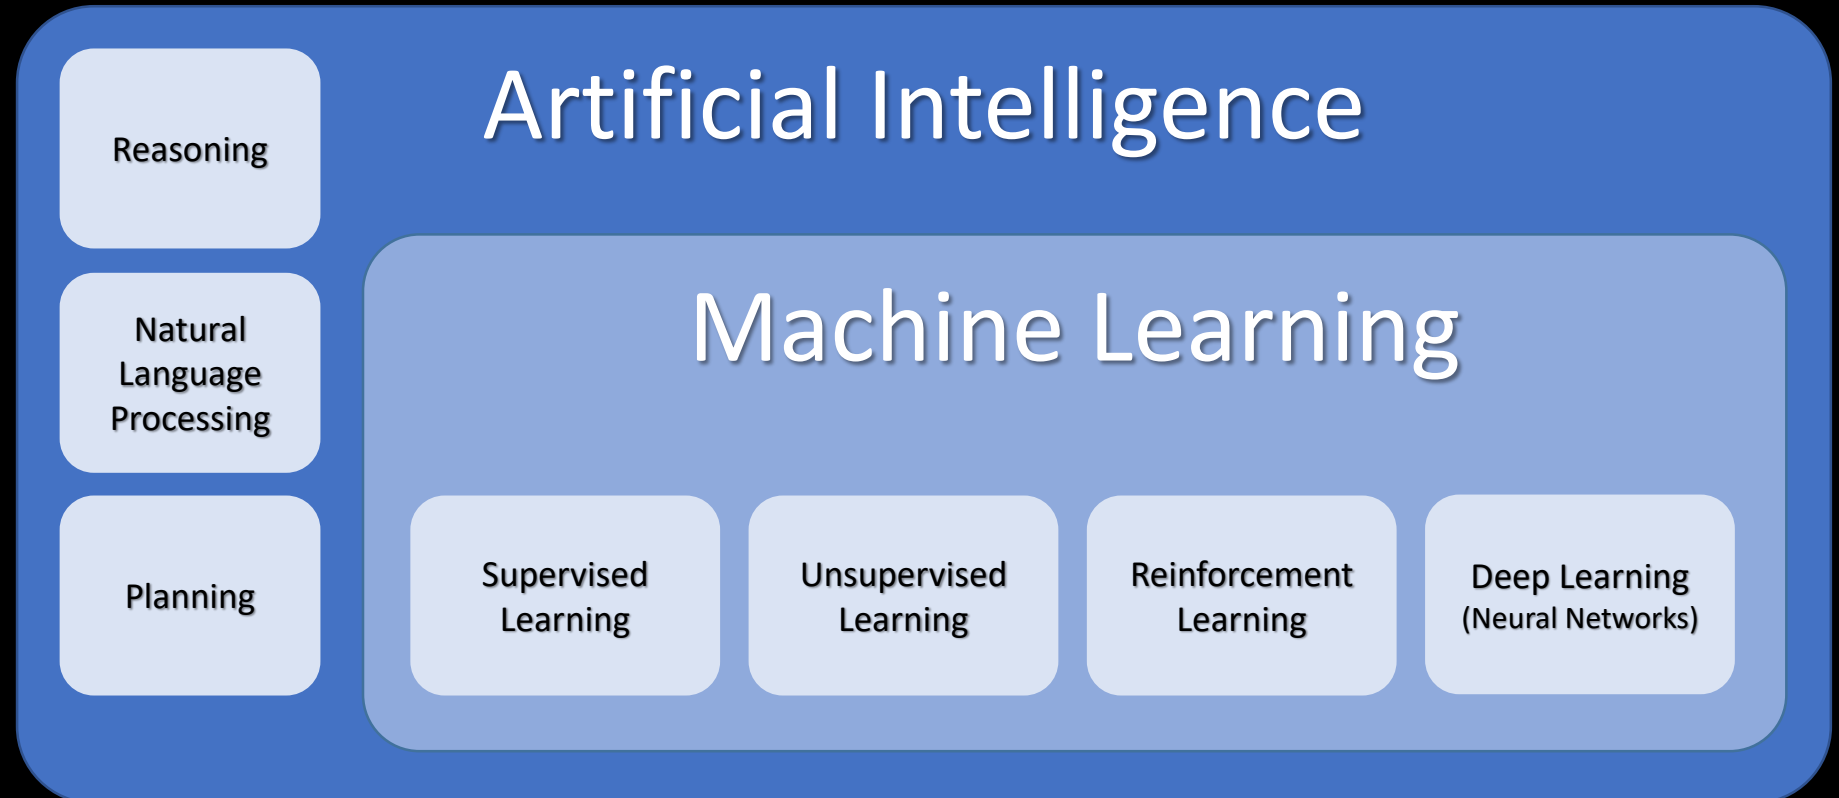


# Machine Learning



(1) <https://towardsdatascience.com/coding-deep-learning-for-beginners-types-of-machine-learning-b9e651e1ed9d>


# Machine Learning and Artificial Intelligence




(1) <https://www.ibm.com/downloads/cas/GB8ZMQZ3>

# Supervised Learning

## Linear Regression




## Logistic Regression



# Unsupervised Learning

## K-means Clustering




(1) <https://medium.com/@imparth/k-means-clustering-algorithm-34807a7cec71>




# Reinforcement Learning

## Generic Reinforcement Learning Model



# Deep Learning


## Artificial Neural Network



(1) <https://vishalyadav831874.medium.com/industry-use-cases-of-neural-network-cccc4b333621>

# Time Series Analysis

## Time Series Data




For example: Db2 performance metrics, such as QBSTGET - Total number of GETPAGE requests at every 5 seconds, for a particular application, throughout the week.

(1) <https://towardsdatascience.com/the-complete-guide-to-time-series-analysis-and-forecasting-70d476bfe775>

# Anomaly Detection

## Detecting an Anomaly

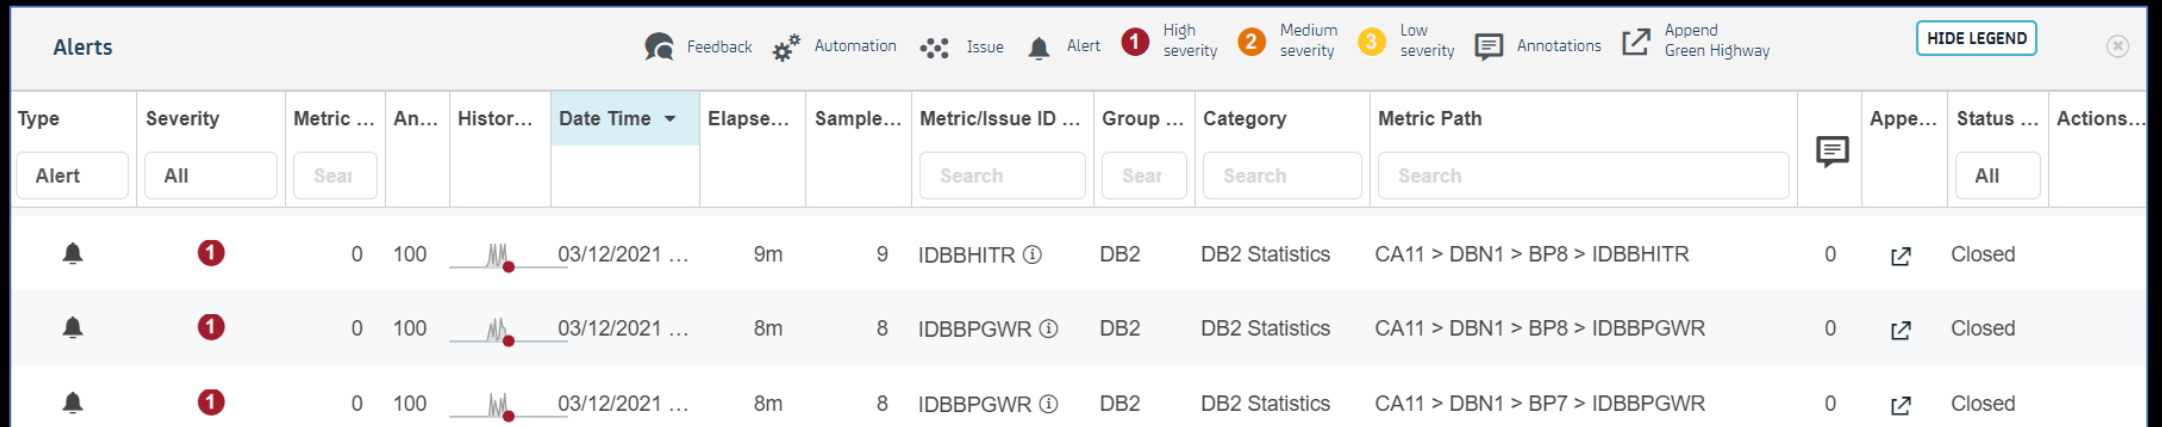


For example: Db2 performance metrics, such as **QBSTGET - Total number of GETPAGE requests at every 5 seconds**, for a particular application, throughout the week.

(1) <https://towardsdatascience.com/real-time-time-series-anomaly-detection-981cf1e1ca13>

# Pattern Discovery

## Discovering Patterns



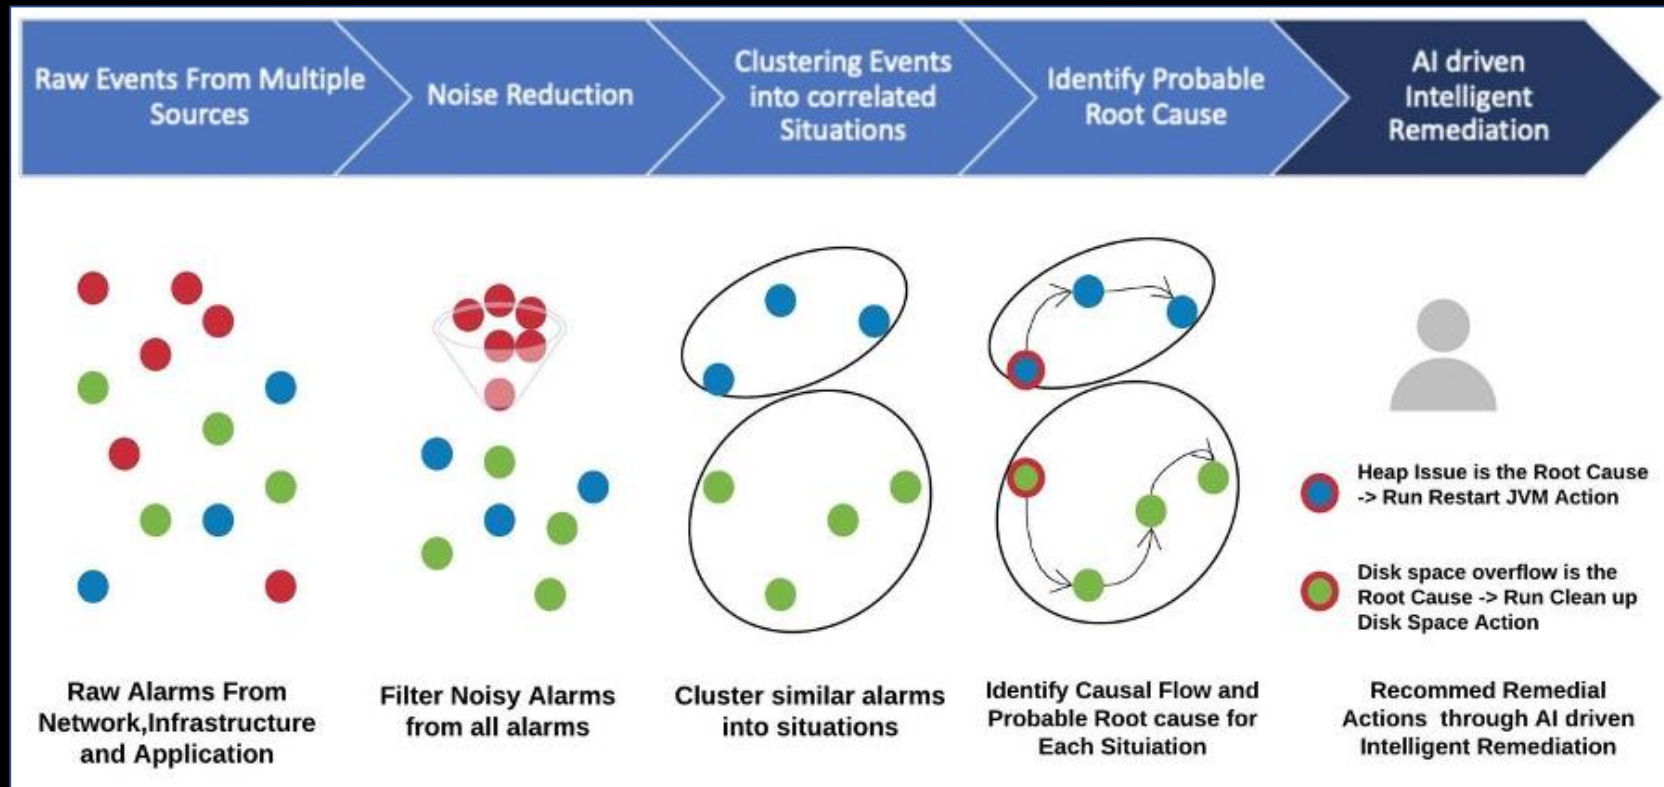
Type	Severity	Metric ...	An...	Histor...	Date Time	Elapse...	Sample...	Metric/Issue ID ...	Group ...	Category	Metric Path	Appe...	Status ...	Actions...
Alert	All	Sear						Search	Sear	Search	Search		All	
	1	0	100		03/12/2021 ...	9m	9	IDBBHITR ⓘ	DB2	DB2 Statistics	CA11 > DBN1 > BP8 > IDBBHITR	0		Closed
	1	0	100		03/12/2021 ...	8m	8	IDBBPGWR ⓘ	DB2	DB2 Statistics	CA11 > DBN1 > BP8 > IDBBPGWR	0		Closed
	1	0	100		03/12/2021 ...	8m	8	IDBBPGWR ⓘ	DB2	DB2 Statistics	CA11 > DBN1 > BP7 > IDBBPGWR	0		Closed

Pattern recognition is the process which can detect different categories and get information about particular data, predicting a label of an observation, such as: **high**, **medium** and **low**.

(1) <https://medium.com/edureka/pattern-recognition-5e2d30ab68b9>

# Root Cause Analysis

## Steps to Root Cause Analysis




(1) <https://towardsdatascience.com/how-to-conduct-a-proper-root-cause-analysis-789b9847f84b>

(2) <https://www.broadcom.cn/aiops-blog/reduce-toil-with-AI-driven-intelligent-remediation>



# Dynamic and Adaptive Alerting (1 | 2)

## *ssid*DIST DB2 CPU Utilization



**Analytics-Based Alerts** detect signal from noise

- Subtler than “**human-observed**”
- Subtler than **static thresholds** which may be the “**last defense**”

# Dynamic and Adaptive Alerting (2 | 2)

## Green Highway


Outside the Bands =  
Anomaly / Possible Alert

Metric Values

Rare Band

Probable Band

Common Band





# APPLIED MACHINE LEARNING TO ENHANCE DB2 PERFORMANCE

# Google Colab Environment

The screenshot displays the Google Colab interface. At the top, the notebook is titled "TRIDEX2024.ipynb". The menu bar includes "Arquivo", "Editar", "Ver", "Inserir", "Ambiente de execução", "Ferramentas", "Ajuda", and "Todas as alterações foram salvas". On the right, there are options for "Comentário", "Compartilhar", and a user profile icon. Below the menu, the RAM and Disco usage is shown as 0% used, and the Gemini AI model is selected. The left sidebar shows the file explorer with a tree view of the "drive" folder, including "MyDrive", "Classroom", "Colab Notebooks", and "Data". The "Data" folder contains "IDUG\_Charlotte...", "PUCRS-Python4...", "PUCRS-TCC", "TRIDEX\_2024" (with sub-file "DB2-IDBDCPU..."), "zHackathon 20...", "zOEC\_PyCalc", "IDUG2024.ipynb", "LabAula03.ipynb", and "LabAula04.ipynb". The main area shows the code editor with the following Python code:

```
1 from statsmodels.tsa.statespace.sarimax import SARIMAX
2 from statsmodels.graphics.tsaplots import plot_acf
3 from statsmodels.tsa.stattools import adfuller
4 import matplotlib.pyplot as plt
5 import matplotlib.dates as md
6 import numpy as np
7 import pandas as pd
8
9 import warnings
10 warnings.filterwarnings('ignore')
11 # %matplotlib inline
```

Below the code, there are two code cells. The first cell contains:

```
[ ] 1 plt.rcParams["figure.figsize"] = (11,5)
```

The second cell contains:

```
[ ] 1 df = pd.read_csv('/content/drive/MyDrive/Colab Notebooks/Data/TRIDEX_2024/DB2-IDBDCPUT-0504S.csv')
2 df['TIMESTAMP'] = pd.to_datetime(df['TIMESTAMP'], format="%d/%m/%Y %H:%M:%S")
3 df['DATE'] = df['TIMESTAMP'].dt.strftime('%d/%m/%Y')
4 df['TIME'] = df['TIMESTAMP'].dt.strftime('%H:%M:%S')
5 df.dtypes
```

The output of the second cell shows the data types for each column:

```
TIMESTAMP    datetime64[ns]
IDBDCPUT      float64
DATE          object
TIME          object
dtype: object
```

At the bottom, there is a code cell with:

```
[ ] 1 df.head()
```

The status bar at the bottom indicates "Conectado ao dispositivo: de back-end do Google Compute Engine em Python 3".



+ Código + Texto

```
[1] 1 from statsmodels.tsa.statespace.sarimax import SARIMAX
2 from statsmodels.graphics.tsaplots import plot_acf
3 from statsmodels.tsa.stattools import adfuller
4 import matplotlib.pyplot as plt
5 import matplotlib.dates as md
6 import numpy as np
7 import pandas as pd
8
9 import warnings
10 warnings.filterwarnings('ignore')
11 # %matplotlib inline
```

```
[2] 1 plt.rcParams["figure.figsize"] = (11,5)
```

```
1 df = pd.read_csv('/content/drive/MyDrive/Colab Notebooks/Data/TRIDEX_2024/DB2-IDBDCPUT-0504S.csv')
2 df['TIMESTAMP'] = pd.to_datetime(df['TIMESTAMP'], format="%d/%m/%Y %H:%M:%S")
3 df['DATE'] = df['TIMESTAMP'].dt.strftime('%d/%m/%Y')
4 df['TIME'] = df['TIMESTAMP'].dt.strftime('%H:%M:%S')
5 df.dtypes
```

```
TIMESTAMP    datetime64[ns]
IDBDCPUT      float64
DATE          object
TIME          object
dtype: object
```

```
[ ] 1 df.head()
```



+ Código + Texto

```
[ ] 1 df.head()
```

	TIMESTAMP	IDBDCPUT	DATE	TIME
0	2024-04-05 12:00:00	14.801119	05/04/2024	12:00:00
1	2024-04-05 12:01:00	13.197293	05/04/2024	12:01:00
2	2024-04-05 12:02:00	11.946425	05/04/2024	12:02:00
3	2024-04-05 12:03:00	19.998539	05/04/2024	12:03:00
4	2024-04-05 12:04:00	23.805432	05/04/2024	12:04:00

```
[ ] 1 df.tail()
```

	TIMESTAMP	IDBDCPUT	DATE	TIME
356	2024-04-05 17:56:00	5.115353	05/04/2024	17:56:00
357	2024-04-05 17:57:00	2.414358	05/04/2024	17:57:00
358	2024-04-05 17:58:00	5.489411	05/04/2024	17:58:00
359	2024-04-05 17:59:00	3.846366	05/04/2024	17:59:00
360	2024-04-05 18:00:00	8.396365	05/04/2024	18:00:00

```
[ ] 1 df.describe()
```





+ Código + Texto

```
[ ] 1 df.describe()
```

	TIMESTAMP	IDBDCPUT
<b>count</b>	361	361.000000
<b>mean</b>	2024-04-05 15:00:00.000000256	12.902163
<b>min</b>	2024-04-05 12:00:00	0.042191
<b>25%</b>	2024-04-05 13:30:00	5.489411
<b>50%</b>	2024-04-05 15:00:00	10.305060
<b>75%</b>	2024-04-05 16:30:00	19.980080
<b>max</b>	2024-04-05 18:00:00	53.991751
<b>std</b>	NaN	8.627902

```
[ ] 1 fig, ax = plt.subplots()
2
3 ax.plot(df['TIMESTAMP'], df['IDBDCPUT'])
4 ax.set_xlabel('Time Stamp')
5 ax.set_ylabel('Db2 CPU Total')
6 xfmt = md.DateFormatter('%Y-%m-%d %H:%M:%S')
7 ax.xaxis.set_major_formatter(xfmt)
8
9 # Convert the string arguments to datetime objects
10 start_date = np.datetime64('2024-04-05 12:00:00')
11 end_date = np.datetime64('2024-04-05 18:01:00')
```



+ Código + Texto


```
1 fig, ax = plt.subplots()
2
3 ax.plot(df['TIMESTAMP'], df['IDBDCPUT'])
4 ax.set_xlabel('Time Stamp')
5 ax.set_ylabel('Db2 CPU Total')
6 xfmt = md.DateFormatter('%Y-%m-%d %H:%M:%S')
7 ax.xaxis.set_major_formatter(xfmt)
8
9 # Convert the string arguments to datetime objects
10 start_date = np.datetime64('2024-04-05 12:00:00')
11 end_date = np.datetime64('2024-04-05 18:01:00')
12
13 # Use the datetime objects with np.arange()
14 xticks = np.arange(start_date, end_date, 1800)
15
16 # Set the xticks on the plot
17 plt.xticks(xticks)
18
19 fig.autofmt_xdate()
20 plt.tight_layout()
```





+ Código + Texto

```
19 fig.autofmt_xdate()  
20 plt.tight_layout()
```






+ Código + Texto

```
1 ADF_result = adfuller(df['IDBDCPUT'])
2
3 print(f'ADF Statistic: {ADF_result[0]}')
4 print(f'p-value: {ADF_result[1]}')
```

```
ADF Statistic: -2.2537664633853196
p-value: 0.18731280708845338
```

You should get a p-value greater than 0.05, meaning that we fail to reject the null hypothesis and conclude that the series is not stationary.

```
[7] 1 plot_acf(df['IDBDCPUT'], lags=20);
2
3 plt.ylim(-0.2, 1.2)
4 plt.tight_layout()
```






+ Código + Texto

```
[7] 1 plot_acf(df['IDBDCPUT'], lags=20);  
2  
3 plt.ylim(-0.2, 1.2)  
4 plt.tight_layout()
```




Autocorrelation





+ Código + Texto



```

[8] 1 series_diff = np.diff(df['IDBDCPUT'], n=1)

```

```

[9] 1 ADF_result = adfuller(series_diff)
    2
    3 print(f'ADF Statistic: {ADF_result[0]}')
    4 print(f'p-value: {ADF_result[1]}')

```

```


ADF Statistic: -9.70261897812648
p-value: 1.0622400124744452e-16

```

```

[10] 1 plot_acf(series_diff, lags=20);
     2
     3 plt.ylim(-0.6, 1.2)
     4 plt.tight_layout()

```








+ Código + Texto

```
3 plt.ylim(-0.6, 1.2)
4 plt.tight_layout()
```





+ Código + Texto

```
[11] 1 df_diff = pd.DataFrame({'values': series_diff})
      2
      3 train = df_diff[: -100]
      4 test = df_diff[-100:]
      5
      6 print(len(train))
      7 print(len(test))
```

260  
100


```
1 df_diff
```

	values
0	-1.603826
1	-1.250868
2	8.052114
3	3.806893
4	-1.162620
...	...
355	2.351139
356	-2.700995
357	3.075053



+ Código + Texto



```
[ ] 1 fig, (ax1, ax2) = plt.subplots(nrows=2, ncols=1, sharex=True)
2
3 ax1.plot(df['IDBDCPUT'])
4 ax1.set_xlabel('Time Steps')
5 ax1.set_ylabel('Db2 CPU Total')
6 ax1.axvspan(260, 360, color='#808080', alpha=0.2)
7 ax1.set_xlim(0, 360)
8
9 ax2.plot(df_diff['values'])
10 ax2.set_xlabel('Time Steps')
11 ax2.set_ylabel('Value')
12 ax2.axvspan(260, 360, color='#808080', alpha=0.2)
13 ax2.set_xlim(0, 360)
14
15 plt.tight_layout()
```





+ Código + Texto

```
12 ax2.axvspan(260, 360, color= #808080 , alpha=0.2)  
13 ax2.set_xlim(0, 360)  
14  
15 plt.tight_layout()
```





+ Código + Texto

0s [19] 0 50 100 150 200 250 300 350

Time Steps

```
1 def rolling_predictions(df_diff, last_train_value, train_len, horizon, window, method):
2
3     TOTAL_LEN = train_len + horizon
4
5     if method == 'MA':
6         pred_MA = []
7
8         for i in range(train_len, TOTAL_LEN, window):
9             model = SARIMAX(df_diff[:i], order=(0,0,1))
10            res = model.fit(dispatch=False)
11            predictions = res.get_prediction(0, i + window - 1)
12            # oos_pred = predictions.predicted_mean.iloc[-window:]
13            oos_pred = predictions.predicted_mean.iloc[-window:] * -1
14            pred_MA.extend(oos_pred)
15
16            last_train_value = np.array([last_train_value])
17            pred_MA = np.concatenate((last_train_value, pred_MA))
18            pred_MA = pred_MA.cumsum()
19
20            return pred_MA[:100]
```

```
10s [27] 1 pred_df = df[-100:].copy()
2
3 TRAIN_LEN = len(train)
4 HORIZON = len(test)
```



+ Código + Texto

```

18     pred_MA = pred_MA.cumsum()
19
20     return pred_MA[:100]

```

```

[27] 1 pred_df = df[-100:].copy()
      2
      3 TRAIN_LEN = len(train)
      4 HORIZON = len(test)
      5 LAST_TRAIN_VALUE = df.iloc[259].IDBDCPUT
      6
      7 windows = [1]
      8
      9 for window in windows:
10     pred_MA = rolling_predictions(df_diff, LAST_TRAIN_VALUE, TRAIN_LEN, HORIZON, window, 'MA')
11     pred_df[f'pred_MA_{window}'] = pred_MA
12
13 pred_df.head()

```


	TIMESTAMP	IDBDCPUT	DATE	TIME	pred_MA_1
261	2024-04-05 16:21:00	21.287241	05/04/2024	16:21:00	19.135260
262	2024-04-05 16:22:00	20.490201	05/04/2024	16:22:00	20.847703
263	2024-04-05 16:23:00	21.045079	05/04/2024	16:23:00	20.200555
264	2024-04-05 16:24:00	21.206927	05/04/2024	16:24:00	19.449346
265	2024-04-05 16:25:00	21.644510	05/04/2024	16:25:00	19.347264





+ Código + Texto

```
1 fig, ax = plt.subplots()
2
3 ax.plot(df['IDBDCPUT'])
4 ax.plot(pred_df['IDBDCPUT'], 'b-', label='actual')
5 ax.plot(pred_df[f'pred_MA_1'], 'k--', label='MA(1)')
6
7 ax.legend(loc=2)
8 ax.set_xlabel('Time steps')
9 ax.set_ylabel('Db2 CPU Total')
10 ax.axvspan(261, 360, color='#808080', alpha=0.2)
11 ax.set_xlim(150, 360)
12 ax.set_title(f'Moving Average')
13
14 plt.tight_layout()
```







+ Código + Texto

```
12 ax.set_title(f'Moving Average')  
13  
14 plt.tight_layout()
```



### Moving Average





## CUSTOMER USE CASES

# What makes a Business Application?

## Web and Mobile apps


- Distributed App
- Middleware
  - MQ
  - WAS
- Package DISTSERV
- Dynamic SQL Statements
- Tables/Indexes
  - Other DB2 objects

## CICS Transaction

- Transaction ID
- Program
- Plan
- Package
- SQL Statements
- Tables/Indexes
  - Other DB2 objects

## Batch Job

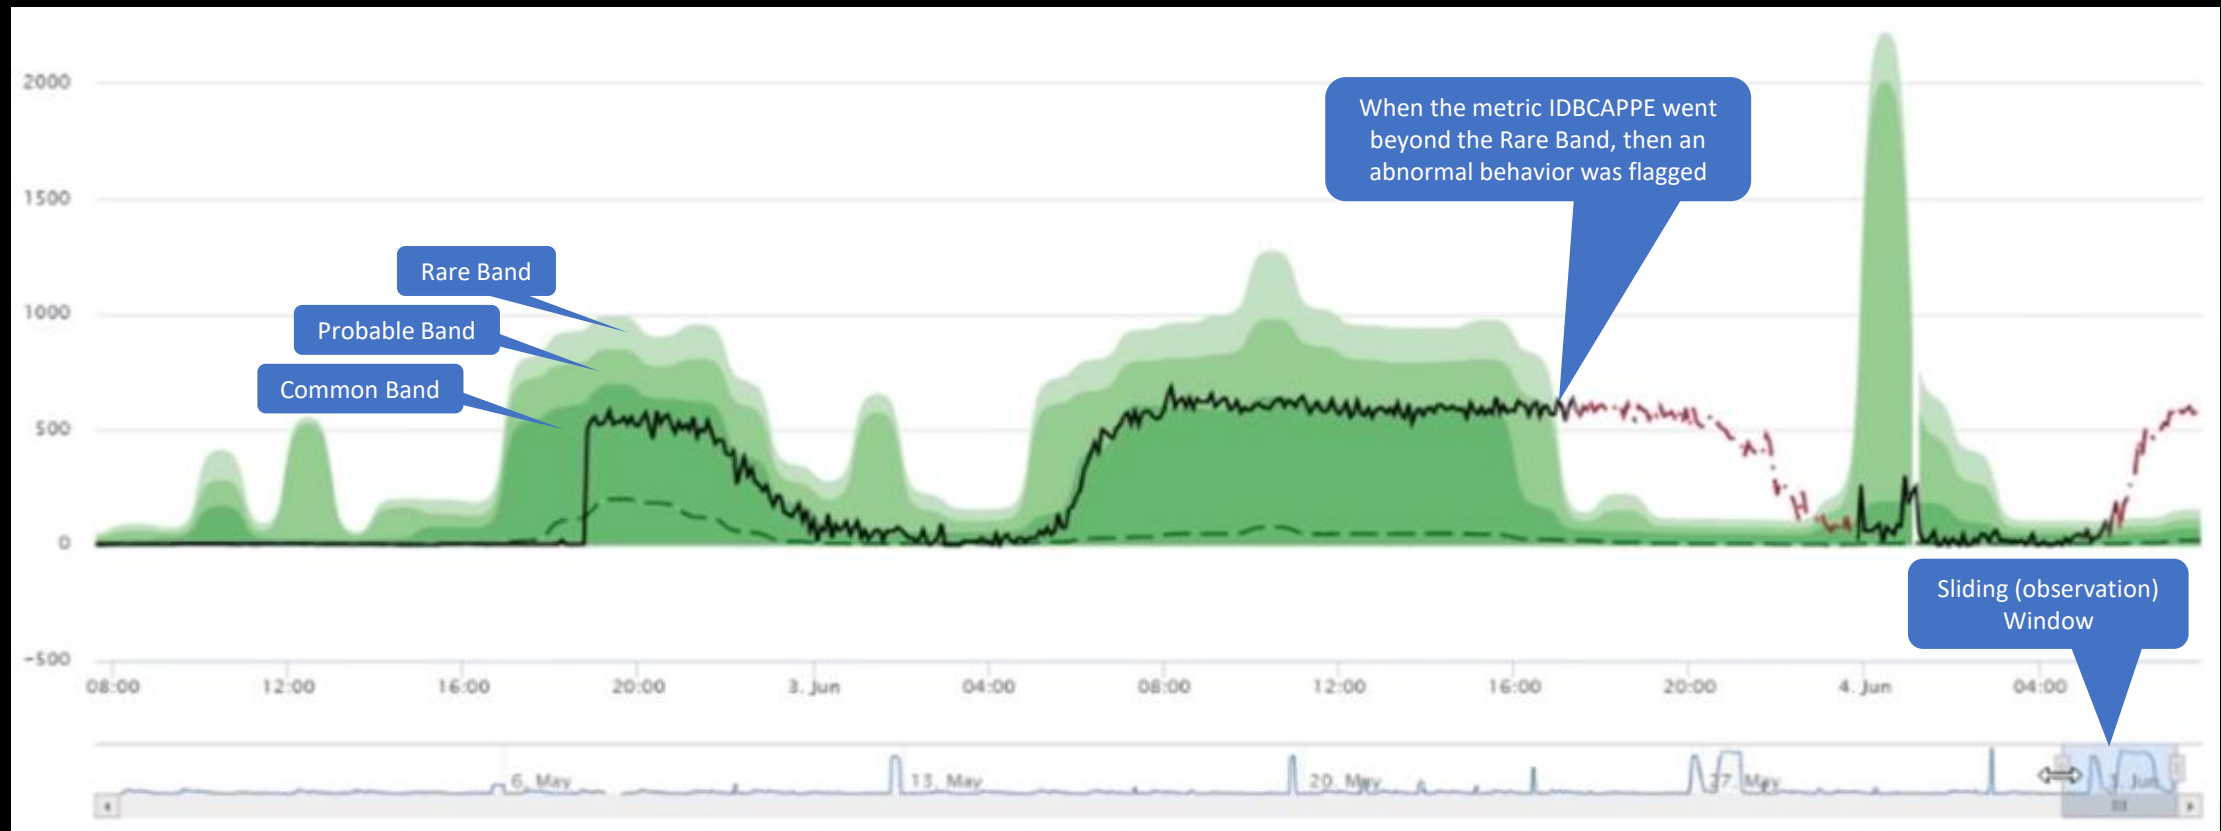
- Program
- Plan
- Package
- SQL Statements
- Tables/Indexes
  - Other DB2 objects



# APPLICATION ELAPSED TIME FOR THE CONNECTION

# Db2 Application Elapsed Time on CICS (1|2)


- Connection with CICS the metric **IDBCAPPE** (IFCID 369 – field QWACESC)






# Db2 Application Elapsed Time on CICS (2 | 2)

- Connection with CICS the metric **IDBCAPPE** (IFCID 369 – field QWACESC)








# HIGH CPU USAGE ON DB2 DIST ADDRESS SPACE

# High CPU usage on Db2 ssidDIST tasks (1 | 5)

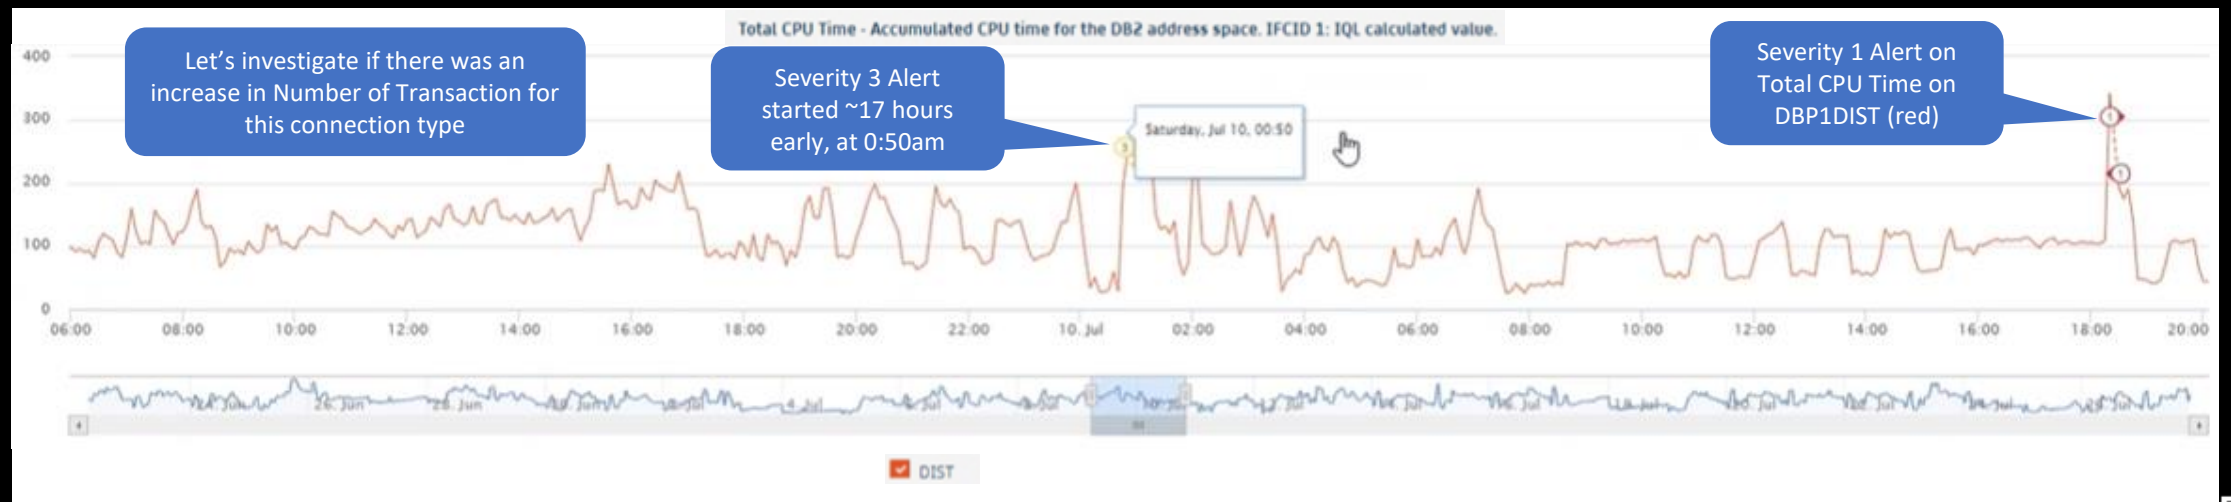
- 5 Production LPARs and 5 Production Db2 Subsystems
- Accumulated CPU time for the Db2 **DIST**ributed Address Space



# High CPU usage on Db2 ssidDIST tasks (2 | 5)

- Zooming in a particular LPAR and Db2 Subsystem
- Accumulated CPU time for the Db2 **DIST**ributed Address Space

Severity	Metric Value	Date	Time	Elapsed Time	Samples	Metric	Metric Path	Status
1	3.078766	07/10/21	18:29	6m	6	IDBDCPUT	S1P1 > <b>DBP1</b> > DIST > IDBDCPUT	Closed
3	1.513442	07/10/21	1:06	16m	16	IDBDCPUT	S1P1 > <b>DBP1</b> > DIST > IDBDCPUT	Closed



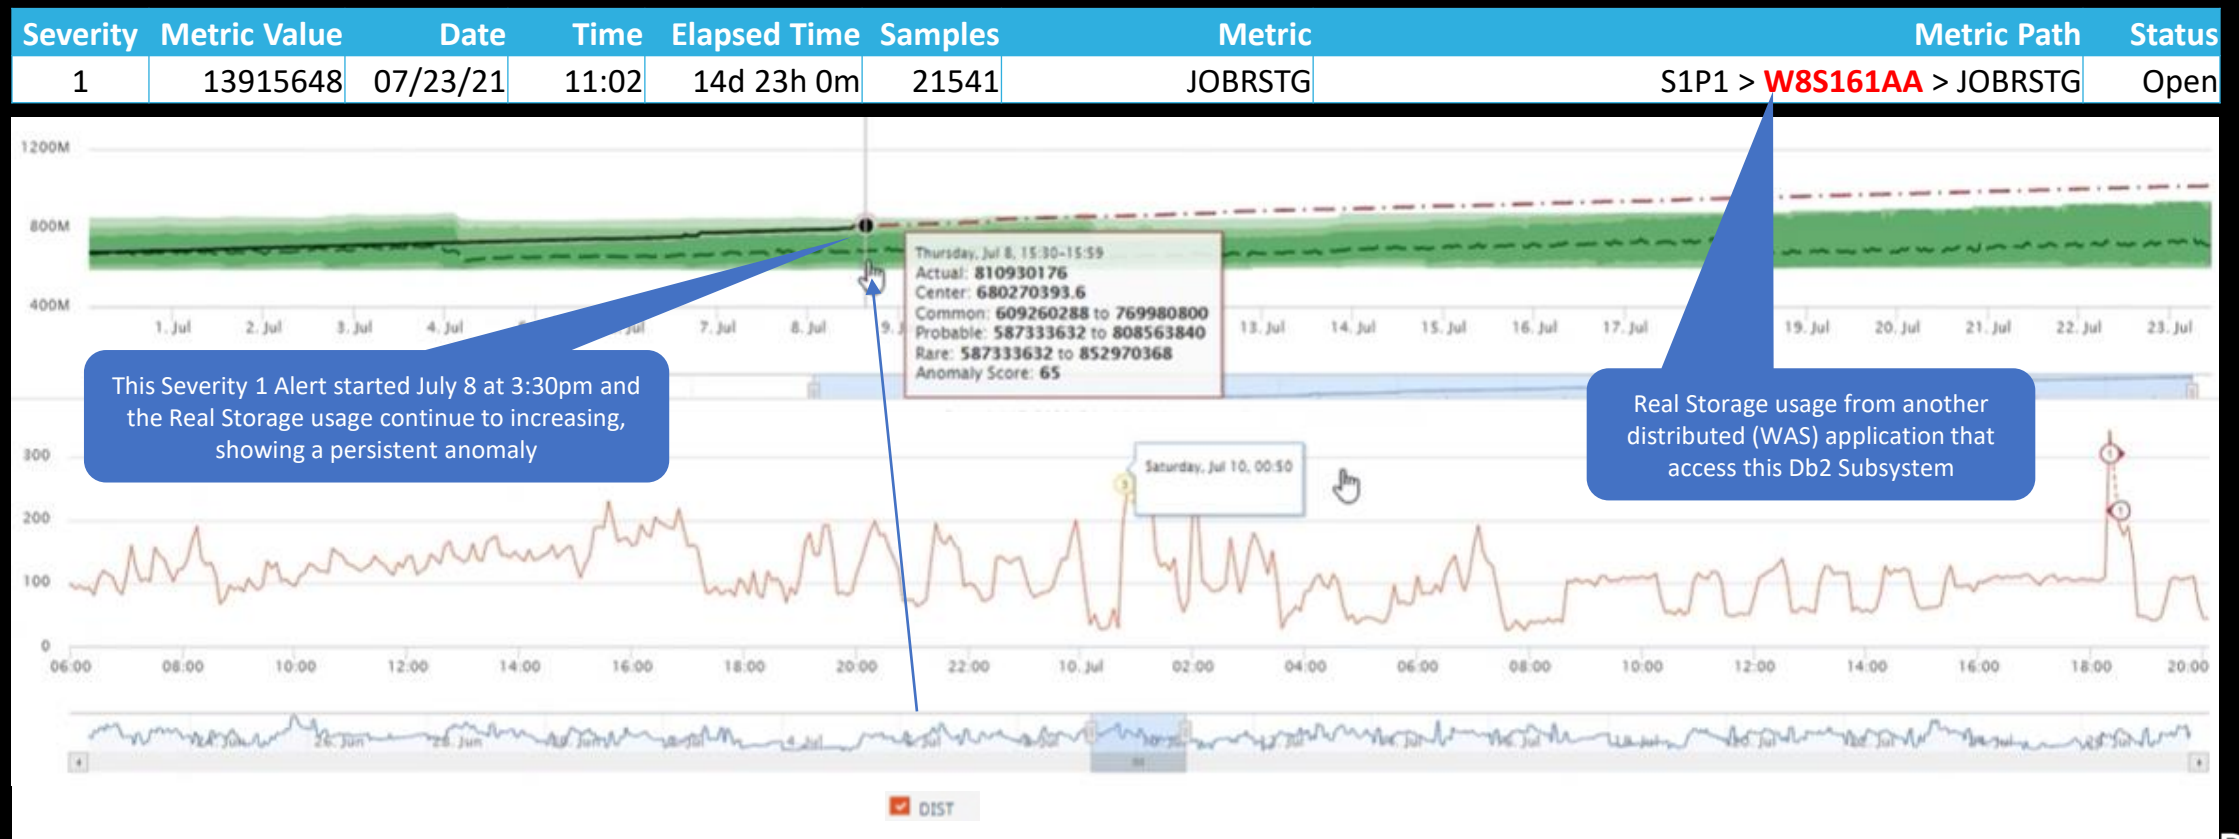
# High CPU usage on Db2 ssidDIST tasks (3 | 5)

- Investigate if there was an increase in Number of Transaction for this connection type



# High CPU usage on Db2 ssidDIST tasks (4 | 5)

- Investigate if there was another distributed application causing this Severity 1 Alert





# High CPU usage on Db2 ssidDIST tasks (5 | 5)

- Investigate if there was another distributed application causing this Severity 1 Alert

Severity	Metric Value	Date	Time	Elapsed Time	Samples	Metric	Metric Path	Status
1	13915648	07/23/21	11:02	14d 23h 0m	21541	JOBRSTG	S1P1 > <b>W8S161AA</b> > JOBRSTG	Open






# DB2 PACKAGE AND PLAN PERFORMANCE INVESTIGATION

# Db2 Package and Plan CPU usage issue

- Analyze the Total CPU Time spent executing a particular PLAN and PACKAGE

Severity	Metric Value	Date	Time	Elapsed Time	Samples	Metric	Metric Path	Status
1	0.721046	08/25/21	17:00	2h 0m	2	DB2_CPU_PKGE	MSTRSVW > D121 > <b>MARBLES</b> > DB2_CPU_PKGE	Closed
1	0.000362	08/25/21	15:00	2h 0m	2	DB2_CPU_AVG_PKGE	MSTRSVW > D121 > <b>MARBLES</b> > DB2_CPU_AVG_PKGE	Closed







# Thank you

**Antonio Couto**

[antonio.couto@broadcom.com](mailto:antonio.couto@broadcom.com)

# GitHub Repository

<https://github.com/AntonioJSCouto/TRIDEX2024Q2>

Antonio Couto

[antonio.couto@broadcom.com](mailto:antonio.couto@broadcom.com)

[https://www.linkedin.com/in/ancouto/  
@antoniojscouto](https://www.linkedin.com/in/ancouto/@antoniojscouto)