

Help! I'm overwhelmed with all these metrics in my Db2 monitor

Room 204B

Session Number:

Monday, February 24, 2020: 11:00 AM - 12:00 PM

Robert Andresen

Broadcom

robert.andresen@broadcom.com





ca[®]
technologies

**A Broadcom
Company**

Agenda

What are the causes of performance problems

How do you detect they are happening

What can you do to prevent them

Use what-if predictive analysis to see what effects a change might have before actually making the change

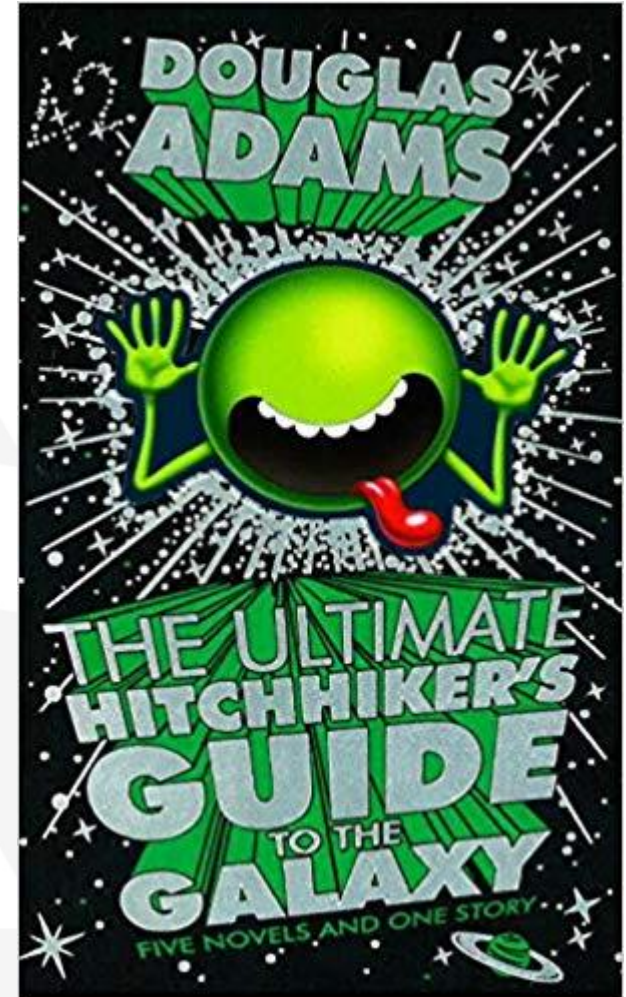
Purpose

My clients want an easy button
I see 1500+ metrics
What should these numbers be?

The answer:

“It depends”

Or 42 according to Douglas Adams



What are performance problems?

- What are we looking for?
- CPU time?
 - MULC are based on the max 4 hour rolling average over a month
 - Db2 application may or may not contribute
- Elapsed time
 - Slow application response
 - Holds resources longer
 - Can cause more contention
 - WLM implications – period 1, 2, 3
- Resources in short supply cause performance problems

When is CPU use a problem?

100% + latent demand

Causing a spike in the 4 hour rolling average

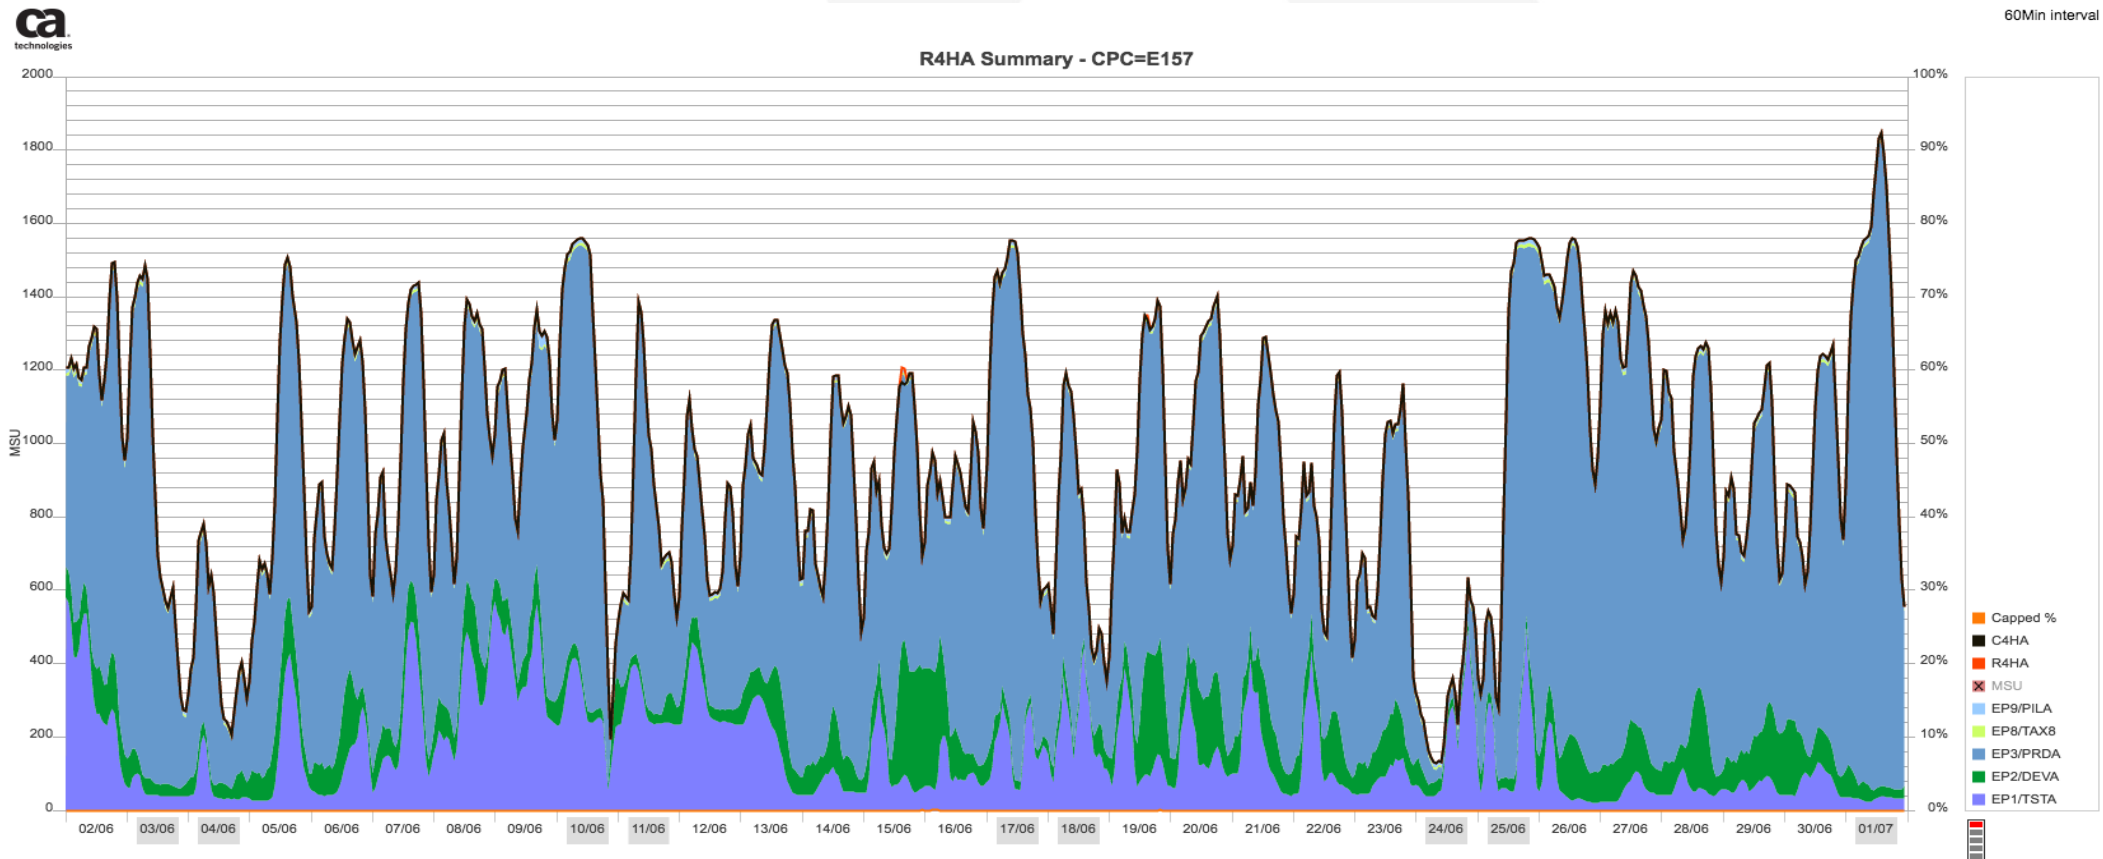
Driving capacity on demand

Failing to meet SLA

Does CPU really matter?

< 95% WLM not active

Contributing to monthly spike?



Where to start?

- Start high, drill down, learn root causes
- Elapsed Time – what are the major contributors?
- Wait times
 - Lock/latch
 - Data sharing
 - Unaccounted for time
- Always slow or only occasionally

Wait times

R/THRDRESP	Thread Response Time		Corr ID	DEMOCD	Wait Times (seconds)				
SORT N/A	Plan	DSNREXX	ID	DEMOCD	1	3	5	7	9
Auth ID ANDR016	Times in HH:MM:SS.TTTT	Evnts	% App Elpsd	% DB2 Elpsd					
Elapsed Time App	12:02.1481				.0	.0	.0	.0	.0
TCB Time Appl	1.9425	N/A	0.3						
Total DB2 Elapsed	27.1141	N/A	3.8						
Appl DB2 Elapsed	27.1141	948		100					
Total DB2 CPU	1.0670	N/A		3.9					
TCB Time DB2	1.0670	N/A		3.9					
Wait for DB2 I/O	5.2315	525		19.3					
Wait Local Lock	0.2494	14		0.9					
Wait Other Read	0.0056	1		0.0					
Latch Wait	0.0000	2		0.0					
Wt Data Shr Msgs	7.8656	554		29.0					
Wait Global Cont	0.0049	15		0.0					
Wt Log Write I/O	0.0198	8		0.1					
Wait DB2 Service	2.3645	62		8.7					
Update Commit	1.3161	34		4.9					
Open/Close/HSM	0.7989	6		2.9					
SYSLGRNG Updats	0.2150	14		0.8					
Other XUS Wait	0.0344	8		0.1					
Async IXL Reqs	0.0000	16		0.0					
Glbl Child L-Lk	0.0098	59		0.0					
Glbl Other L-Lk	1.0607	201		3.9					
Glbl Pgset P-Lk	0.0667	10		0.2					
Glbl Page P-Lk	0.7930	44		2.9					
Glbl Other P-Lk	0.0260	6		0.1					
Other DB2 Time	8.3468	N/A		30.8					

No matter how fast your machine is, they all wait at the same speed

-Peter Enrico

What can cause performance issues?

- The three R's
 - RUNSTATS
 - REBIND
 - REORG
- Indexes
- Bad SQL
- Contention
- Bufferpool issues



DB2 Address spaces

- System services: ssnmMSTR: system-related functions
- Database services: ssnmDBM1: large storage areas above the 2 GB bar
 - Bufferpools
 - EDM pool and Native Stored Procedures
- Distributed data facility: ssnmDIST provides support for remote requests
- IRLMPROC controls DB2 locking
- WLM address spaces for stored procedures
- User address spaces
 - TSO
 - Batch
 - CICS®
 - IMS™ dependent region
 - IMS control region
 - WebSphere®

Measure the effect of changes

- Limit the number of changes
- Repeatable workload
- Starting state of the subsystem
 - Bufferpool contents
 - Back out updates/deletes/inserts



An example

CREATE INDEX

Added INCLUDE with 2 columns

Specifies additional columns to append to the set of index key columns

Not used to enforce uniqueness.

Might improve performance for some queries using index only access.

Metric	Before	After	Delta
AVG DB2 Response Time	0.001065	0.000526	-50.61%
AVG DB2 CPU Usage	0.000281	0.000105	-62.63%
AVG GETPAGES	15.88	5.26	-66.88%

When should we monitor performance?

- The earlier the better
- Mistakes early in a project lifecycle are more expensive to fix
- Make performance tools available to developers/agile teams



Avoid moving bad SQL into Production

- Programmers should explain their packages
 - The earlier a problem is detected the easier they are to fix
 - Copy production statistics to test subsystems
- Run compare analysis before anything gets moved into production
 - Proactively avoid performance problems
 - Before migration weekend, when you have time to fix any problems
- Add a cost analysis step to change management
 - Set return codes
 - Stop the job if non-zero return
 - Last line of defense

SQL Incompatibilities

IFCID 376

Migrating to Db2 12

Applying maintenance

Activating function levels

https://www.ibm.com/support/knowledgecenter/SSEPEK_12.0.0/wnew/src/tpc/db2z_relincompat.html

```
SELECT 'BIND PACKAGE(DUMMYCOL) COPY(' ||  
      COLLID || '.' || NAME || ') ' ||  
      CASE WHEN(VERSION <> '')  
            THEN 'COPYVER(' || VERSION || ') '  
            ELSE '' END ||  
      'EXPLAIN(ONLY) '  
FROM SYSIBM.SYSPACKSTMT  
WHERE STATEMENT LIKE '%SELECT% INTO % UNION %'
```




WHAT IS HAPPENING NOW?

Current and Near Recent Performance

- Subsystem
- Threads
- SQL



Some Interesting IFCIDs

1,2,3	Class 1,2 & 3 response times by plan
7,8	Packages
149, 150	Locking
196	Deadlocks
225	Storage Summary
316, 317	Dynamic SQL Stats

System statistics

001 STATS1 – z/OS System Statistics

002 STATS2 – Buffer Statistics

Subsystem monitor

- Thread metrics - IFCID 0003
 - Response time graph
 - Active and historical threads
 - Class 1, 2 and 3 accounting times
 - Not accounted for time
 - Packages (collection, SP/UDF, trigger)
 - Active and historical threads
 - Breaks out resource use by object

003 ACCT record:

DB2 accounting records

SQL statistics

Subtotals of DB2 above counters

Buffer pool statistics

Lock and resource limiter statistics

Distributed unit of work statistics

MVS account code and DDF information

IFI statistics

Package statistics

Package buffer pool statistics

Package locking and resource limit statistics

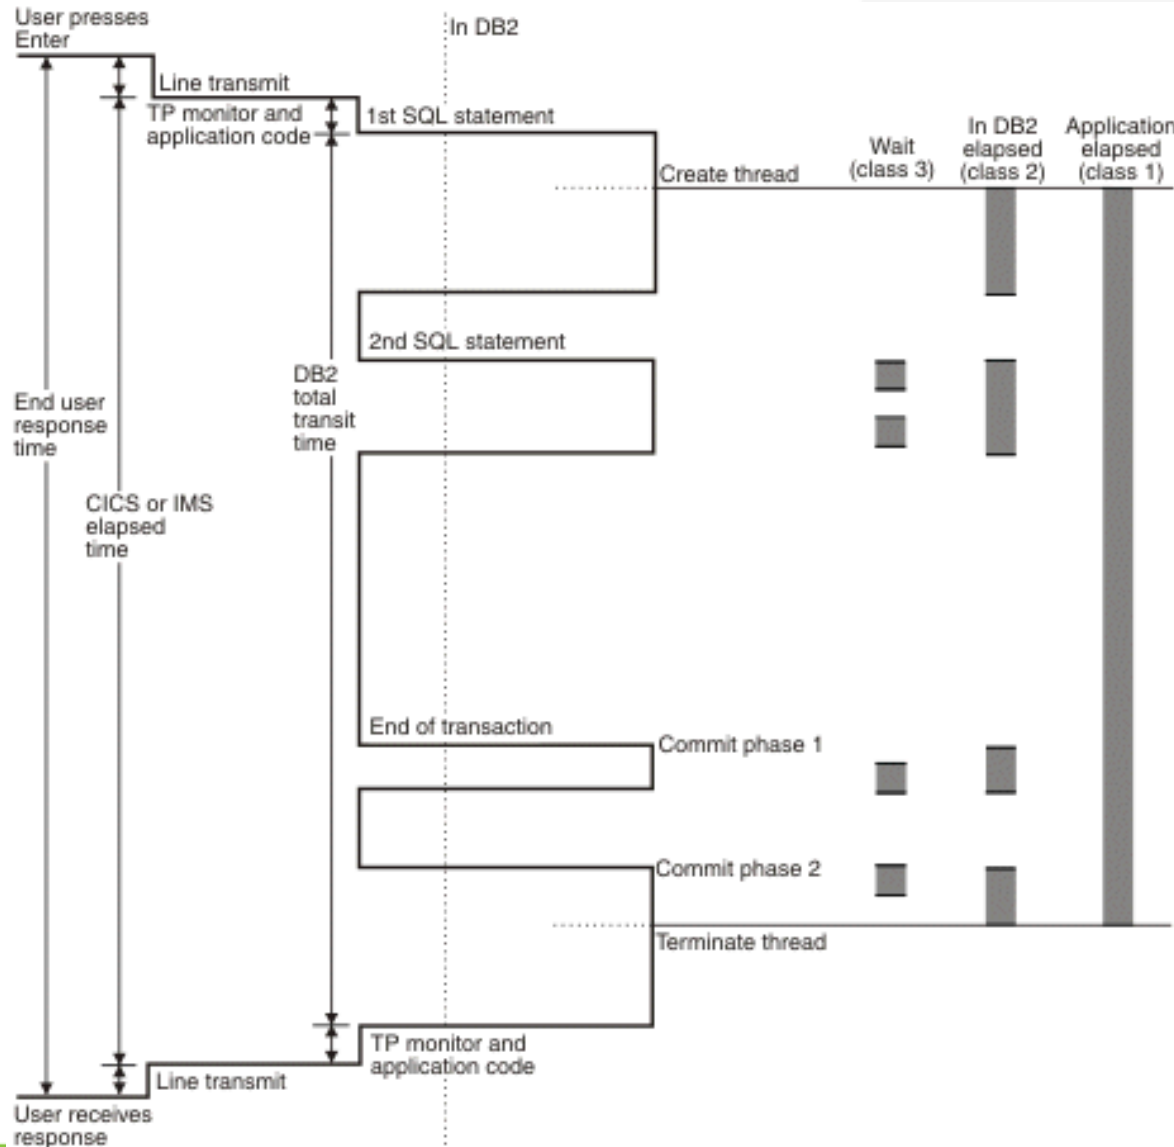
Package SQL statistics

Group buffer pool statistics

Accounting data for data sharing

Global locking statistics

Db2 response time



- Class 1 elapsed time includes application and Db2 elapsed times
- Class 2 elapsed is the elapsed time in Db2.
- Class 3 is elapsed wait time in Db2

Source:

http://www.ibm.com/support/knowledgecenter/SSEP_EK_10.0.0/perf/src/tpc/db2z_responsetimereports.html

Accumulated Thread statistics

CICS

- ACCOUNTREC({NONE|TASK|TXID|UOW})
- Specifies the minimum amount of DB2® accounting required for transactions using this DB2 entry
https://www.ibm.com/support/knowledgecenter/en/SSGMCP_4.2.0/com.ibm.cics.ts.resourcedefinition.doc/resources/db2entry/dfha4_thrdopattr.html

DDF and RRSF

- ACCUMACC(NO,integer):
https://www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/inst/src/tpc/db2z_ipf_accumacc.html
- ACCUMUID determines the aggregation fields that are to be used for DDF and RRSF accounting rollup.
https://www.ibm.com/support/knowledgecenter/en/SSEPEK_11.0.0/inst/src/tpc/db2z_ipf_accumuid.html

Subsystem monitor

- Subsystem history metrics
- Bufferpool Page Residency Time
- See which threads were running during an interval the system was performing poorly
 - Accounting summary by interval
 - Go to thread list for a specific connection type for an interval
- Thread history
 - Search for threads that exhibit specific performance metrics
 - Any class 1, 2 or 3 accounting times
 - Specific bufferpool use
 - Any metric in IFCID 3

Bufferpool option recommendations

Autosize: In DB2 11, the AUTOSIZE option can limit the range within VPSIZEMIN and VPSIZEMAX.

Pagefix

Framesize

$VPSEQT = \text{MAXIMUM BUFFERS ON SLRU (HWM)} / VPSIZE$

Random page residency time is calculated as the maximum of:

- $VPSIZE / \text{total pages read per second (including pages prefetched)}$
- $VPSIZE * (1 - VPSEQT/100) / \text{random synchronous IO/sec}$
- More detail in DB2 9 for z/OS: Buffer Pool Monitoring and Tuning, REDP-4604 by Mike Bracey

With Db2 12 the EDM Pool may need to be increased

Identify Buffer Pool problems

IBM DB2 Bufferpool Suggestion:

- **BP0 – DB2 Catalog**
- **BP1 – Small-sized Reference Tablespaces**
- **BP2 – Small-sized Reference Indexspaces**
- BP3 to BP6, BP8 to BP9 – expansion/isolation for performance bottlenecks in BP1 and BP2
- **BP7 – Sort DSNDB07**
- BP10 to BP19 – Tablespace buffer pools
- BP10 – Medium-sized Sequential Access
- BP11 – Medium-sized Random Access
- BP12 – Large-sized Sequential Access
- BP13 – Large-sized Random Access
- BP14 to BP19 – expansion/isolation for performance bottlenecks
- BP20 to BP29 – Indexspace buffer pools
- BP20 – Medium-sized Sequential Access
- BP21 – Medium-sized Random Access
- BP22 – Large-sized Sequential Access
- BP23 – Large-sized Random Access
- BP24 to BP29 – expansion/isolation for performance bottlenecks

Alerting on current problems

- Automated alerting/corrective action
 - Automate close to the problem
 - Review Thresholds
 - eMail/SNMP Notifications
 - On call support
 - Open a trouble ticket
 - Post on enterprise console
 - OPS/MVS API or WTO
 - Cancel threads
 - Start traces
 - z/OS commands

Get the red out

If your exception monitor always shows twelve alerts when there is no problem – no one will notice the thirteenth

Turn it off or tune the thresholds

Get rid of test alerts

Take what you learn from problem resolution to set up new alerts

SQL Performance

Collect performance data for each SQL statement that runs

Standard collection aggregates the metrics over some interval

Summaries by Plan, Package, SQL Statement

Optionally by User, Correlation ID, Connection, Work station values, etc.

Exceptions can be set up to capture metrics for a single instance if a threshold is breached

Homegrown SQL monitoring

- IFCIDs:
 - 316: Reports on the contents of the prepared SQL statement cache.
 - 317: Provides the SQL statement text for an individual SQL statement in the prepared SQL statement cache.
 - 318: Generates no external data, but causes DB2 to collect additional statistics that are reported in the IFCID 316 record. Start a request with this record (DYNSTATS) to provide the additional statistics.
 - 401: Static statements that are tracked in the EDM pool. Use the statement identifier in this record to query the catalog to identify its statement text. IFCID 400 must be enabled to collect data.
- Monitor class 29
 - Turns on IFCIDS 316, 318. 400 and 401

Object maintenance

- Utilities
 - REORG
 - RUNSTATS/REBIND
 - Image Copy
- Object Selection profiles
 - Uses RTS
 - Include/exclude databases, tablespaces
- Prioritize Object Profiles
 - Prioritize objects in action JCL so that the most critical objects will be processed first
 - Based on ranking criteria out of the box
 - You can customize

- Get recommendations for when to reorganize, image copy, or update statistics for table spaces or index spaces
- Identify when a data set has exceeded a specified threshold for the number of extents that it occupies.
- Identify whether objects are in restricted states
- Externalize in-memory statistics to the RTS tables before calling the stored procedure.
 - STATSINT subsystem parameter.
 - ACCESS DATABASE command and specify the MODE(STATS) option
- DSNACCOX runs with isolation UR to avoid lock contention
- If the real-time statistics tables contain information for only a small percentage of your Db2 subsystem, the recommendations that DSNACCOX makes might not be accurate



FINDING CAUSES OF PAST PERFORMANCE PROBLEMS

"THOSE WHO CANNOT REMEMBER THE PAST ARE CONDEMNED TO REPEAT IT."

-GEORGE SANTAYANA

Finding causes of past performance problems

- Looking for problems after the fact
- The more history the better
- Similar to the concept of a data lake

Db2 subsystem monitor

- Online history
 - System history – default size holds several months
 - Thread history – default size may fill up quickly
 - Increase to several days to a week if you can afford the space
 - Online search is very efficient
- Files
 - Offload history
 - By percent or by time
 - Use date time system parameters as part of the file name
 - These files can be brought back online
 - User started requests
 - Traces can be saved to files

Batch capabilities

- PM style reports
- Batch versions of many online panels
- TOP(x) on CPU, elapsed, etc
- Archive extract to load to performance warehouse
- CSV output format
 - Customize new or modify existing requests
 - Limit the fields from the IFCID to a manageable few
 - As you learn what's important you can add additional metrics
 - Easy to load into Excel and graph, compare,

Performance Warehouse

- Database of performance metrics derived from IFCIDS
- Application
 - By thread
 - By Package
 - DDF
- Subsystem

Vendor supplied Thread monitor DB2 tables

• CA

- APPL_BP_DETAIL
- APPL_DDF_DETAIL
- APPL_GBP_DETAIL
- APPLICATION_DETAIL
- APPL_PGM_DETAIL
- APPL_BP_DAILY
- APPL_DDF_DAILY
- APPL_GBP_DAILY
- APPLICATION_DAILY
- APPL_PGM_DAILY
- SUBSYS_BP_DETAIL
- SUBSYS_DDF_DETAIL
- SUBSYS_GBP_DETAIL
- SUBSYSTEM_DETAIL

• IBM

- DB2PMFACCT_BUFFER
- DB2PMFACCT_DDF
- DB2PMFACCT_GBUFFER
- DB2PMFACCT_GENERAL
- DB2PMSACCT_BUFFER
- DB2PMSACCT_DDF
- DB2PMSACCT_GBUFFER
- DB2PMSACCT_GENERAL
- DB2PMSACCT_PROGRAM
- DB2PM_STAT_BUFFER
- DB2PM_STAT_DDF
- DB2PM_STAT_GBUFFER
- DB2PM_STAT_GENERAL

Vendor Supplied SQL Monitor DB2 Tables

- CA
 - PDT_DYNAMREQ
 - PDT_DYNAMTXT
 - PDT_ERRORTXT
 - PDT_ERRORVAR
 - PDT_HOSTVARS
 - PDT_OBJECT
 - PDT_SQLERROR
 - PDT_STANDARD
 - PDT_STANTEXT
 - PSA_BP
 - PSA_DB
 - PSA_VOL
- IBM
 - CQMTOOLS database

Performance Warehouse tips

- Extracting everything can use a lot of space
 - Limit by application
 - Limit by kind of data
 - Detail vs. Summary
 - Application vs. System
 - Only load data when you want to research a problem
 - Load everything and put it into the IDAA
- Usually have one warehouse
 - Extract data from multiple sources



PREDICTING THE FUTURE

“I WAS NOT PREDICTING THE FUTURE, I WAS TRYING TO PREVENT IT.”
— RAY BRADBURY

Predicting the Future Preventing the Future

- What-if analysis
- Finding trends and correlations



Bufferpool simulation

- -ALTER BUFFERPOOL (BPx) SPSIZE(y)
- Run an application
- -Display Bufferpool
- -ALTER BUFFERPOOL (BPx) SPSIZE(0)
- Back out changes
- Run the same application
- -Display Bufferpool

Caveats:

- 2 or 3 bufferpools max
- Turn off when done testing
- Db2 overhead

Virtual Indexes

- *Table name*
 - virtually create an index definition
 - virtually drop an existing catalog index
- Rebind Compare Analysis
 - Access path, cost changes
 - Make part of change control
- Statistics Migration
 - Migrate statistics from one subsystem to another
 - Create statistics for an environment when production statistics do not exist for modeling.
 - Automatically calculate regression and future statistics.
 - Propagate new cardinality to all related objects (such as indexes and tablespaces).

Virtual indexes

- Allows you to simulate the impact of creating or dropping an index
- Only affects EXPLAIN; not BIND/REBIND or dynamic SQL
- Considered by the DB2 Optimizer with existing indexes

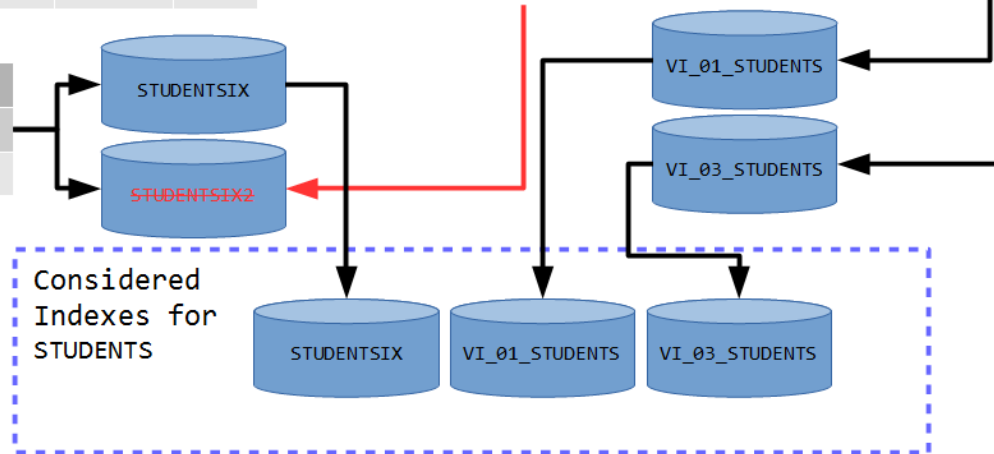
DSN_VIRTUAL_INDEXES

TBNAME	IXNAME	ENABLE	MODE
STUDENTS	VI_01_STUDENTS	Y	C
STUDENTS	VI_02_STUDENTS	N	C
STUDENTS	VI_03_STUDENTS	Y	C
STUDENTS	STUDENTSIX2	Y	D

→ CREATE INDEX VI_01_STUDENTS ON STUDENTS...
→ CREATE INDEX VI_03_STUDENTS ON STUDENTS...
→ DROP INDEX STUDENTSIX2;

SYSINDEXES

TBNAME	IXNAME
STUDENTS	STUDENTSIX
STUDENTS	STUDENTSIX2



Analytics

- Data science
- Learn from your historical data
- Correlations
- Normal behavior

What is Normal?



Deeper Analysis may be needed

- Problem: MAXDBAT is being hit
- Simple solution: Increase MAXDBAT
- But why are you hitting MAXDBAT?
 - Running more threads?
 - More remote locations?
 - Longer running threads?
 - More contention?
- Increasing MAXDBAT if the real problem is contention may increase contention
- Problems are symptoms
- Find and fix the underlying cause

Using Analytics and Machine Learning

- Free
 - Excel Analytics ToolPak
 - R
 - Python
- Commercial
 - Splunk
 - CA Mainframe Operational Intelligence

Excel

- Built-in statistics functions

- Formulas
- Subtotals

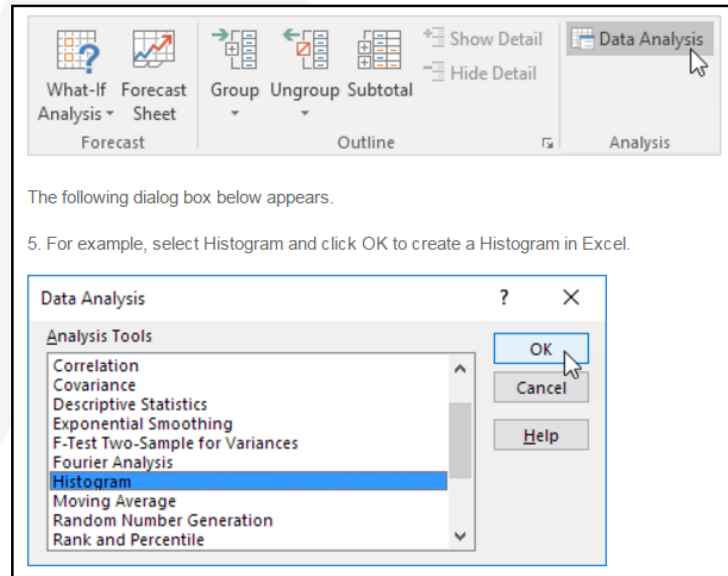
- Data Analysis ToolPak

- complex statistical analyses

- <https://support.office.com/en-us/article/Load-the-Analysis-ToolPak-6a63e598-cd6d-42e3-9317-6b40ba1a66b4>

- Graphics

- Chart Wizard
- Save charts as template
- Create a histogram tutorial: <https://support.office.com/en-us/article/Create-a-histogram-85680173-064b-4024-b39d-80f17ff2f4e8?ui=en-US&rs=en-US&ad=US>



Get the correlation value

APPLSUM.xlsx - Excel Andresen, Robert D

File Home Insert Page Layout Formulas Data Review View Tell me what you want to do

Clipboard Font Alignment Number Styles Cells Editing WebEx

SUM $\text{=CORREL(L1229:L1513,N1229:N1513)}$

	I	J	K	L	M	N	O	P	Q	R	S	T	U
339	DSNREXX	DB2CALL	1	24080556	353811	143372	18867872	0	0		Correlations		
341	DSNREXX	DB2CALL	1	61998474	50650	2191	61524876	0	0		Elapsed Time vs. I/O time		
344	DSNREXX	DB2CALL	1	204968216	22943624	1539729	1514361	0	0		:N1513)		
345	DSNREXX	DB2CALL	1	1299984	25560	5223	1114887	0	0				
346	DSNREXX	DB2CALL	1	241399106	24201817	1392802	4361135	0	0				
1184	DSNREXX	DB2CALL	1	1043031869	84326508	9030071	224	0	0				
1229	DSNREXX	DB2CALL	1	1370134809	80518678	8120078	629	0	0				
1390	DSNREXX	DB2CALL	1	200403667	20995762	629418	1371231	0	0				
1393	DSNREXX	DB2CALL	1	213684479	20996250	569622	413	0	0				
1405	DSNREXX	DB2CALL	1	190768102	20815721	428923	2574284	0	0				
1407	DSNREXX	DB2CALL	1	503033121	20714405	387415	313760931	0	0				
1408	DSNREXX	DB2CALL	2	1092466817	39289306	707184	713054280	0	0				
1409	DSNREXX	DB2CALL	3	1342919365	59977019	904452	804695107	0	0				
1410	DSNREXX	DB2CALL	2	360537599	23824408	15591227	149741501	0	0				
1411	DSNREXX	DB2CALL	2	878051814	39415791	602689	512552454	0	0				
1413	DSNREXX	DB2CALL	1	196693623	20453262	358920	2486645	0	0				
1414	DSNREXX	DB2CALL	3	984063143	60041005	772515	492646431	0	0				
1416	DSNREXX	DB2CALL	2	916440208	39725142	402564	609106104	0	0				
1417	DSNREXX	DB2CALL	1	419747373	20668402	209082	265531914	0	0				

Correlation Raw Data Charts

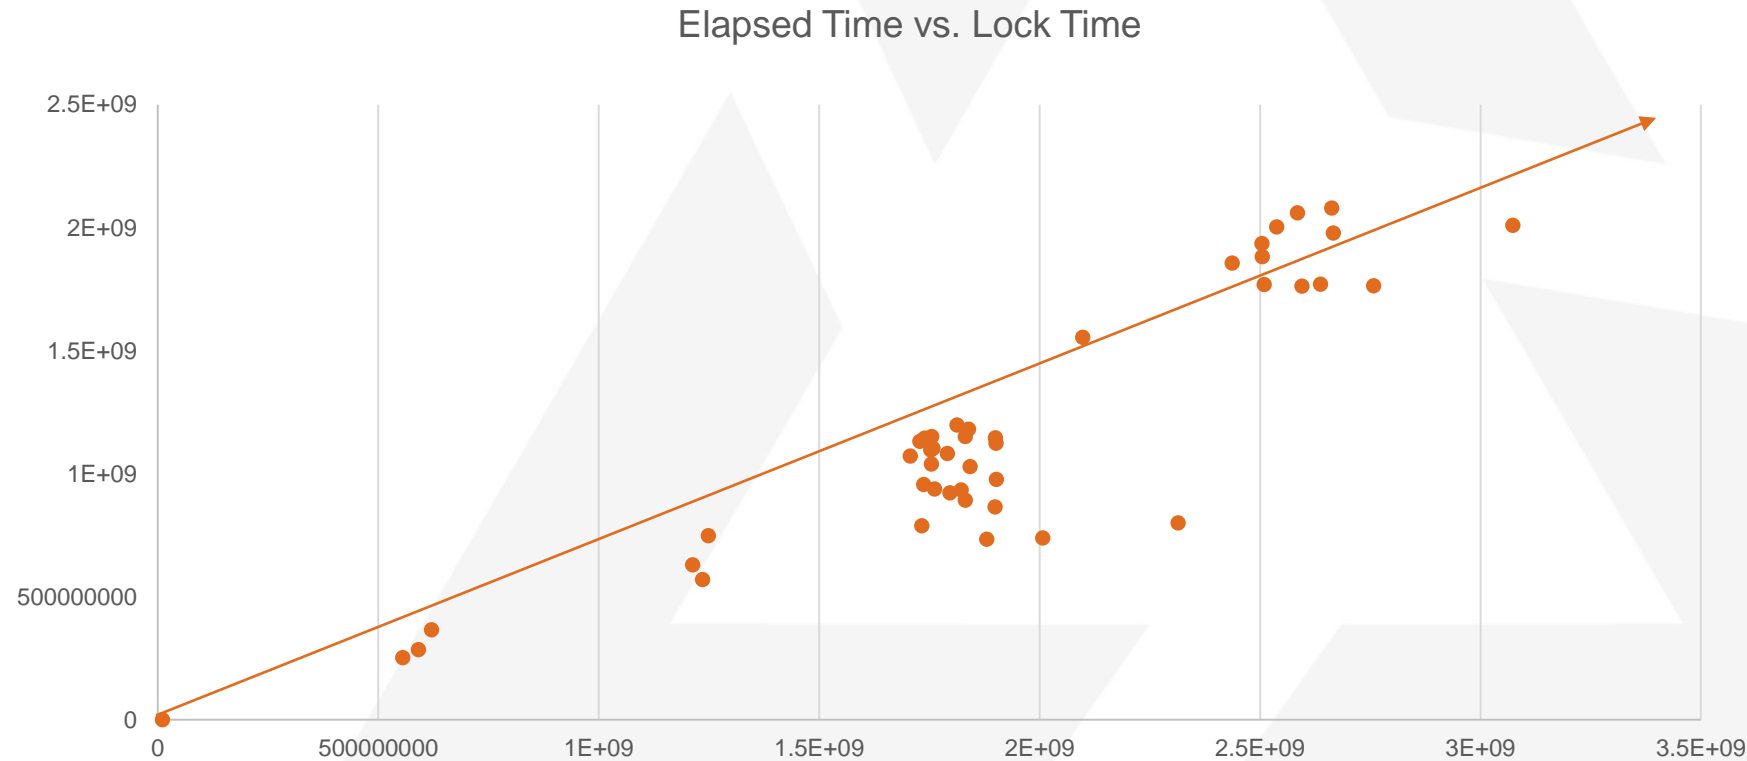
Edit 100%

Excel provides built-in functions
 $\text{=CORREL(L1229:L1513,N1229:N1513)}$
Result here is 0.541086

Scatter Plot

Highlight the two sets of values you want to see correlation

Insert -> Charts -> Scatter plot



R

- R is a free software environment for statistical computing and graphics. It compiles and runs on a wide variety of UNIX platforms, Windows and MacOS.
- The simplest Google search: R
- Traditionally been used in academics and research
- Documentation and samples on the web



Python

- Interpreted, object-oriented, high-level programming language
- Easy to learn
- Quick to develop
- Runs on Linux, MacOS and Windows
- Gaining popularity with data scientists



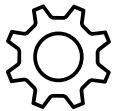
Splunk: Critical Insights

SOURCES



Order Processing

Customer ID Order ID Product ID
ORDER, 2016-05-21T14:04:12.484, 10098213, 569281734, 67.17.10.12, 43CD1A7B8322, SA-2100



Middleware Error

MAY 21 14:04:12.996 wl-01.acme.com Order 569281734 failed for customer 10098213.
Exception follows: weblogic.jdbc.extensions.C (Order ID DeadSQLException Customer ID
weblogic.common.resourcepool.ResourceDeadException: Could not create pool connection. The
DBMS driver exception was: [BEA][Oracle JDBC Driver] Error establishing socket to host and port:
ACMEDB-01:1521. Reason: Connection refused



Care IVR

05/21 16:33:11.238 [CONNEVENT] Ext 1207130 (0192033): Event 20111, CTI Num:ServID:Type
0:19:9, App 0, ANI T7998#1, DNIS 5555685981, SerID 40489a07-7f6e-4251-801a-
Time waiting on hold k T451.16
05/21 16:33:11.242 [SCREENPOPEVENT] SerID 40489a07-7f6e-4251-801a-13ae51a6d092
CUSTID 10098213 Customer ID
05/21 16:37:49.732 [DISCEVENT] SerID 40489a07-7f6e-4251-801a-13ae51a6d092



Twitter

{actor:{displayName: "Go Boys!!",followersCount:1366,friendsCount:789,link:
http://dallascowboys.com/,location:{displa Twitter ID llas, TX",objectType:"place Customers Tweet
objectType:"person",preferredUsername:"B0ysF@n80",statusesCount:6072},body: "Can't buy
this device from @ACME. Site doesn't work! Called, gave up on waiting for them to answer! RT if
you hate @ACME!!",objectType:"activity",postedTime:"2016-05-21T16:39:40.647-0600"}

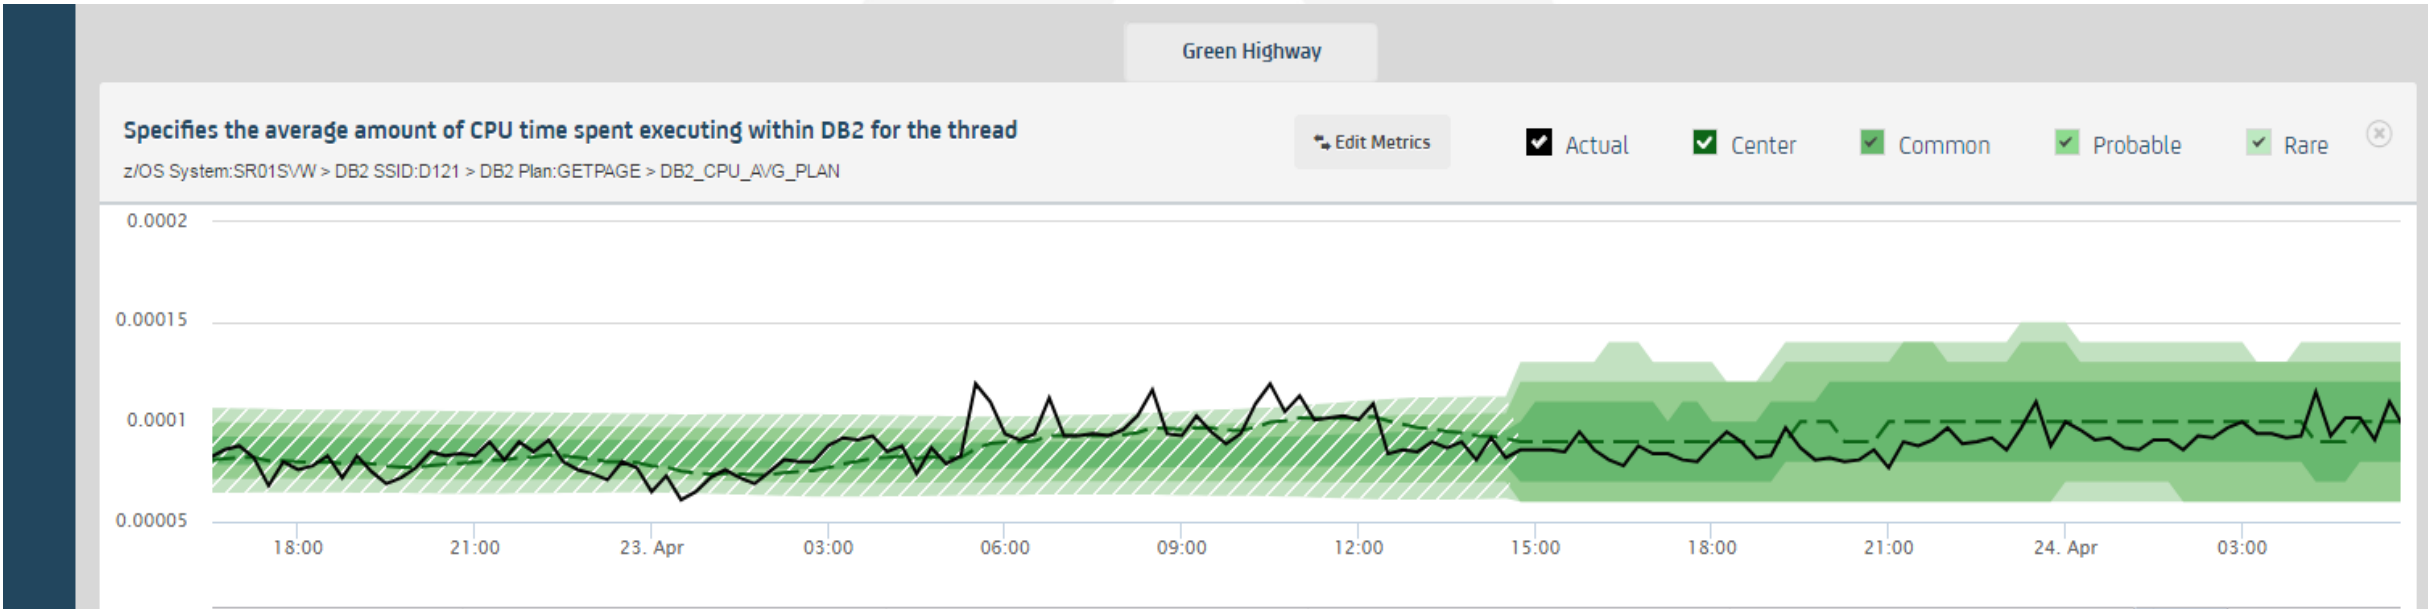
Company's Twitter ID

Source: Slideshare: analyzing machine data with Splunk

Mainframe Operational Intelligence



Idea: Leverage machine learning to understand normal verses abnormal behaviors



Outcome: Mature green highways

Thank you

Your feedback is important!

Submit a session evaluation for each session you attend:

SHARE mobile app

-or-

www.share.org/evaluation

