IDUG

Leading the Db2 User Community for 35 Years

IDUG

2023 EMEA Db2 Tech Conference

It’s AI Jim, but not as we know it!

Roy Boxwell

Software Engineering GmbH

*Session Code **E12***

#IDUGDb2

Platform: Db2 for z/OS

Agenda

- What is Artificial Intelligence?
- Db2 13 and Data Insights - AI for free?
- SEG and AI
- Q&A

Agenda

- What is Artificial Intelligence?
- Db2 13 and Data Insights - AI for free?
- SEG and AI
- Q&A

What is Artificial Intelligence?

It is ***the*** buzzword of the moment!

Everyone has heard of, and probably already used, ChatGPT, MidJourney, Prezo, etc. etc.

According to industry gurus, it will make a lot of people very rich and/or unemployed.

There are lots of web sites out there where you can play with it and see what it can do (or not do!)

What is Artificial Intelligence?

Sadly, it is all bogus!

There is no such thing as AI… however, wherever you look, „There it is!“

Why is this?

We, as programmers, are all acutely aware that it is a predictive program for the next word, based on a massive input of 36 Billion data points. It looks and feels like an AI but all it is doing is *guessing* at what comes next, based only on what/how it happened before.

What is Artificial Intelligence?

With pictures, it gets even better, or worse, depending on how you look at it!

I asked some of the leading graphic AIs to generate pictures for me based upon the keywords:

AI COBOL Credit card

What did I get?

What is Artificial Intelligence?

This is not too bad:

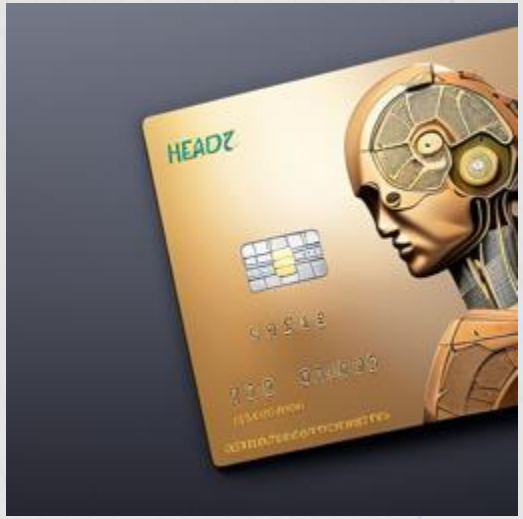

What is Artificial Intelligence?

Now it gets a bit wobbly:

What is Artificial Intelligence?

Next up was AI COBOL DB2:

What is Artificial Intelligence?

Next up was AI COBOL DB2:

What is Artificial Intelligence?

Next up was AI COBOL DB2:

Hmmm…

What is Artificial Intelligence?

Next up was AI COBOL DB2:

Hmmm…

and all in black and white!

What is Artificial Intelligence?

Then I tried „Cluster Detection“:

What is Artificial Intelligence?

Then I tried „Cluster Detection“:

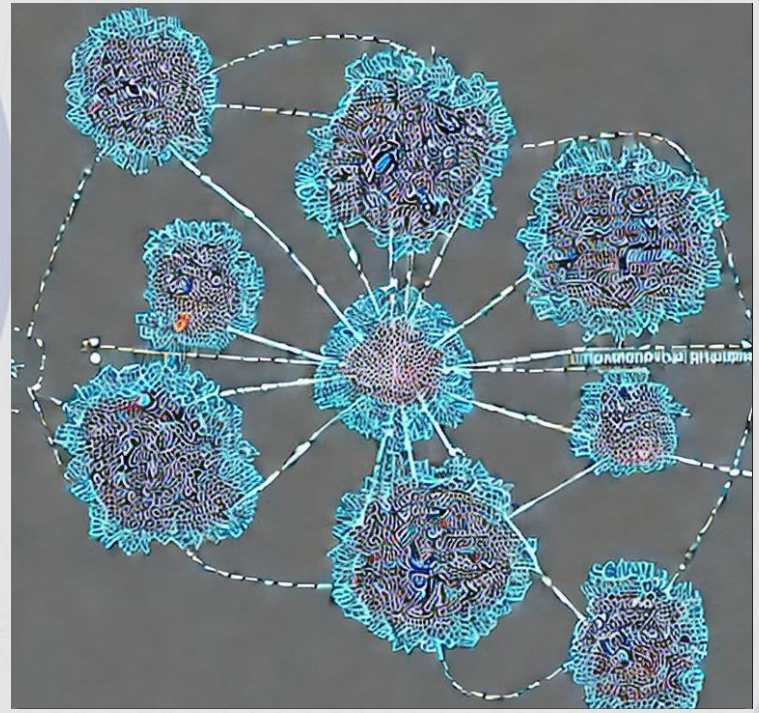

What is Artificial Intelligence?

Then I tried „Cluster Detection“:

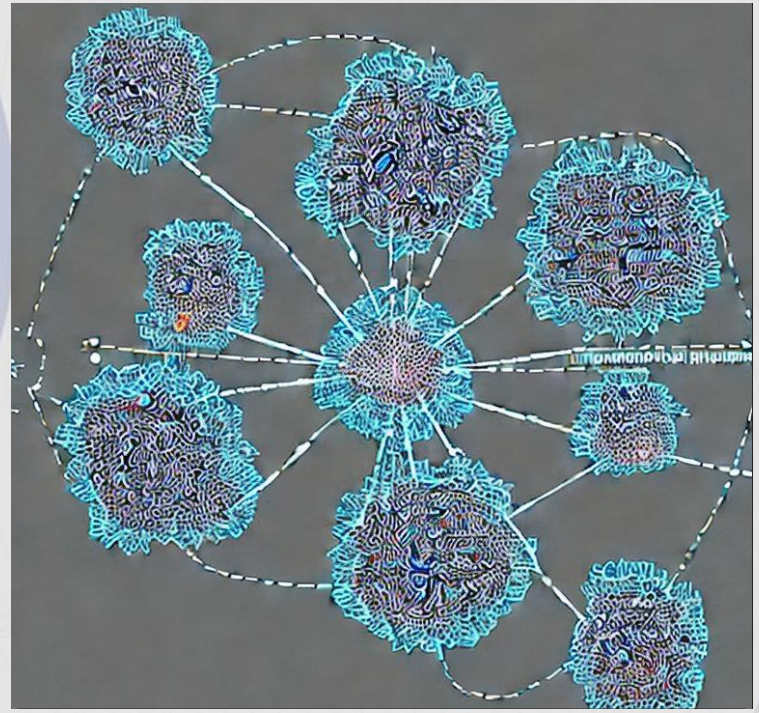

And Anomaly Detection:

What is Artificial Intelligence?

Then I tried „Cluster Detection“:

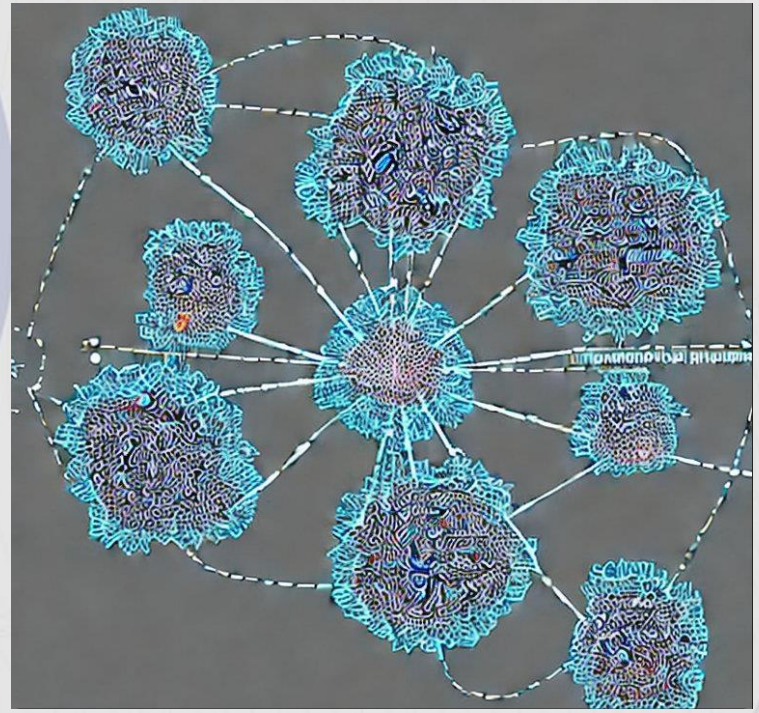

And Anomaly Detection:

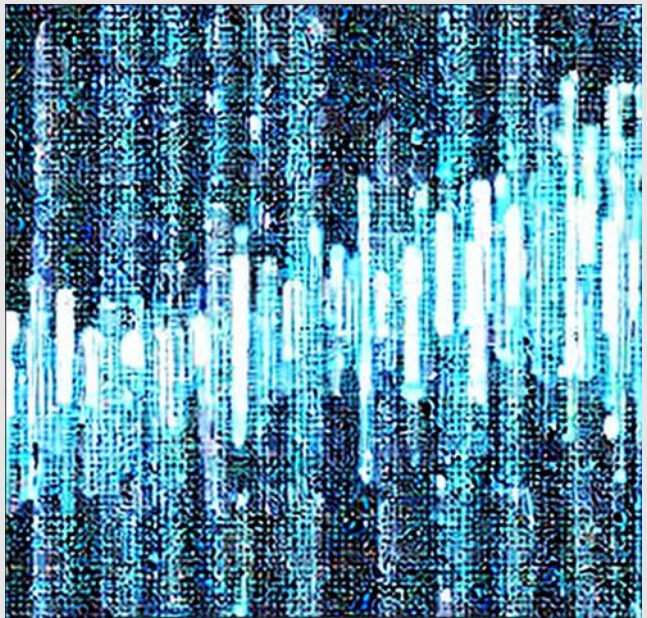

What is Artificial Intelligence?

Then I tried a different generator with „COBOL AI DB2“:

What is Artificial Intelligence?

Then I tried a different generator with „COBOL AI DB2“:

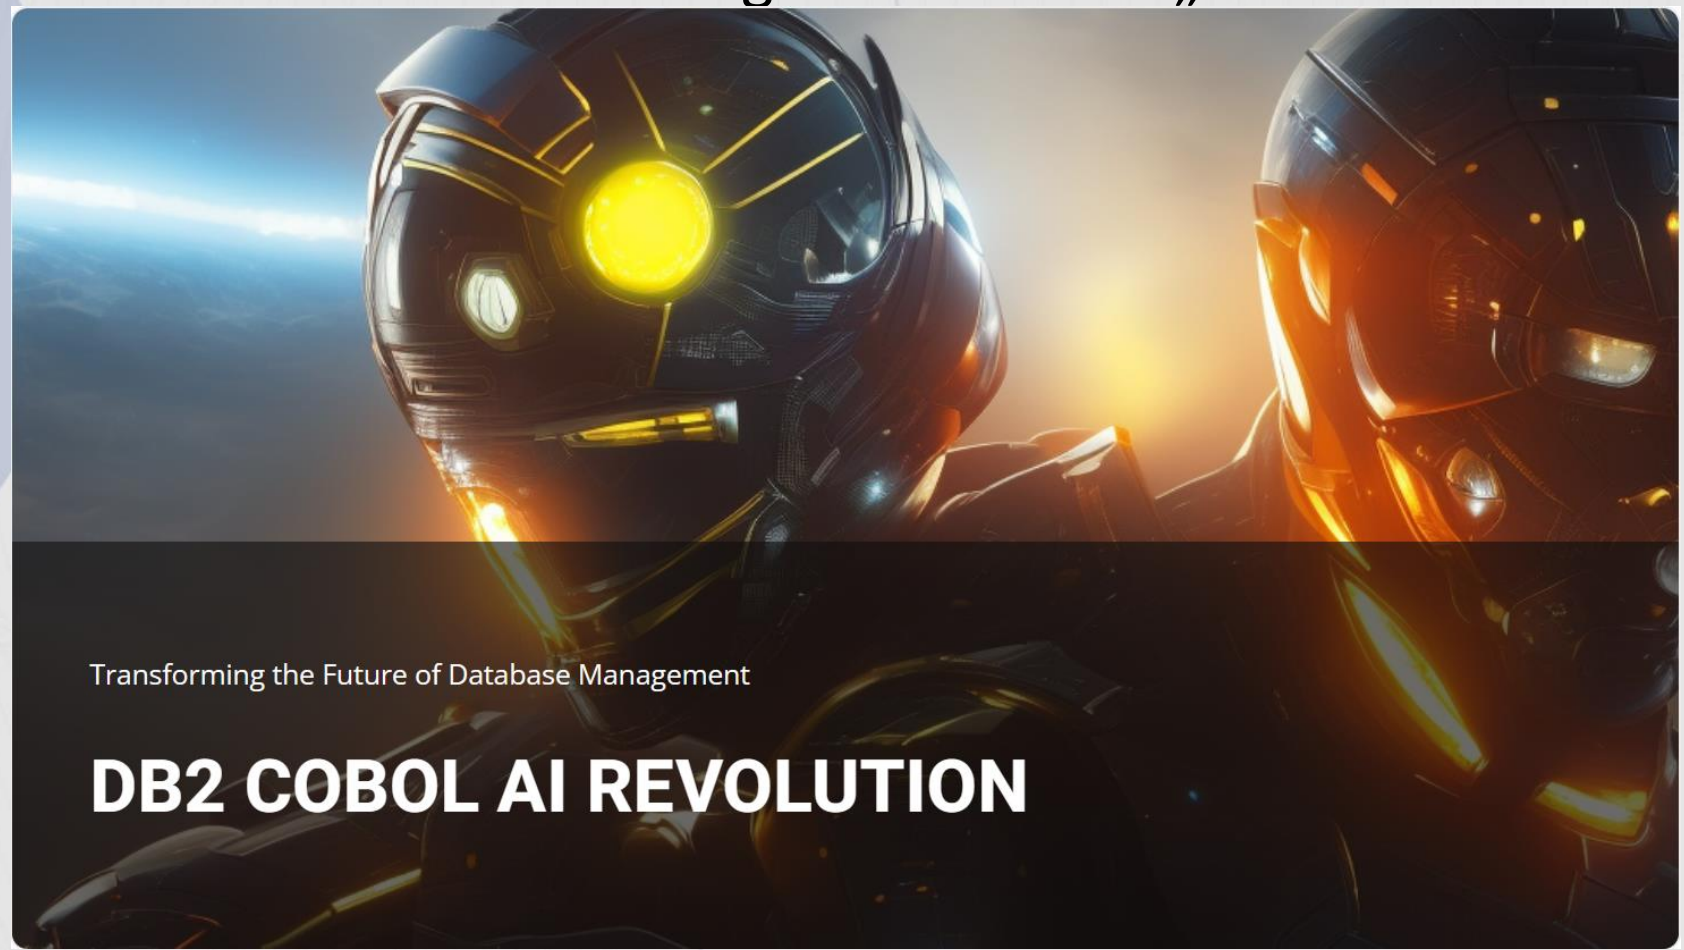

What is Artificial Intelligence?

WOW! But it went off on one as well:

What is Artificial Intelligence?

WOW! But it went off on one as well:

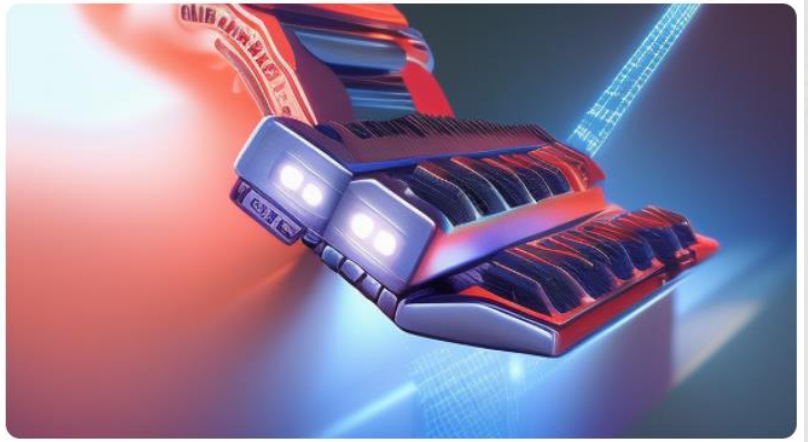

WHAT IS DB2 COBOL?

DB2 COBOL is a powerful, high-performance database management system that integrates with the COBOL programming language, providing a robust and reliable platform for data storage and retrieval.

DB2 COBOL combines a powerful database with the COBOL language for a solid data management solution.

What is Artificial Intelligence?

The last page of the generated presentation looked like:

What is Artificial Intelligence?

The last page of the generated presentation looked like:

EMBRACE THE DB2 COBOL AI REVOLUTION

By adopting the powerful combination of DB2 COBOL and artificial intelligence, businesses can harness the potential of this cutting-edge technology to revolutionize their operations and achieve greater success in the digital era.

What is Artificial Intelligence?

The latest trend is AI-written books that can actually end up killing you!

AI-written books on mushrooms and fungi and the grave danger they pose

August 30, 2023 By Sovan Mandal — 0 Comments

What is Artificial Intelligence?

The website decrypt.co listed several such suspect titles, which include *The Ultimate Mushroom Books Field Guide of the Southwest*, or the *Wild Mushroom Cookbook For Beginner*. Interestingly, both the titles seem to have been removed which is a good thing considering that the information it contained was largely inaccurate and even fictitious.

It is still a dangerous trend given that the stakes are as high as someone's life. The risks are even higher for children or senior citizens. What is even more dangerous, is that there is practically no way one can determine such books to have been written by AI tools that haven't been adequately proofread. Hence, they can be of considerable danger to the unsuspecting reader.

Maybe with the advent of AI tools such as ChatGPT, we now need a tool that can accurately differentiate content written by humans and machines.

What is Artificial Intelligence?

Scary!

What is Artificial Intelligence?

Scary!

What is Artificial Intelligence?

So you can see, it might write you a love letter, an interpretation of „Der Gärtner“ from Joseph Freiherr von Eichendorff and translate code snippets between Python and C++, but it is ***not*** Intelligent!

Fingers anyone?

What is Artificial Intelligence?

So you can see, it might write you a love letter, an interpretation of „Der Gärtner“ from Joseph Freiherr von Eichendorff and translate code snippets between Python and C++, but it is ***not*** Intelligent!

Fingers anyone?

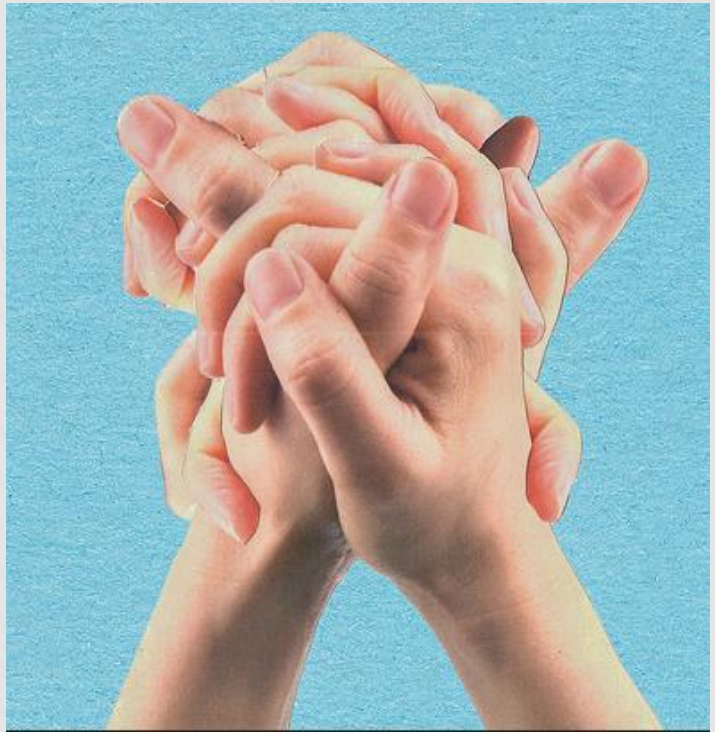

What is Artificial Intelligence?

Xing article from 2023-04-18 by Thomas Knüwer:

Künstliche Intelligenz – mal nüchtern betrachtet

Wie intelligent sind die aktuellen AI-Programme?

Das ist eine Frage der Definition von Intelligenz. Weshalb ich diese Vokabel im Zusammenhang mit Software auch für schwierig halte. Eigentlich müsste man sagen:

> **Die aktuellen KI-Projekte haben eine Intelligenz von null.**

Denn: Sie verstehen ja nicht, was sie da gerade tun. Sie arrangieren Worte und Bilder um eine vorgegebene Aufgabe möglichst gut zu erfüllen.

Doch wie gesagt: Ich halte Intelligenz für die falsche Maßgröße im Rahmen der Debatte. Sinnvoller ist aus meiner Sicht die Frage:

> **Wie leistungsfähig sind die aktuellen AI-Programme?**

Bei den meisten derzeit öffentlich zugänglichen Programmen kommt erst Begeisterung auf und dann recht schnell Ernüchterung. Meine jährliche Trendprognose habe ich zum Beispiel diesmal mit KI-Bildern illustriert und die Ergebnisse fand ich eher wenig euphorisierend.

https://www.xing.com/news/article/kuenstliche-intelligenz-mal-nuechte-1

What is Artificial Intelligence?

Xing article from 2023-04-18 by Thomas Knüwer:

Artificial intelligence – a sober look

How intelligent are current AI programs?

This is a question of the definition of intelligence. Which is why I also find this vocabulary difficult in the context of software. Actually, one would have to say:

Current AI projects have an intelligence of zero.

After all, they don‘t understand what they are doing. They are arranging words and images to fulfill a given task as well as possible.

But as I said, I think intelligence is the wrong measure for this debate. More meaningful, in my view, is the question:

How powerful are the current AI programs?

With most of the programs currently available to the public, enthusiasm arises at first and then disillusionment quite quickly. This time, for example, I illustrated my annual trend forecast with AI images and found the results rather less euphoric.

https://www.xing.com/news/article/kuenstliche-intelligenz-mal-nuechte-1

Agenda

- What is Artificial Intelligence?
- Db2 13 and Data Insights - AI for free?
- SEG and AI
- Q&A

Db2 13 and Data Insights – AI for free?

Db2 13 FL500 brings three new Scalar Built-in Functions (BiFs). These are the new SQL Data Insights (DI) BiFs.

First, you need to make sure you have all the prereqs in place. These are basically a couple of APARs for the IBM Z Deep Neural Network Library (zDNN), the z/OS Supervisor, IBM OpenBLAS, z/OS OpenSSH and IBM 64-bit SDK for z/OS Java.

zDNN and OpenBLAS come with z/OS 2.4/2.5, but without the required APARs, the libraries may be empty.

Db2 13 and Data Insights – AI for free?

SQL Data Insights (SQL DI) is a kind of no-charge add-on to Db2 13, so you need to order and install it separately (FMID HDBDD18).

Then you need to install and customize SQL DI, starting with the definition of a (technical) user along with its appropriate authorization.

Db2 13 and Data Insights – AI for free?

They ask for 100 GB of storage for the zFS home directory, but I think you'll only need that when you start to run AI model training on vast amounts of data.
For my first tiny steps into the world of Db2 AI, it worked with a tenth of that without any problems.

The requirements listed for CPU and system memory aren't much smaller, and I'm experiencing a very measurable CPU consumption whenever the Model training on an object starts.

Db2 13 and Data Insights – AI for free?

Next up is a bunch of RACF defs and the first occasion where the dreaded “certificates” requirement raises its ugly head!

While waiting for the RACDCERT to be sorted out, you can now execute the DDL in the DSNTIJAI member in the SDSNSAMP lib. This creates the pseudo-catalog tables and other artifacts that SQL DI requires.

Db2 13 and Data Insights – AI for free?

You also have to make sure that WLM has the correct Utility DD Cards and the NUMTCB=1 override for the LOAD.

It is ***very*** annoying to wait one hour and then have the LOAD fail due to a DD Card being missing!

Pro tip: Start ***small*** and increase slowly!

Db2 13 and Data Insights – AI for free?

Now the real install process begins with the “standard” SMP/E install of SQL DI and the execution of the Install Script in the USS environment.

Be very careful with the values you use for all of these and stick with them to the bitter end!

Finally, you can create an STC for SQL DI which is highly recommended. Conversely, the STC for Spark I would not bother with, as it is fully under the control of SQL DI anyway.

Db2 13 and Data Insights – AI for free?

This is how DI looks as a whole:

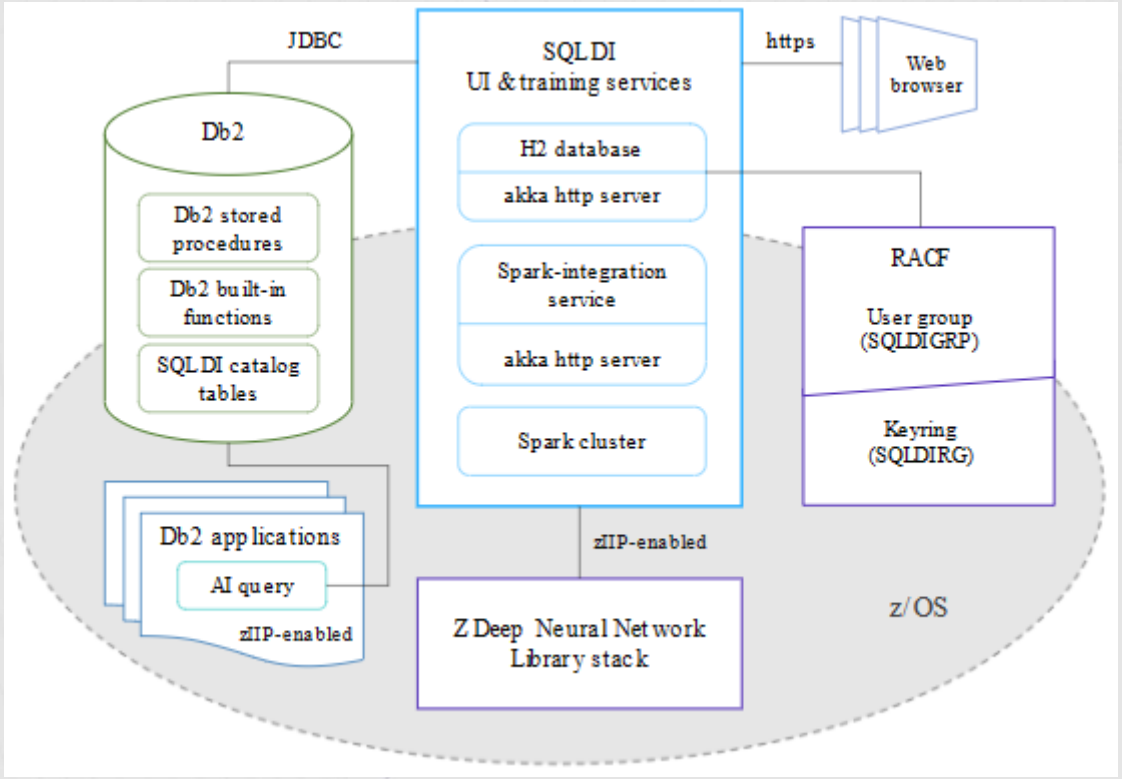

https://www.ibm.com/docs/en/db2-for-zos/13?topic=running-ai-queries-sql-data-insights

Db2 13 and Data Insights – AI for free?

After having started the STC you can sign-on through the web interface:

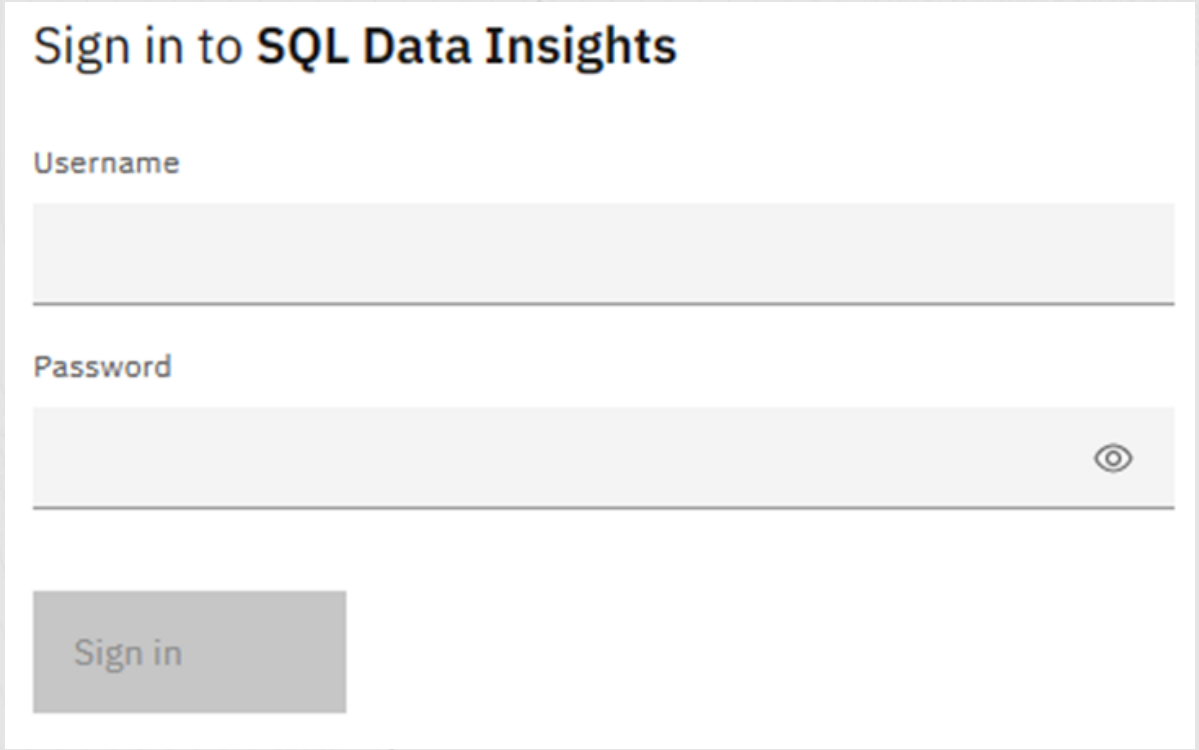

Db2 13 and Data Insights – AI for free?

The “tricky” part is choosing the “Usual Suspects”. Which columns should be used as input to the Model? You only get one “key” column and so I chose STMT_ID and then 16 other columns from WLX:

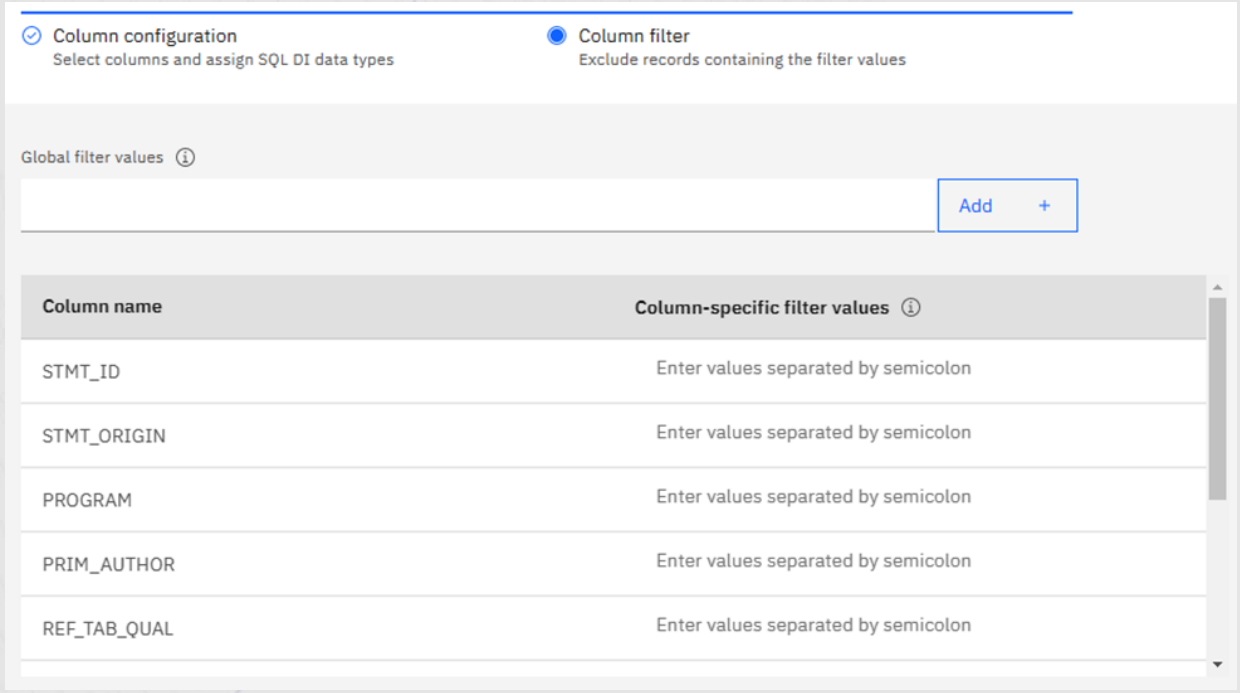

Db2 13 and Data Insights – AI for free?

Once you've done all that, you start the Model on its merry way…

JOBNAME	StepName	ProcStep	JobID	Owner	C	Pos	DP	Real	Paging	SIO	CPU%
SQLDSPKX	SQLDSPKX		STC09599	SQLDIID		IN	DE	12T	0.00	907.18	3.81
SQLDSPKX	SQLDSPKX		STC09598	SQLDIID		IN	DE	13T	0.00	940.71	4.08
SQLDSPKD	SQLDSPKD		STC09597	SQLDIID		IN	DE	39T	0.00	2177.2	64.75
SQLDSPKX	SQLDSPKX		STC09596	SQLDIID		IN	DE	11T	0.00	892.00	4.08
SQLDSPKX	SQLDSPKX		STC09595	SQLDIID		IN	DE	11T	0.00	940.73	4.20
SQLDAPPS	SQLDAPPS	STEP1	STC08994	SQLDIID		LO	FF	407	0.00	0.00	0.00
SQLDAPPS	SQLDAPPS		STC08972	SQLDIID		IN	DE	85T	0.00	2.56	5.37
SQLDAPPS	SQLDAPPS		STC08877	SQLDIID		LO	FF	3401	0.00	0.00	0.00
SQLDSPKM	SQLDSPKM		STC07682	SQLDIID		IN	DE	35T	0.00	0.85	0.26
SQLDAPPS	SQLDAPPS		STC07665	SQLDIID		LO	FF	409	0.00	0.00	0.00
SQLDSPKW	SQLDSPKW		STC07659	SQLDIID		IN	DE	32T	0.00	110.48	1.15

Yes, you can get multiple cups of coffee as you break all local records for SIO, CPU and zIIP CPU usage at your site!

After 40 minutes elapsed, the LOAD Utility got triggered and I was finished!

Db2 13 and Data Insights – AI for free?

Once you've done all that, you start the Model on its merry way…

JOBNAME	StepName	ProcStep	JobID	Owner	C	Pos	DP	Real	Paging	SIO	CPU%
SQLDSPKX	SQLDSPKX		STC09599	SQLDIID		IN	DE	12T	0.00	907.18	3.81
SQLDSPKX	SQLDSPKX		STC09598	SQLDIID		IN	DE	13T	0.00	940.71	4.08
SQLDSPKD	SQLDSPKD		STC09597	SQLDIID		IN	DE	39T	0.00	2177.2	64.75
SQLDSPKX	SQLDSPKX		STC09596	SQLDIID		IN	DE	11T	0.00	892.00	4.08
SQLDSPKX	SQLDSPKX		STC09595	SQLDIID		IN	DE	11T	0.00	940.73	4.20
SQLDAPPS	SQLDAPPS	STEP1	STC08994	SQLDIID		LO	FF	407	0.00	0.00	0.00
SQLDAPPS	SQLDAPPS		STC08972	SQLDIID		IN	DE	85T	0.00	2.56	5.37
SQLDAPPS	SQLDAPPS		STC08877	SQLDIID		LO	FF	3401	0.00	0.00	0.00
SQLDSPKM	SQLDSPKM		STC07682	SQLDIID		IN	DE	35T	0.00	0.85	0.26
SQLDAPPS	SQLDAPPS		STC07665	SQLDIID		LO	FF	409	0.00	0.00	0.00
SQLDSPKW	SQLDSPKW		STC07659	SQLDIID		IN	DE	32T	0.00	110.48	1.15

Yes, you can get multiple cups of coffee as you break all local records for SIO, CPU and zIIP CPU usage at your site!

After 40 minutes elapsed, the LOAD Utility got triggered and I was finished!

This was with only 64,390 rows…

Db2 13 and Data Insights – AI for free?

Once you've done all that, you start the Model on its merry way…

JOBNAME	StepName	ProcStep	JobID	Owner	C	Pos	DP	Real	Paging	SIO	CPU%
SQLDSPKX	SQLDSPKX		STC09599	SQLDIID		IN	DE	12T	0.00	907.18	3.81
SQLDSPKX	SQLDSPKX		STC09598	SQLDIID		IN	DE	13T	0.00	940.71	4.08
SQLDSPKD	SQLDSPKD		STC09597	SQLDIID		IN	DE	39T	0.00	2177.2	64.75
SQLDSPKX	SQLDSPKX		STC09596	SQLDIID		IN	DE	11T	0.00	892.00	4.08
SQLDSPKX	SQLDSPKX		STC09595	SQLDIID		IN	DE	11T	0.00	940.73	4.20
SQLDAPPS	SQLDAPPS	STEP1	STC08994	SQLDIID		LO	FF	407	0.00	0.00	0.00
SQLDAPPS	SQLDAPPS		STC08972	SQLDIID		IN	DE	85T	0.00	2.56	5.37
SQLDAPPS	SQLDAPPS		STC08877	SQLDIID		LO	FF	3401	0.00	0.00	0.00
SQLDSPKM	SQLDSPKM		STC07682	SQLDIID		IN	DE	35T	0.00	0.85	0.26
SQLDAPPS	SQLDAPPS		STC07665	SQLDIID		LO	FF	409	0.00	0.00	0.00
SQLDSPKW	SQLDSPKW		STC07659	SQLDIID		IN	DE	32T	0.00	110.48	1.15

Yes, you can get multiple cups of coffee as you break all local records for SIO, CPU and zIIP CPU usage at your site!

After 40 minutes elapsed, the LOAD Utility got triggered and I was finished!

This was with only 64,390 rows…but our machine is only 14 MSUs

Db2 13 and Data Insights – AI for free?

You may now look into the SYSAIDB.* tables to review what it has saved but it is not really required. As long as the LOAD has completed you are good to go!

Now you can use the RUN QUERY button in the WEB UI or you can simply write SQL in SPUFI to see what the AI BiFs have given you.

Db2 13 and Data Insights – AI for free?

One last performance point: 500,000 rows took five hours to model before being cancelled as it had taken 99% of local paging datasets and our machine was about to die.

Looking in the logs it was only half way through the analysis…

Db2 13 and Data Insights – AI for free?

One last performance point: 500,000 rows took five hours to model before being cancelled as it had taken 99% of local paging datasets and our machine was about to die.

Looking in the logs it was only half way through the analysis…

Pro tip: Start ***small*** and increase slowly!

Db2 13 and Data Insights – AI for free?

Naturally the question hanging in the air is:

Db2 13 and Data Insights – AI for free?

Naturally the question hanging in the air is:

Was that really all free?

Db2 13 and Data Insights – AI for free?

OK, we have now our “model” – What can we do with it?

Db2 13 and Data Insights – AI for free?

OK, we have now our “model” – What can we do with it?

SQL DI comes with three BiFs out-of-the-box:

Db2 13 and Data Insights – AI for free?

OK, we have now our “model” – What can we do with it?

SQL DI comes with three BiFs out-of-the-box:

1) AI_ANALOGY
2) AI_SEMANTIC_CLUSTER
3) AI_SIMILARITY

Db2 13 and Data Insights – AI for free?

OK, we have now our “model” – What can we do with it?

SQL DI comes with ~~three~~ four BiFs out-of-the-box:

1) AI_ANALOGY
2) AI_SEMANTIC_CLUSTER
3) AI_SIMILARITY
4) AI_COMMONALITY – OA64845 zADE v1.2.0 – UJ93339 – 2023-08-02

Agenda

- What is Artificial Intelligence?
- Db2 13 and Data Insights - AI for free?
- SEG and AI
- Q&A

SEG and AI

The AI_ANALOGY is linking various arguments to then see what hidden links are also there.

The classic is: “Roy” is to “Beer” what “Andre” is to ???

In Db2 terms you supply two pairs of source and target values, which then returns a double-precision number between -1 and +1 with +1 being a very good analogy.

SEG and AI

For AI_ANALOGY I wanted to see which tables are being used analogous to BOXWEL3.IQATW001 with creator KKKK:

```
SELECT AI_ANALOGY('BOXWEL3'  USING MODEL COLUMN PRIM_AUTHOR,
                  'IQATW001' USING MODEL COLUMN REF_TABLE  ,
                  'KKKK'     USING MODEL COLUMN PRIM_AUTHOR,
                  REF_TABLE ) AS AI_VALUE
      ,A.WLX_TIMESTAMP
      ,A.STMT_ID
      ,A.STMT_TIMESTAMP
      ,SUBSTR(A.PRIM_AUTHOR , 1 , 8 ) AS PRIM_AUTHOR
      ,SUBSTR(A.PROGRAM     , 1 , 8 ) AS PROGRAM
      ,SUBSTR(A.REF_TABLE   , 1 , 18) AS REF_TABLE
      ,A.EXECUTIONS
      ,A.GETP_OPERATIONS
      ,A.ELAPSE_TIME
      ,A.CPU_TIME
      ,A.STMT_TEXT
FROM IQA061QB.IQATW001 A
WHERE A.PRIM_AUTHOR = 'KKKK'
  AND  AI_ANALOGY('BOXWEL3'    USING MODEL COLUMN PRIM_AUTHOR,
                  'IQATW001'   USING MODEL COLUMN REF_TABLE  ,
                  'KKKK'       USING MODEL COLUMN PRIM_AUTHOR,
                  REF_TABLE ) > 0.5
ORDER BY 1 DESC -- SHOW BEST FIRST
--ORDER BY 1    -- SHOW WORST FIRST
FETCH FIRST 2000 ROWS ONLY ;
```

SEG and AI

The AI_SEMANTIC_CLUSTER clusters data together enabling you to check if a given new value is within the generated cluster.

The classic is: “Roy Beer Count” , 1 , 5 , 20 [trending to ∞]

In Db2 terms you supply one to three cluster values and then it tests whether the first value is within this cluster, which then returns a double-precision number between -1 and +1 with +1 being a very good fit within the cluster.

SEG and AI

For AI_SEMANTIC_CLUSTER I am looking for any programs that are in the cluster SYSLN200 and DSNATYPU. I know that these are dynamic SQL issuing programs so I am also looking for any other dynamic SQL issuing programs:

```
SELECT AI_SEMANTIC_CLUSTER (PROGRAM, 'SYSLN200', 'DSNATYPU')
                                    AS SIMILARITY
      ,A.STMT_ORIGIN
      ,SUBSTR(A.STMT_TEXT , 1 , 60) AS STMT_TEXT
      ,SUBSTR(A.PROGRAM , 1 , 8)    AS PROGRAM
      ,A.STMT_ID
      ,A.EXECUTIONS
      ,A.GETP_OPERATIONS
      ,A.ELAPSE_TIME
      ,A.CPU_TIME
FROM DAIN0610.IQATW001 A
WHERE 1 = 1
  AND A.PROGRAM              NOT IN ('SYSLN200', 'DSNATYPU')
  AND AI_SEMANTIC_CLUSTER (PROGRAM, 'SYSLN200', 'DSNATYPU')
                          > 0.3
ORDER BY SIMILARITY DESC
FETCH FIRST 20 ROWS ONLY;
```

This only worked with > 0.3 while it found none with > 0.5

SEG and AI

The semantic cluster did find just Dynamic SQL in the cluster of those two Dynamic SQL Packages:

SIMILARITY	STMT_ORIGIN	STMT_TEXT	PROGRAM
+0.44718432426452 64E+00	D	SELECT DISTINCT PG.COLLID FROM SYSIBM.SYSPACKLIST PL, SYSIBM	DSNCLINF
+0.44718432426452 64E+00	D	SELECT * FROM SYSIBM.SYSDBRM WHERE NAME = '[redacted]' AND PLNAM	DSNCLINF
+0.44718432426452 64E+00	D	SELECT STMT, STMTNO FROM SYSIBM.SYSPACKSTMT WHERE LOCATION =	DSNCLINF
+0.44718432426452 64E+00	D	SELECT COLLID, OWNER, CREATOR, VERSION, PDSNAME, QUALIFIER,	DSNCLINF
+0.44718432426452 64E+00	D	SELECT STMT, STMTNO FROM SYSIBM.SYSPACKSTMT WHERE LOCATION =	DSNCLINF
+0.44718432426452 64E+00	D	SELECT COLLID, OWNER, CREATOR, VERSION, PDSNAME, QUALIFIER,	DSNCLINF
+0.44718432426452 64E+00	D	SELECT DISTINCT PG.COLLID FROM SYSIBM.SYSPACKLIST PL, SYSIBM	DSNCLINF
+0.44718432426452 64E+00	D	SELECT * FROM SYSIBM.SYSDBRM WHERE NAME = '[redacted]' AND PLNAM	DSNCLINF
+0.44718432426452 64E+00	D	SELECT * FROM SYSIBM.SYSPLAN WHERE NAME = '[redacted]' FOR FET	DSNCLINF
+0.44718432426452 64E+00	D	SELECT * FROM SYSIBM.SYSPLAN WHERE NAME = '[redacted]' FOR FET	DSNCLINF

However, if the program was not in the training data you get back the NULL value…

SEG and AI

The AI_SIMILARITY reports on how similar two columns or expressions are.

The classic is: “Roy & Beer” , “Roy & Whisky”

In Db2 terms you supply two values and it returns a double-precision number between -1 and +1 with +1 being a very similar and -1 being very dissimilar (So “Roy & Water” in this case!)

SEG and AI

For AI_SIMILARITY I am searching for any programs that behave like our main Dynamic SQL program IQADBACP:

```
SELECT AI_SIMILARITY( PROGRAM,
                      'IQADBACP') AS AI_VALUE
      ,A.STMT_ORIGIN
      ,SUBSTR(A.STMT_TEXT , 1 , 60) AS STMT_TEXT
      ,SUBSTR(A.PROGRAM , 1 , 8) AS PROGRAM
      ,A.WLX_TIMESTAMP
      ,A.STMT_ID
      ,A.STMT_TIMESTAMP
      ,A.EXECUTIONS
      ,A.GETP_OPERATIONS
      ,A.ELAPSE_TIME
      ,A.CPU_TIME
FROM DAIN0610.IQATW001 A
WHERE 1 = 1
  AND NOT A.PROGRAM = 'IQADBACP'
  AND AI_SIMILARITY( PROGRAM,
                     'IQADBACP') IS NOT NULL
ORDER BY 1 DESC -- SHOW BEST FIRST
--ORDER BY 1    -- SHOW WORST FIRST
FETCH FIRST 10 ROWS ONLY;
```

SEG and AI

This is the result of looking for similar programs to our main dynamic SQL program:

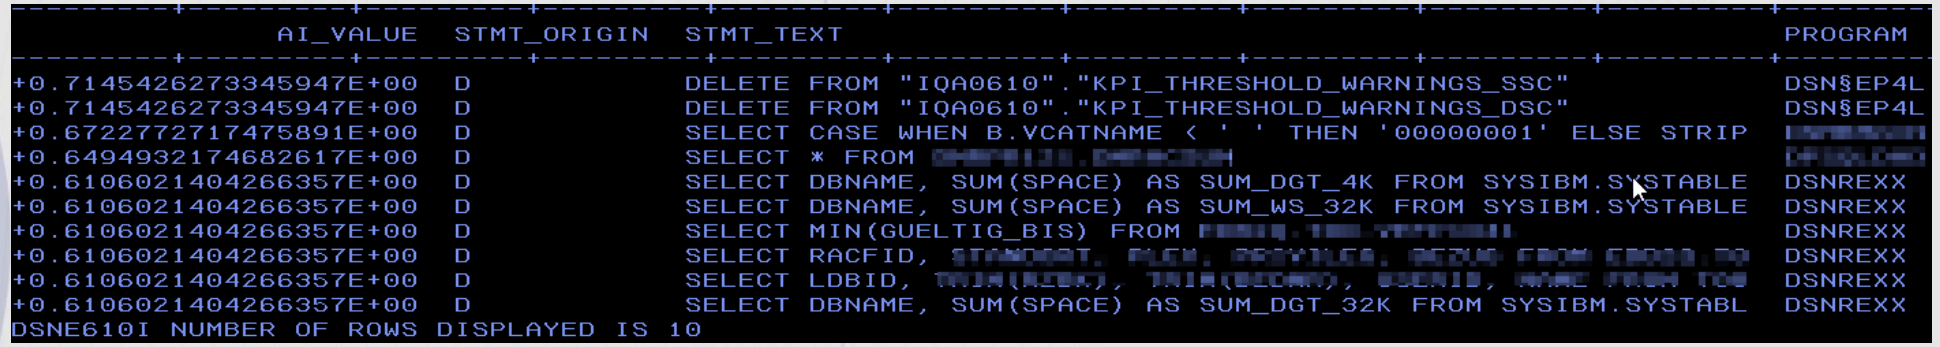

AI_VALUE	STMT_ORIGIN	STMT_TEXT	PROGRAM
+0.7145426273345947E+00	D	DELETE FROM "IQA0610"."KPI_THRESHOLD_WARNINGS_SSC"	DSN§EP4L
+0.7145426273345947E+00	D	DELETE FROM "IQA0610"."KPI_THRESHOLD_WARNINGS_DSC"	DSN§EP4L
+0.6722772717475891E+00	D	SELECT CASE WHEN B.VCATNAME < ' ' THEN '00000001' ELSE STRIP	[illegible]
+0.6494932174682617E+00	D	SELECT * FROM [illegible]	[illegible]
+0.6106021404266357E+00	D	SELECT DBNAME, SUM(SPACE) AS SUM_DGT_4K FROM SYSIBM.SYSTABLE	DSNREXX
+0.6106021404266357E+00	D	SELECT DBNAME, SUM(SPACE) AS SUM_WS_32K FROM SYSIBM.SYSTABLE	DSNREXX
+0.6106021404266357E+00	D	SELECT MIN(GUELTIG_BIS) FROM [illegible]	DSNREXX
+0.6106021404266357E+00	D	SELECT RACFID, [illegible]	DSNREXX
+0.6106021404266357E+00	D	SELECT LDBID, [illegible]	DSNREXX
+0.6106021404266357E+00	D	SELECT DBNAME, SUM(SPACE) AS SUM_DGT_32K FROM SYSIBM.SYSTABL	DSNREXX

DSNE610I NUMBER OF ROWS DISPLAYED IS 10

Yes indeed – It did find Dynamic SQL issuing programs!

SEG and AI

Finally using a not similar version of AI_SIMILARITY:

```
SELECT * FROM
    (SELECT AI_SIMILARITY(PROGRAM, 'SYSLN200') AS SIMILARITY
      ,C.STMT_ORIGIN
      ,SUBSTR(C.STMT_TEXT , 1 , 60) AS STMT_TEXT
      ,SUBSTR(C.PROGRAM , 1 , 8) AS PROGRAM
      ,C.STMT_ID
      ,C.END_USERID
      ,C.STMT_TIMESTAMP
      ,C.EXECUTIONS
      ,C.GETP_OPERATIONS
      ,C.ELAPSE_TIME
      ,C.CPU_TIME
     FROM DAIN0610.IQATW001 C
     WHERE PROGRAM <> 'SYSLN200')
   WHERE SIMILARITY < 0.5
   ORDER BY SIMILARITY   ASC
FETCH FIRST 200 ROWS ONLY
;
```

Notice that this one is doing dissimilar!

SEG and AI

And it works! It truly finds in this case ***just*** Static SQL doing stuff nowhere like the Dynamic SQL Package we used in the previous example:

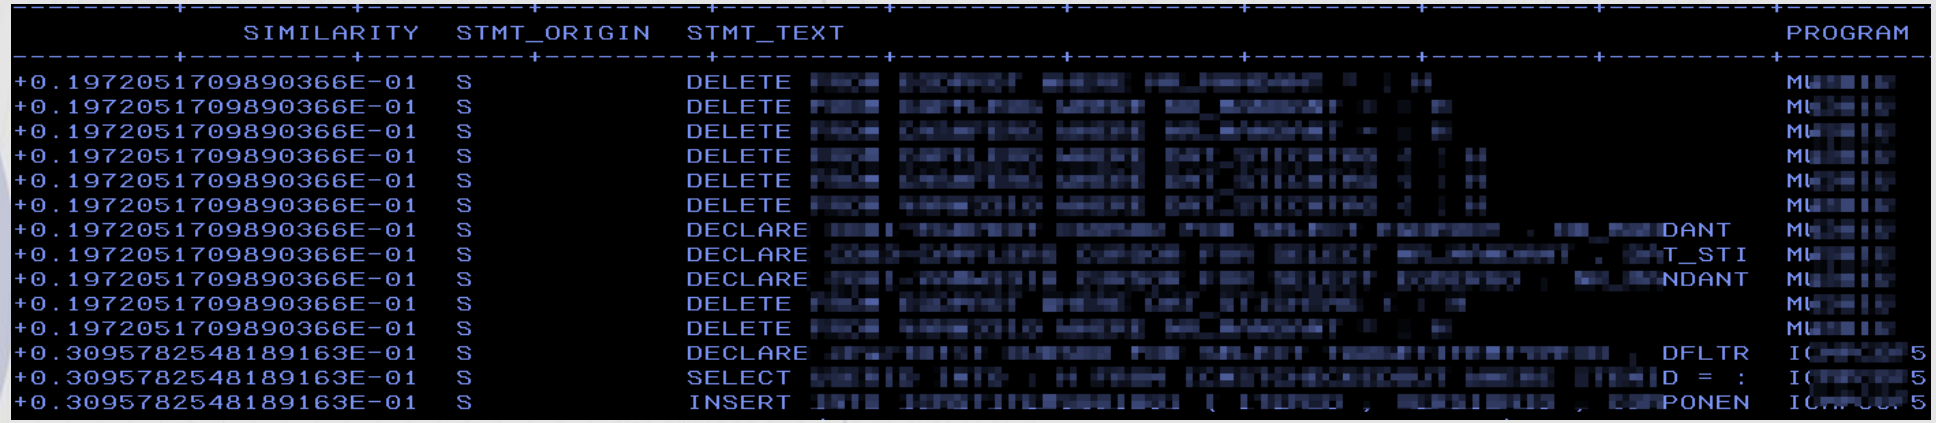

SIMILARITY	STMT_ORIGIN	STMT_TEXT	PROGRAM
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DECLARE DANT	MU
+0.1972051709890366E-01	S	DECLARE T_STI	MU
+0.1972051709890366E-01	S	DECLARE NDANT	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.1972051709890366E-01	S	DELETE	MU
+0.3095782548189163E-01	S	DECLARE DFLTR	IC 5
+0.3095782548189163E-01	S	SELECT D = :	IC 5
+0.3095782548189163E-01	S	INSERT PONEN	IC 5

This is actually pretty neat!

SEG and AI

The AI_COMMONALITY is so new I have ***no*** idea what it does! However, in an IBM presentation with all the usual disclaimers and a **BOLD** heading "**Future**", was this little nugget:

Pattern Queries	**Detect global patterns**	**Identify entities that are outliers or with most common behavior**

See also APAR OA64845 zADE v1.2.0

PTF UJ93339 – Closed 2023-08-02

SEG and AI

As you have seen these queries actually ***do*** find “hidden patterns” as long as you have an idea about what is ***in*** the data.

All three (and possibly four) AI BiFs come with nasty little notes like:

> “The result can be null; if any argument is null, the result is the null value. If the arguments to the function contain values that **were not seen during model training**, and the model column is trained as categorical, the result is the null value.”

This is, for me, a bit of a killer…

SEG and AI

The next problem is that to find an abnormal result you must **know** the abnormal result first.

Further, and worse in my opinion, is that in the Model data used you ***cannot*** have any **abnormal** data within, otherwise the Model will not "see" that as being abnormal in a query.

How can you possibly know this?

[Naturally here we have the possibility that AI_COMMONALITY will enable outlier detection but as of now this is not known to me.]

SEG and AI

The following list of APARs must all be applied:

```
HIPER UI80601 PH45358: WHEN RUNNING SQL DATA INSIGHTS Z16 WITH ZAIU GET
                       ABEND=U4039
      UI80917 PH46488: ABEND MIGHT OCCUR WHEN CALLING DB2 SQL DATA INSIGHTS
                       BUILT-IN FUNCTIONS (AI_SIMILARITY, AI_ANALOGY,
                       AI_SEMANTIC_CLUSTER)
      UI83744 PH49781: SQLCODE404 OR SQLCODE104 FROM DSNTIJAI JOB FOR STEP
                       DSNTIAI3 WHEN RUNNING WITH JAPANESE KATAKANA CCSID OR
                       DECP DECIMAL=COMMA [Space within table name and DROP
                       problems]
      UI90547 PH51892: NEW FUNCTION [Love these titles! This is a performance
                       booster for AI_SEMANTIC_CLUSTER]
HIPER UI91462 PH53069: ABEND05E RC28 AND DB2 ABNORMAL TERMINATION RC00E50054
                       FOR SQL DATA INSIGHTS QUERY RUNNING FROM A DDF THREAD
                       [Db2 crash!]
      UI93117 PH55212: NEW FUNCTION [Db2 new function for SQLDI query]
```

SEG and AI

So, after all that, where is AI in the real world of Db2 on z/OS?

SEG and AI

So, after all that, where is AI in the real world of Db2 on z/OS?

SEG and AI

But lets look on the brighter side of life shall we?

SEG and AI

But lets look on the brighter side of life shall we?

We decided to write our own “Looking into the Future” style of AI.

SEG and AI

But lets look on the brighter side of life shall we?

We decided to write our own “Looking into the Future” style of AI.

So we added a capability to look into the future into our Space Assurance Expert product SAX which then, based upon the history of the Tablespace or Index, projects its future growth and looks pretty cool in Zowe!

SEG and AI

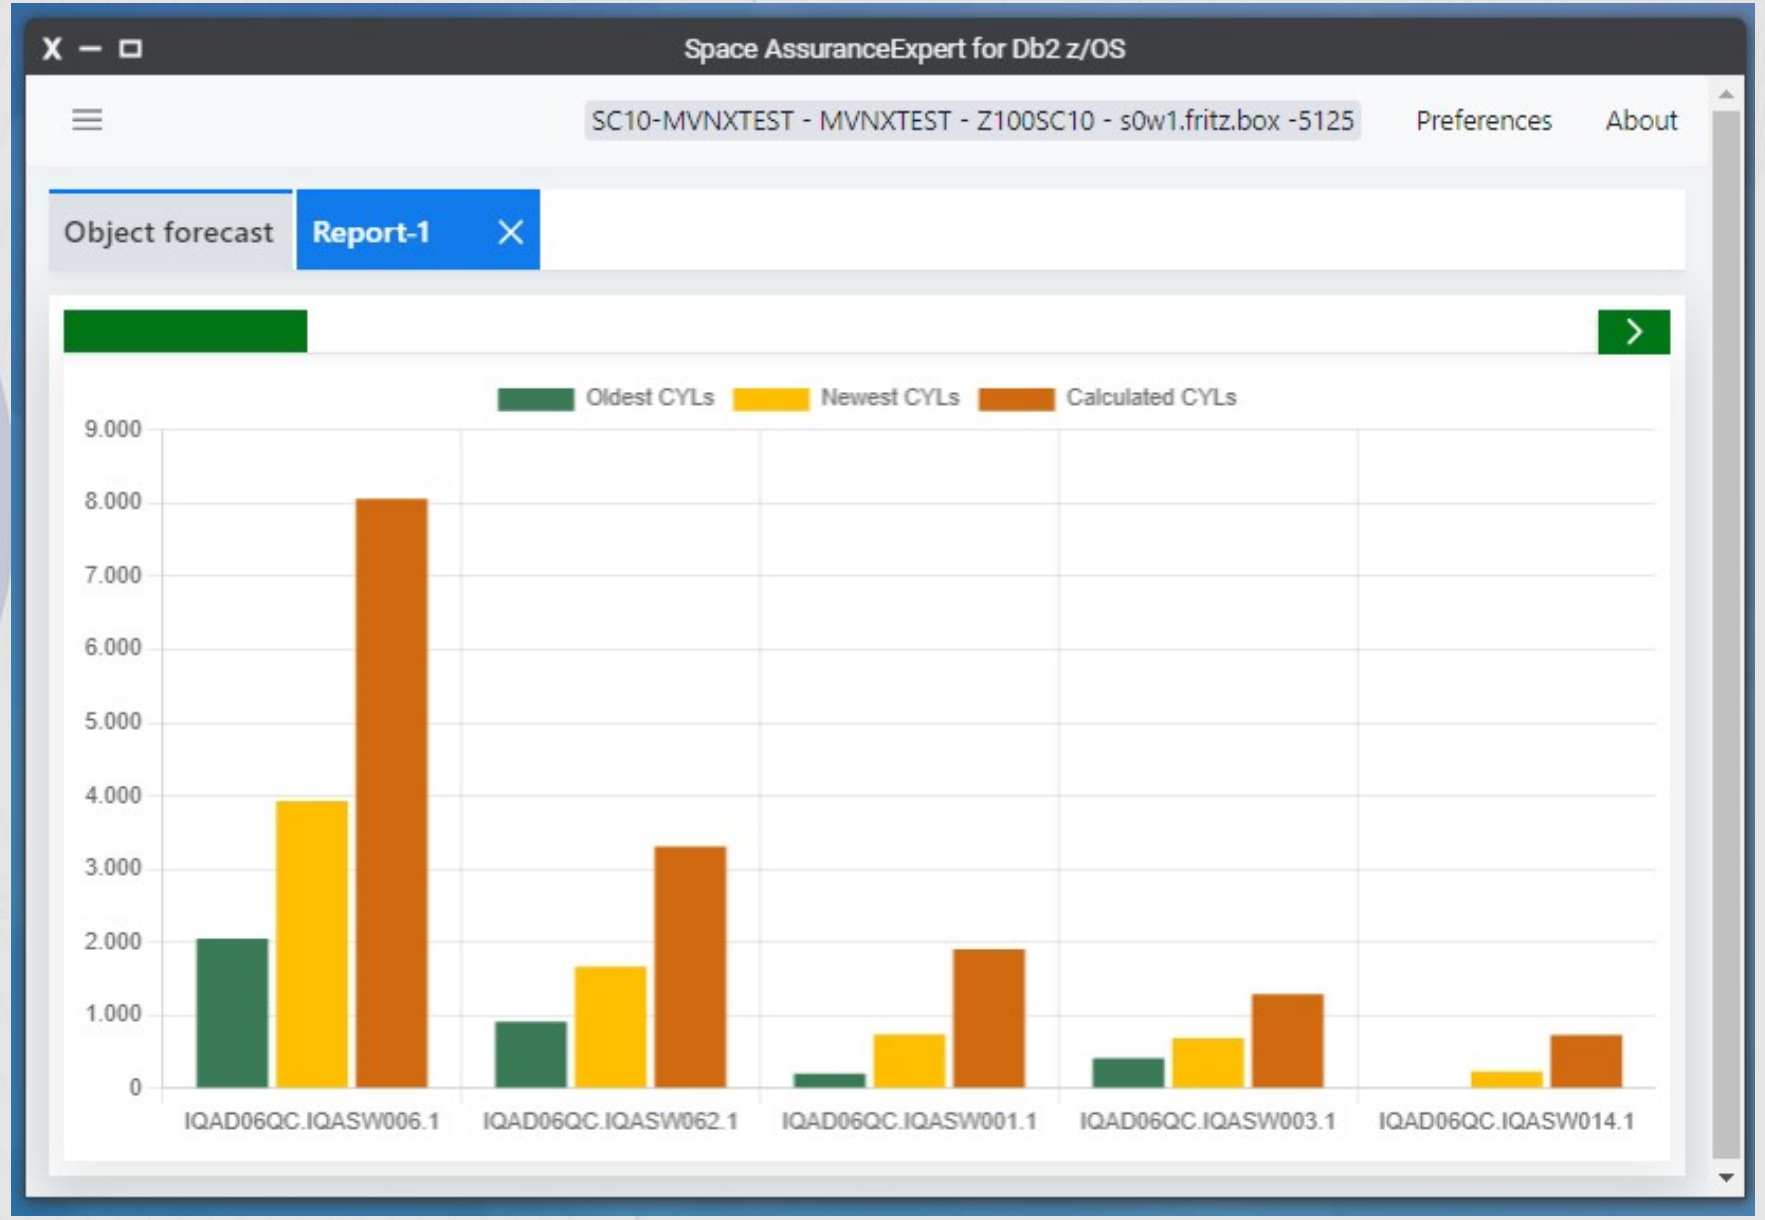
Space AssuranceExpert for Db2 z/OS
SC10-MVNXTEST - MVNXTEST - Z100SC10 - s0w1.fritz.box -5125
Preferences
About
Object forecast
Report-1
Oldest CYLs
Newest CYLs
Calculated CYLs
9.000
8.000
7.000
6.000
5.000
4.000
3.000
2.000
1.000
0
IQAD06QC.IQASW006.1
IQAD06QC.IQASW062.1
IQAD06QC.IQASW001.1
IQAD06QC.IQASW003.1
IQAD06QC.IQASW014.1

SEG and AI

We decided to then write an Anomaly Identification program (Clever name huh? ☺) that does indeed find the bad guys in all of the WLX data.

```
  <PROFILE>
    <ALIAS-CREATOR                          VALUE="IQA061QC" />
    <MIN-TOT-CPU                            VALUE="1000" />
    <THRS-AVG-CPU                           VALUE="25" />
    <LESS-MORE                              VALUE="BOTH" />
    <WTO-MSG                                VALUE="WTO" />
    <ALARM-RC                               VALUE="16" />
    <XPLNSUPPR-PRAUTHID                     VALUE="N" />
  </PROFILE>
  <PROCESS>
    <COMPARE-AGAINST                        VALUE="WEEK" />
    <COMPARE-HOURS                          VALUE="12" />
  </PROCESS>
</WLX>
```

SEG and AI

We decided to then write an Anomaly Identification program (Clever name huh? ☺) that does indeed find the bad guys in all of the WLX data.

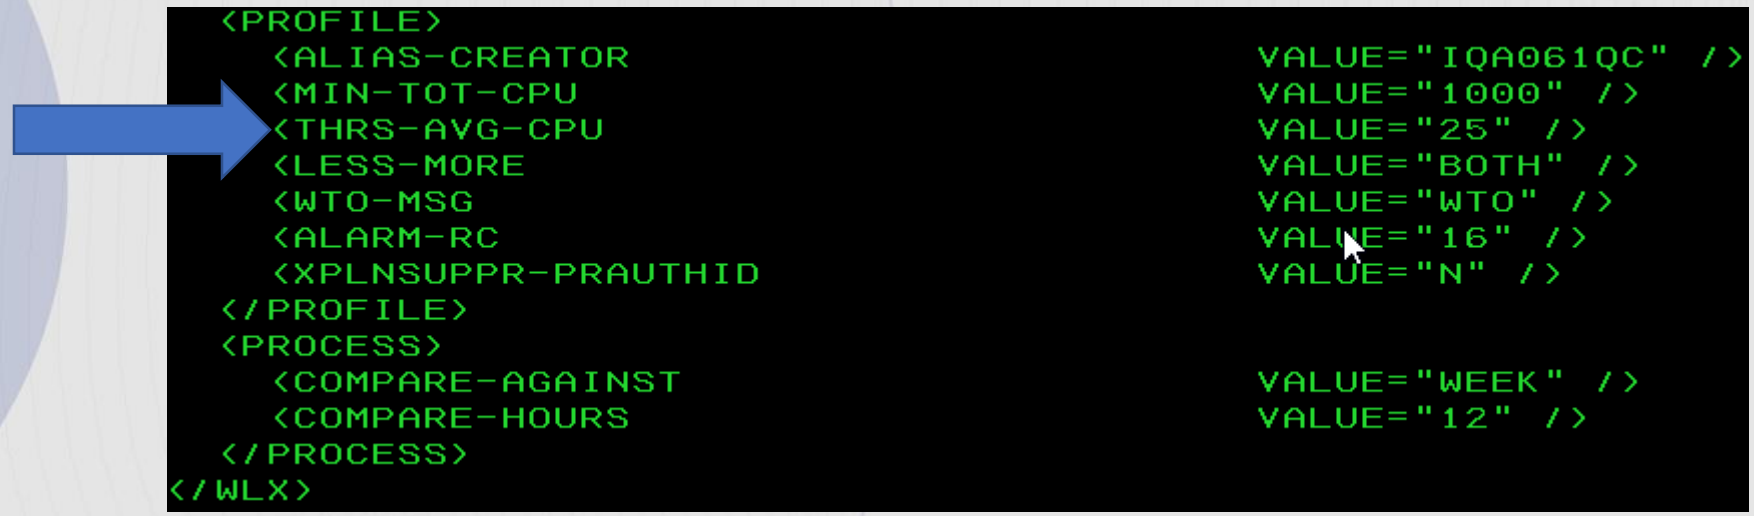

SEG and AI

We decided to then write an Anomaly Identification program (Clever name huh? ☺) that does indeed find the bad guys in all of the WLX data.

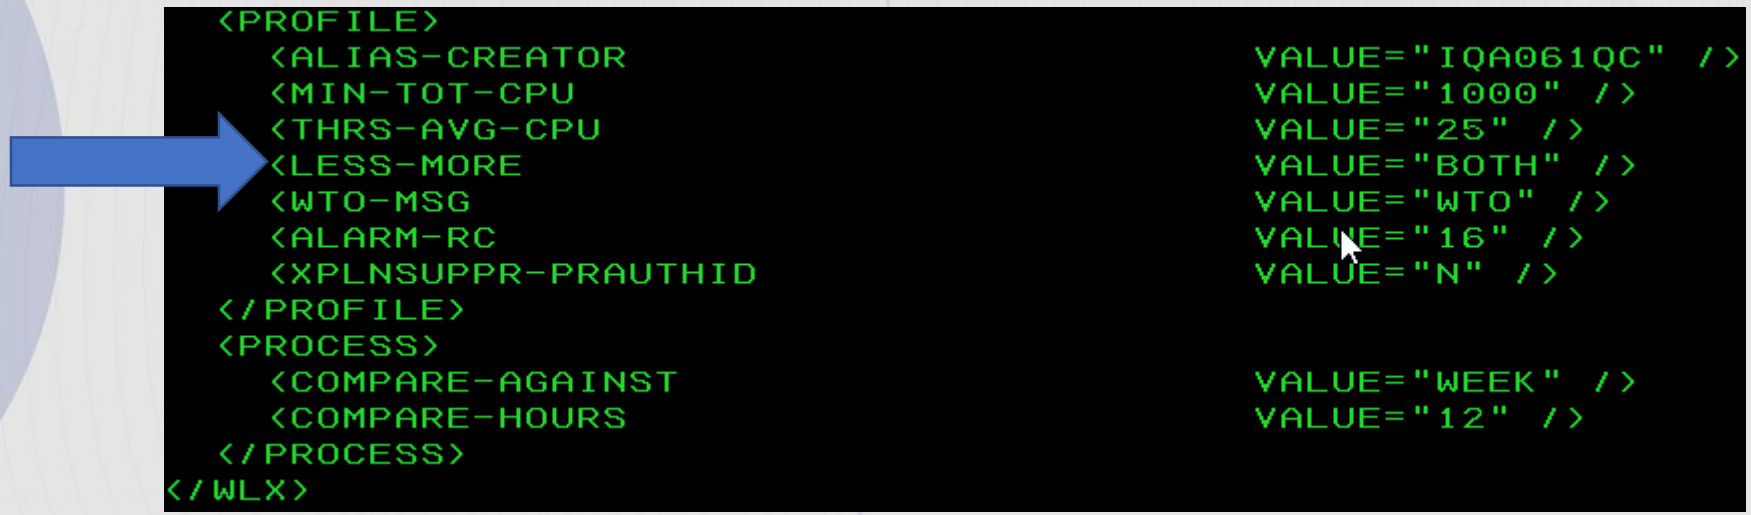

SEG and AI

We decided to then write an Anomaly Identification program (Clever name huh? ☺) that does indeed find the bad guys in all of the WLX data.

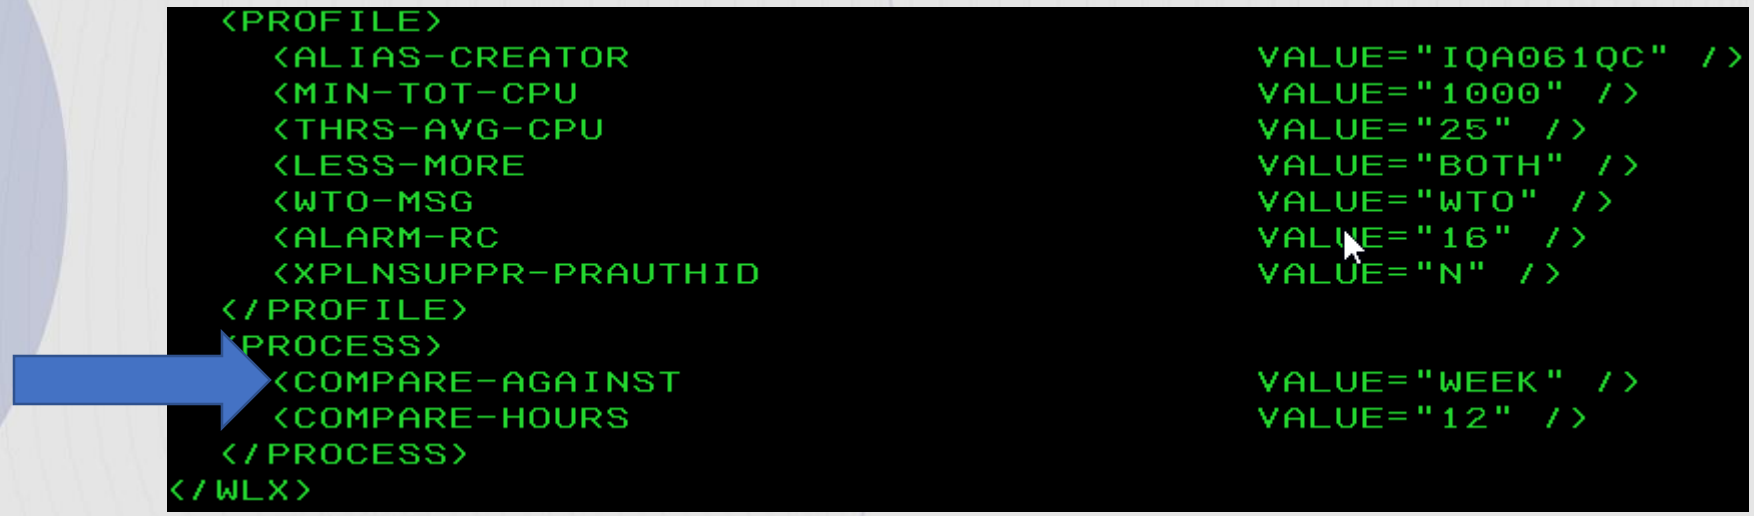

SEG and AI

Output is just in batch at this time but it is written to a Db2 table which we will soon integrate into our Zowe.

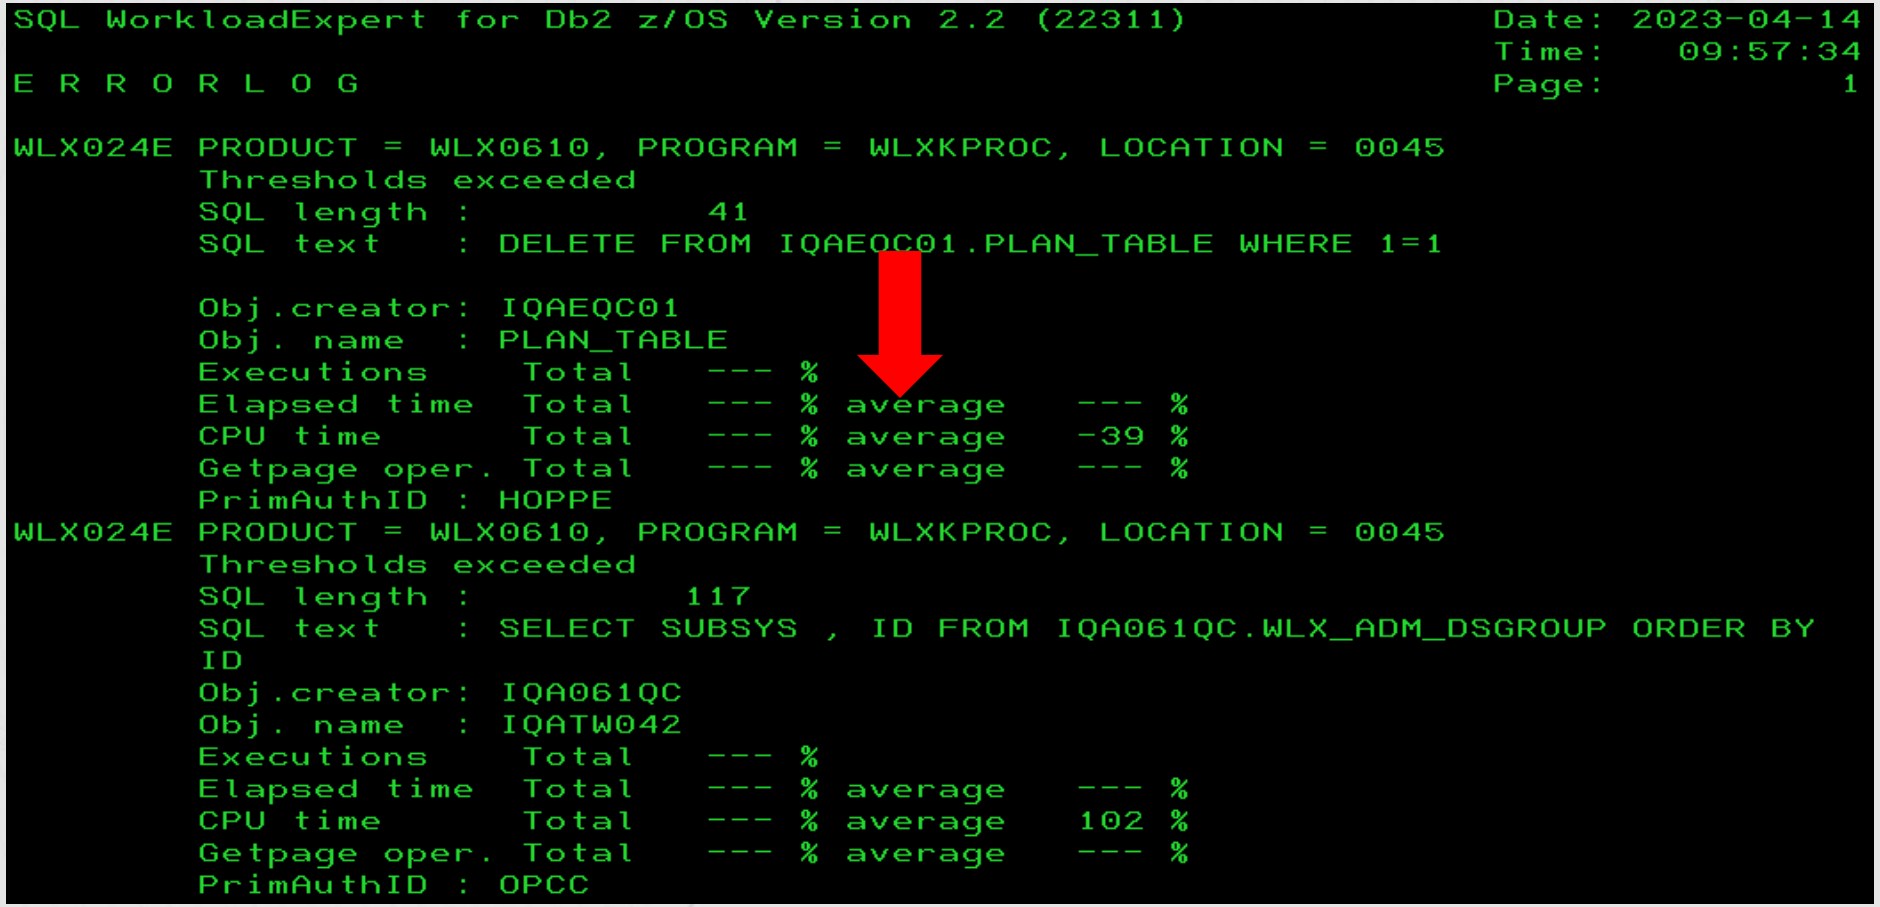

```
SQL WorkloadExpert for Db2 z/OS Version 2.2 (22311)                Date: 2023-04-14
                                                                   Time:    09:57:34
E R R O R L O G                                                    Page:           1

WLX024E PRODUCT = WLX0610, PROGRAM = WLXKPROC, LOCATION = 0045
        Thresholds exceeded
        SQL length :           41
        SQL text   : DELETE FROM IQAEQC01.PLAN_TABLE WHERE 1=1

        Obj.creator: IQAEQC01
        Obj. name  : PLAN_TABLE
        Executions    Total    --- %
        Elapsed time  Total    --- % average    --- %
        CPU time      Total    --- % average    -39 %
        Getpage oper. Total    --- % average    --- %
        PrimAuthID : HOPPE
WLX024E PRODUCT = WLX0610, PROGRAM = WLXKPROC, LOCATION = 0045
        Thresholds exceeded
        SQL length :          117
        SQL text   : SELECT SUBSYS , ID FROM IQA061QC.WLX_ADM_DSGROUP ORDER BY
        ID
        Obj.creator: IQA061QC
        Obj. name  : IQATW042
        Executions    Total    --- %
        Elapsed time  Total    --- % average    --- %
        CPU time      Total    --- % average    102 %
        Getpage oper. Total    --- % average    --- %
        PrimAuthID : OPCC
```

SEG and AI

Output is just in batch at this time but it is written to a Db2 table which we will soon integrate into our Zowe.

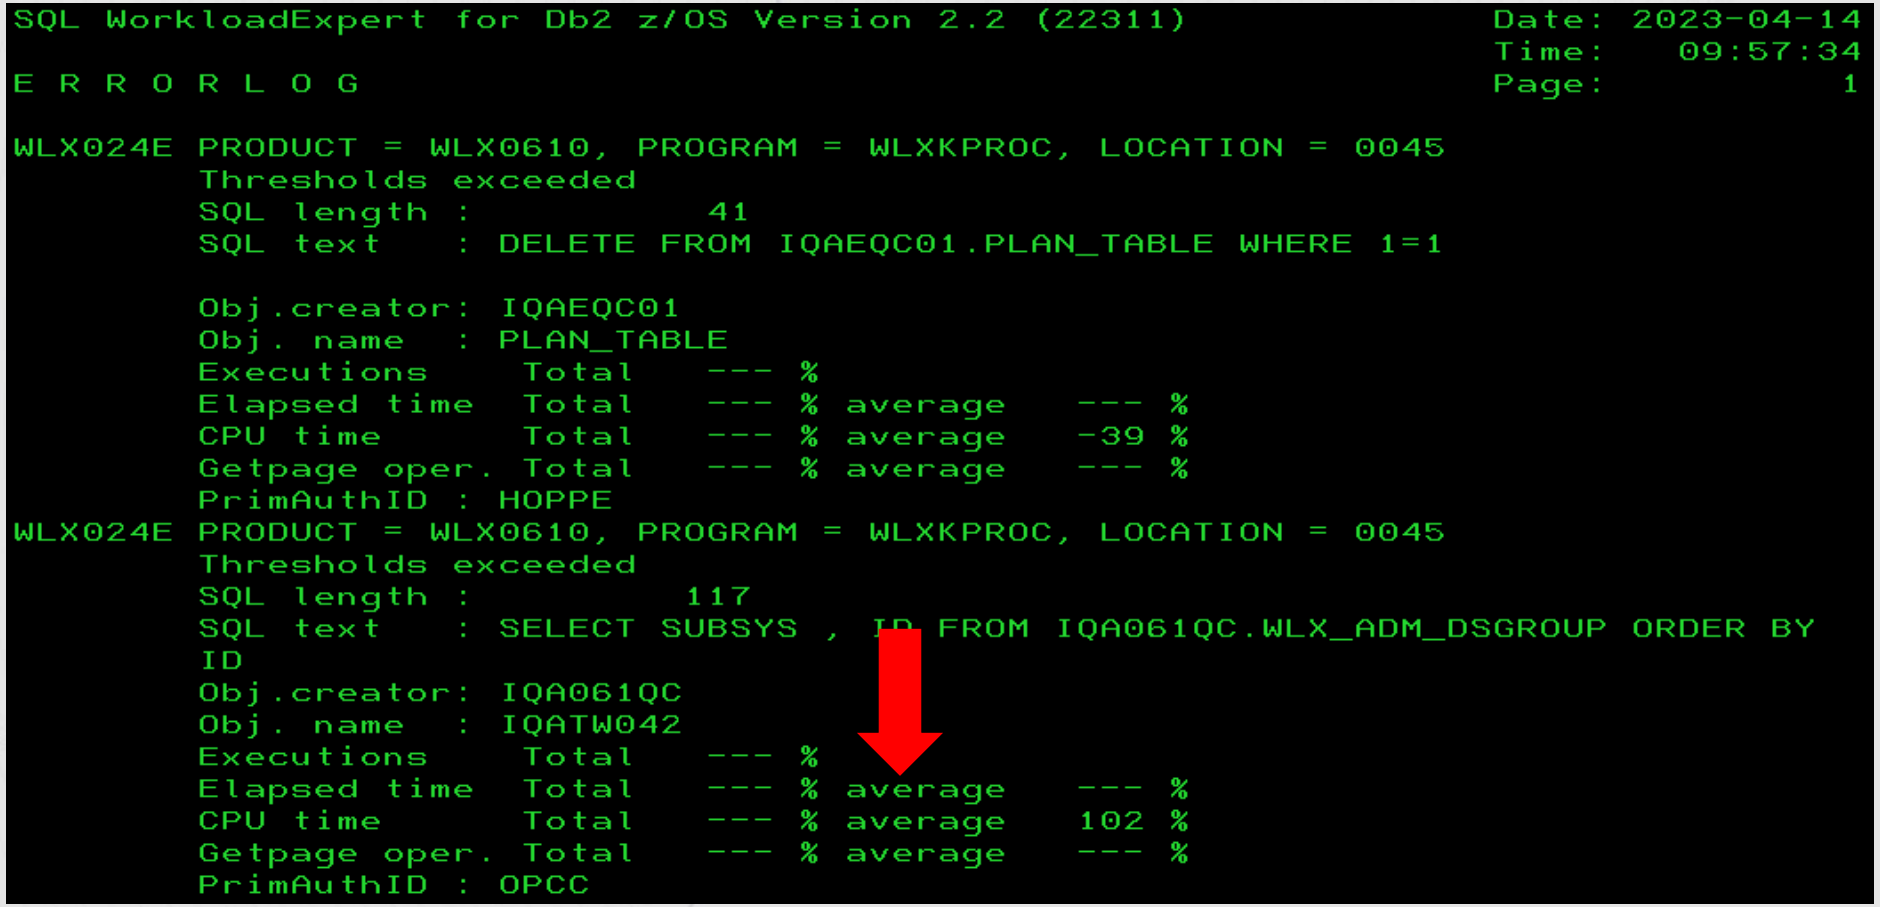

SEG and AI

Sadly, it does not use any SQL DI BiFs, however, it works a treat and does not require half your machine to run!

We are still running and testing the DI BiFs but, as of this time, we see no real world chance of actually ever using them.

This can, and hopefully will, change of course!

SEG and AI

So, after all that, where is AI in the real world of Db2 on z/OS?

SEG and AI

So, after all that, where is AI in the real world of Db2 on z/OS?

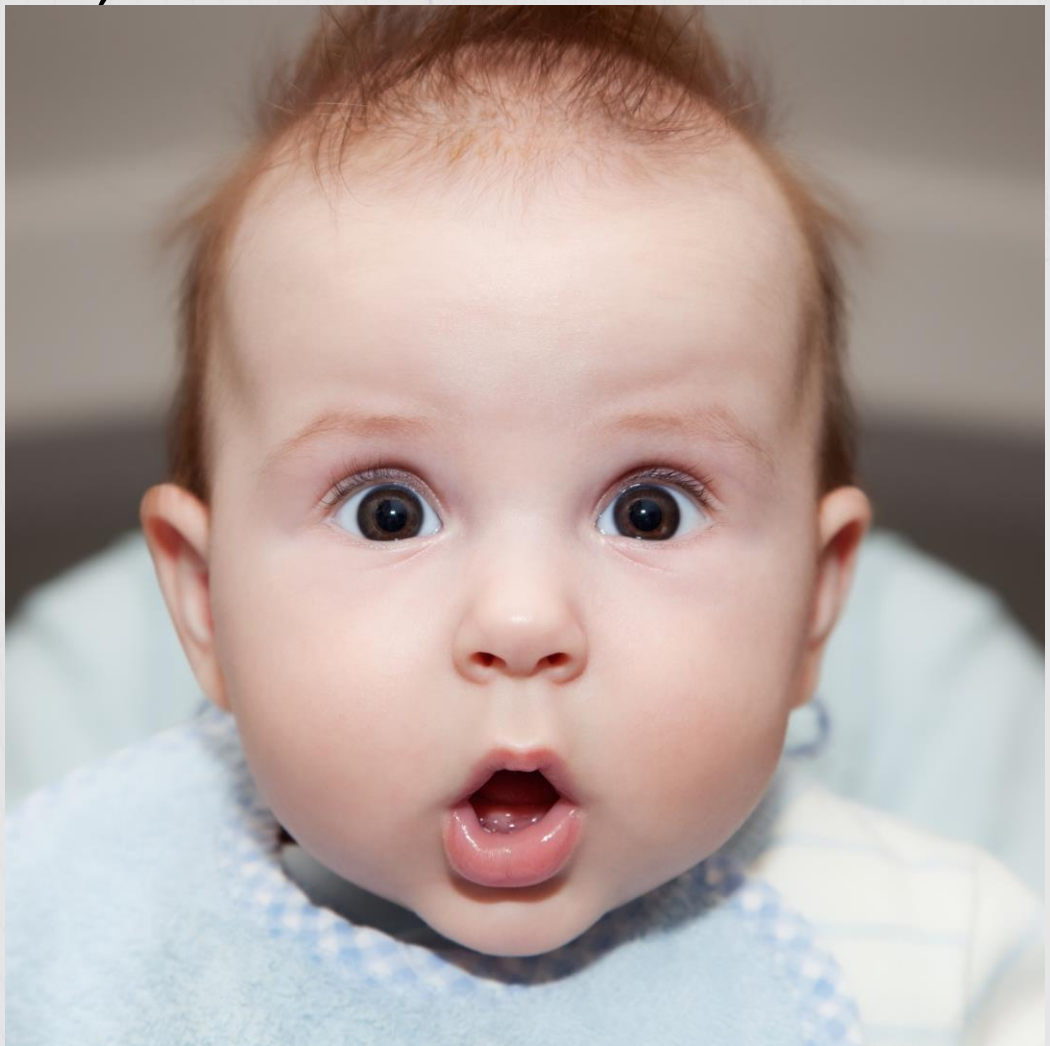

Agenda

- What is Artificial Intelligence?
- Db2 13 and Data Insights - AI for free?
- SEG and AI
- Q&A

Questions & Answers

IDUG

2023 EMEA Db2 Tech Conference

It’s AI Jim, but not as we know it!

Roy Boxwell

Software Engineering GmbH

r.boxwell@seg.de

*Session Code **E12***

Please fill out your session evaluation!