

**DB2 for z/OS – Ultimate Database for Cloud, Analytics and Mobile**

Industry-leading performance, security, scale and reliability



# “Guilty Until Proven Innocent”

Anthony Ciabattoni

Db2 for z/OS Development SWAT Team

[aciabattoni@ibm.com](mailto:aciabattoni@ibm.com)



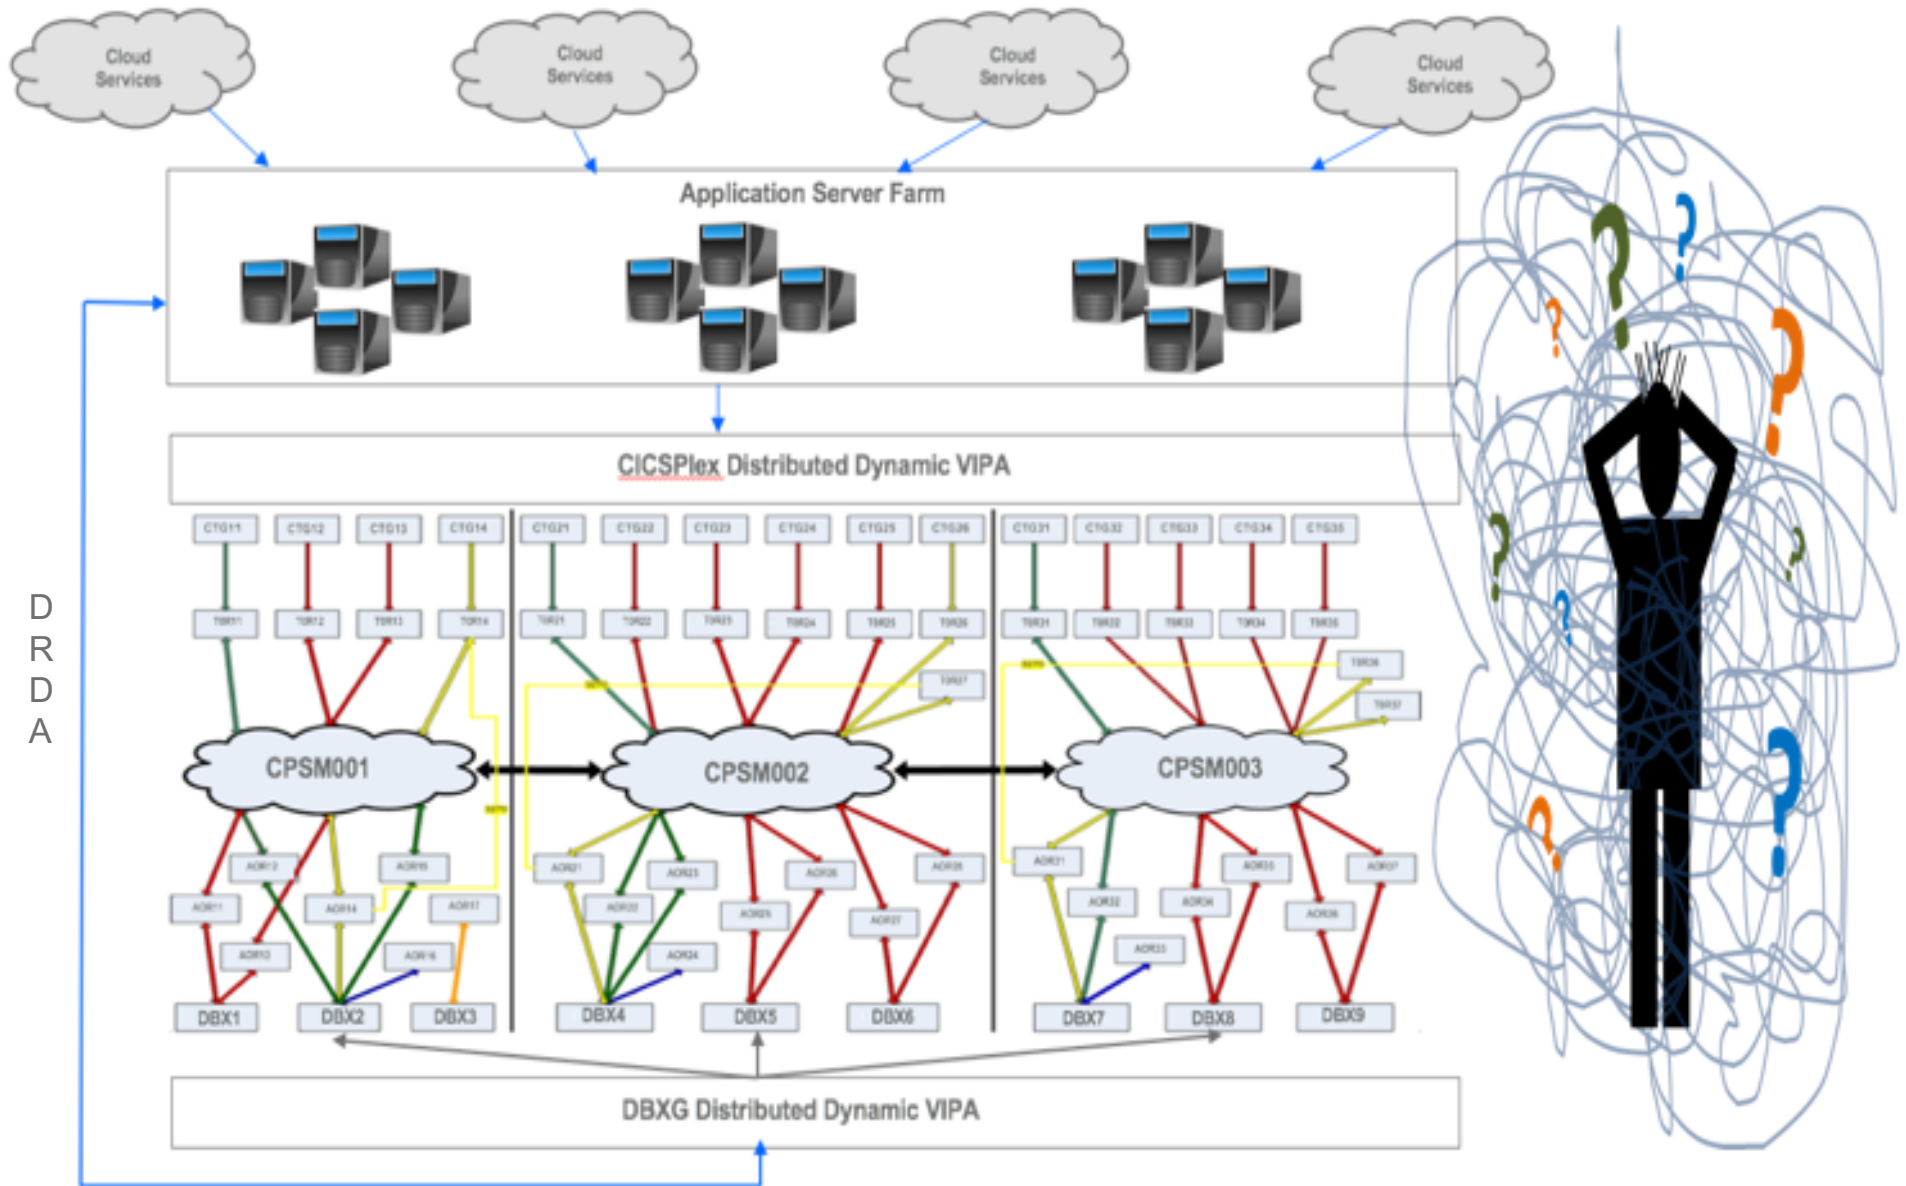
## Agenda

- **Infrastructure Investment**
- **Problem Resolution Procedures**
- **Application Transactional Flow**
- **Db2 Accounting Records and KPIs**
- **Troubleshooting Methodology**
- **Root cause analysis**
- **Summary**

## The Need for High Available Applications

- **The evolution of society and technology have driven organizations to design, implement and maintain high available, efficient, function rich and easy to use applications.**
  - Drive's growth & revenue
  - Customer Loyalty
  - Customer Confidence
  - Efficient systems leads to decreased operating costs
- **Organizations need to invest in order to achieve and exceed customers expectations**
  - Infrastructure and technology financial investment
  - Proper resource allocation
    - Subject Matter Experts (SMEs)
    - Number of skilled resources
  - Business & Executive commitment

# High Available Application Infrastructure



## Application Availability

- **“Uptime”/Availability Definition.**

- Uptime is a computer industry term for the time during which a computer application/infrastructure meets or exceeds pre-defined System Level Agreements (SLA's). The industry standard for uptime is Six Sigma – Six 9s, 99.9999 availability
  - SLA's can be external or internal measured availability
  - Often times SLAs are tied to monetary penalties

- **Measuring Availability**

- Each organization typically has their own standards for measuring and defining what a negative event maybe as well as penalties for not meeting SLAs.

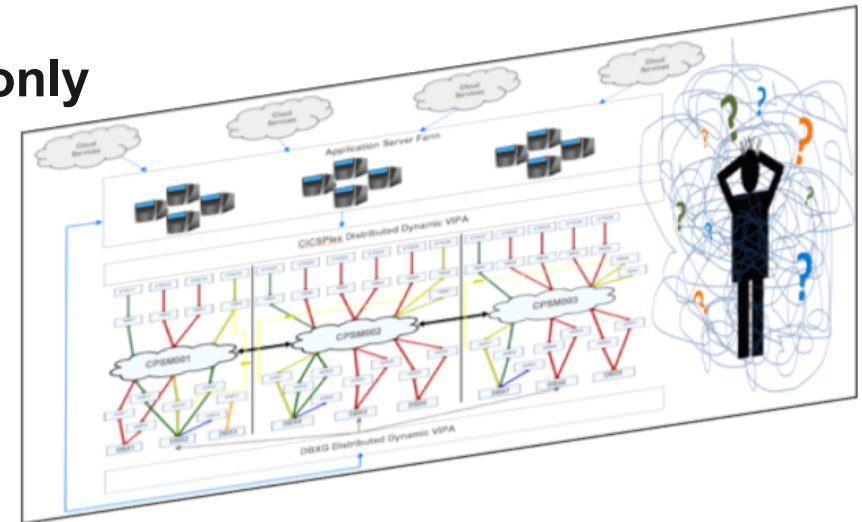
- System or Component outage

- Failed transactions
- Slow Application response
- Application meltdowns leading to system or component outage
- Failure to process a given deliverable within a specified time

- Negative events that impact SLAs often times receive extreme visibility and scrutiny to top executives within an organization

## Once an “Event” Occurs....

- An “All hands on deck” approach is commonly used to resolve an event
  - Subject Matter Experts (SME’s)
  - Management & Executives
- War room procedures
- Intense pressure to restore normal activity
  - Quick and correct actions are expected



## Db2 is often Pronounced Guilty!!!!

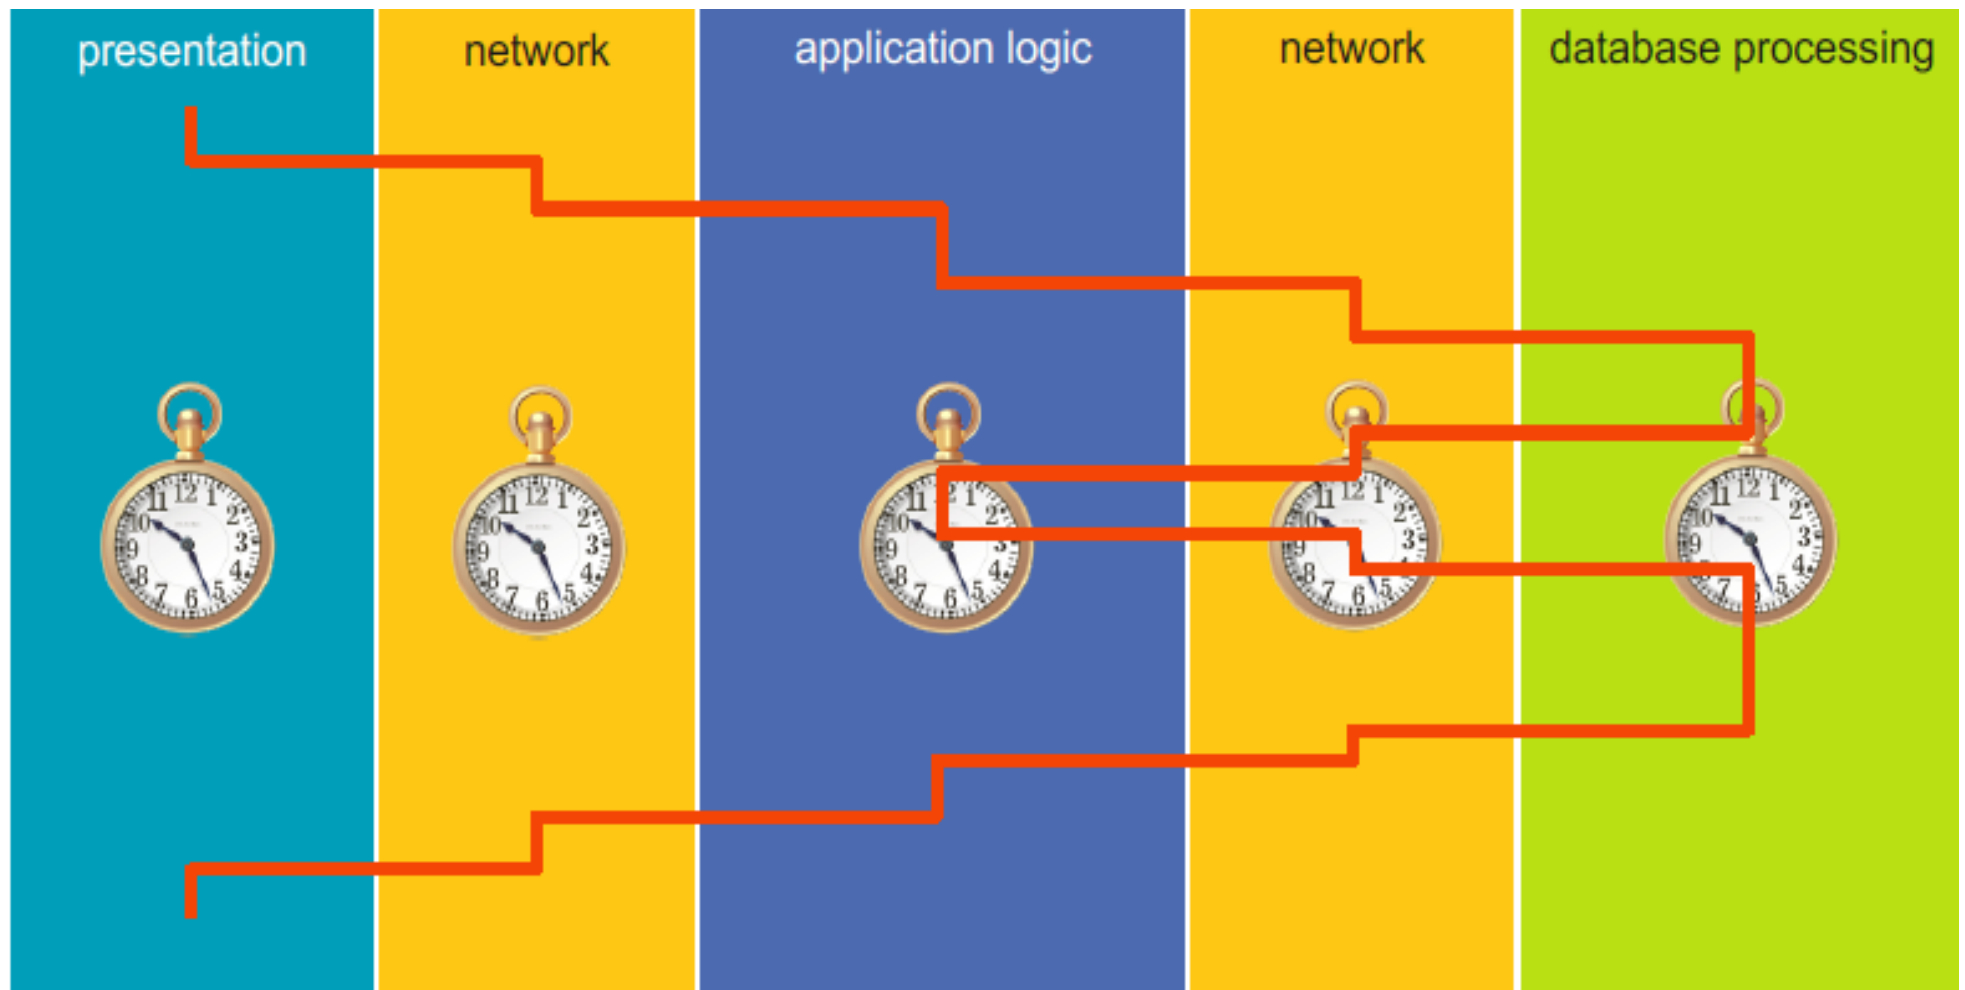




- **Why is z/OS and specifically Db2 an initial easy target of initial blame?**




# Application Infrastructure – Generic Flow



## Enterprise Instrumentation Data

- **Each infrastructure component has the potential of producing instrumentation data**
  - Cloud & Distributed application server farms
  - Network
  - z/OS
  - CICS
  - Storage
  - Db2 (Accounting & Statistics)
- **Various levels of instrumentation**
  - Turned off
  - Minimal
  - Default Settings
  - Robust
- **Applications can be built with additional instrumentation data**
  - Trouble shooting
  - Operational Analytics
  - Customer behavior
  - Configurable



## Db2 Instrumentation Data

### *Accounting Records*

- **Db2 Accounting trace records**

- Contain information about the activity performed by an individual Db2 thread or 'transaction' (or x threads in the case of roll-up accounting)
- Used as a prime indicator for thread-related problems

- **IFCIDS**

- **IFCID 3** – Plan level information
- IFCID 239 – Package level information
  - Since V8, IFCID is written as a separate record, but cut at same time

- **Accounting Classes**

- **Class 1 data – IFCID 3**

- Total times (elapsed & CPU)
- SQL statements counter
- Dynamic statement cache
- RID list processing
- Stored Procedures/triggers/UDFs
- Locking
- Data Sharing and Global Contention
- Query Parallelism
- Claim and drain activity
- Logging
- Resource limit facility (RLF)
- Buffer pool and GBP activity
- Distributed data facility (DDF)

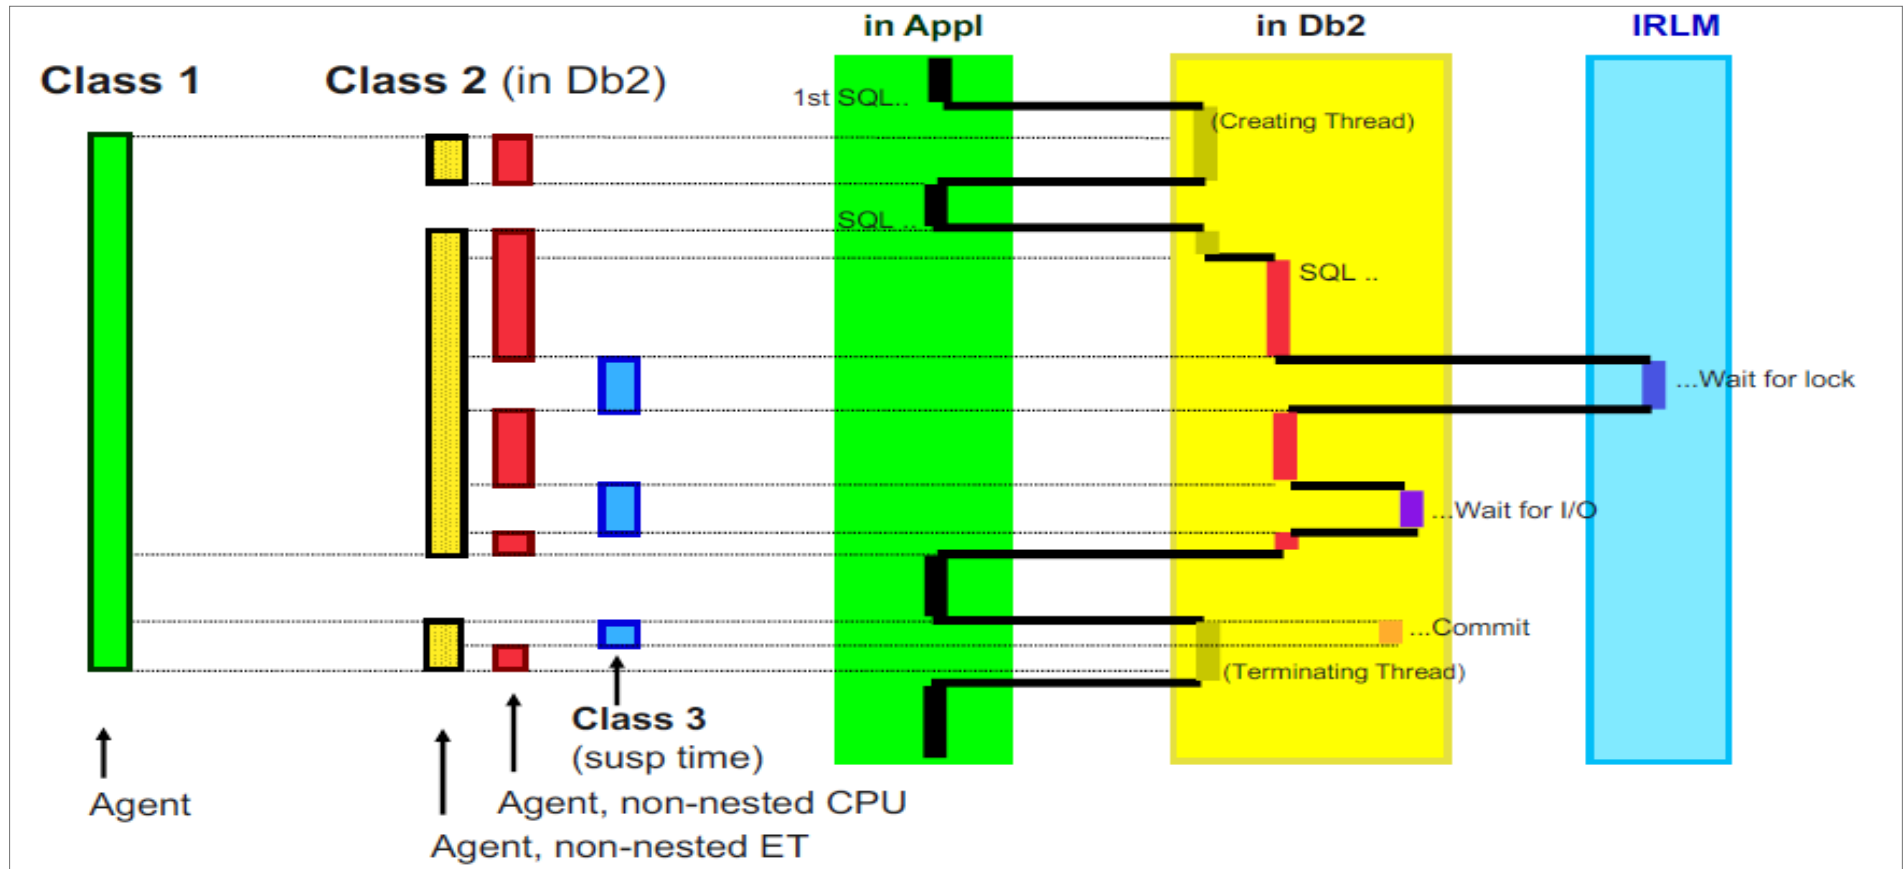
## Db2 Instrumentation Data

### *Accounting Records*

- **Db2 Accounting Classes ... trace records**
  - **Class 2 data**
    - Adds times in Db2 (elapsed and CPU) to existing **IFCID 3** record
  - **Class 3 data**
    - Activates all Db2 suspension counters and times and adds to existing IFCID 3 records
  - **Class 5 data**
    - Activates IFI counters
  - **Class 7 data → IFCIDF 239**
    - Package level info (similar to class 2 but at package level)
  - **Class 8 data**
    - Adds package level wait counters to existing IFCID 239 records
  - **Class 10 data**
    - Adds detailed package accounting (SQL/Locking/BP info) to existing IFCID 239 record
  - **Class 11 data**
    - To 'separate' writing of IFCID 3 from 239 (V11 only)
    - -STA TRA(A) Class(11) DEST(SMF) if you never want package info written to SMF

# Db2 Accounting Class 1, 2 & 3

## Illustration



## Db2 Accounting Class 1, 2 & 3

### *Description*

Thread activity time	=	Class 1 elapsed
└─ Elapsed time spent out of Db2 (in application, network, idle..)	=	Class 1 elapsed – Class 2 elapsed
└─ Elapsed time spent in Db2	=	Class 2 elapsed
└─ Processing time	=	Class 2 CPU (GCP and zIIP)
└─ Waiting time	=	Class 2 elapsed – Class 2 CPU
└─ Suspended time	=	Class 3 suspension time
└─ Not accounted time	=	Waiting time – Suspended time

## Breaking down time spent in Db2...

### *Accounting Class 3*

- **Suspended in Db2 for a known event (Class 3 suspension times)**
  - Many different types of suspensions (21 in V10)
- **Class 3 accounting strongly recommended**
- **Negligible overhead except when high internal Db2 latch contention (> 10,000/sec)**
- **Provides information about “known” Db2 suspensions, for suspensions Db2 reports:**
  - The number of suspensions
  - The total time it was suspended for those
- **Always look at both # of suspensions and suspension time**
  - Double accounting is avoided

## Db2 Accounting Class 3

		V5	V6	V7,8	V9	V10,V11	
Class 3 Suspension Types	I/O	Synchronous read/write & log write	✓				
		Synchronous read/write		✓	✓	✓	✓
		Log write		✓	✓	✓	✓
		Other agents' read	✓	✓	✓	✓	✓
		Other agents' write	✓	✓	✓	✓	✓
		Force-at-commit database writes (LOG NO LOBs only)		✓	✓	✓	✓
		TCP/IP LOB XML				✓	✓
	Locking	IRLM lock/latch & Db2 internal latch	✓	✓	✓	✓	
		IRLM lock/latch					✓
		Db2 internal latch					✓
		Page latch	✓	✓	✓	✓	✓
		Drain lock	✓	✓	✓	✓	✓
		Claim release	✓	✓	✓	✓	✓
	Sync.EU Switch	Synchronous Execution Unit Switch total	✓				
		Open/Close		✓	✓	✓	✓
		Define/Extend/Delete		✓	✓	✓	✓
		SYSLGRNX recording		✓	✓	✓	✓
		Commit		✓	✓	✓	✓
Archiving	Other services		✓	✓	✓	✓	
	Archive Log command	✓	✓	✓	✓	✓	
Scheduling	[Archive] log read	✓	✓	✓	✓	✓	
	Stored procedures	✓	✓	✓	✓	✓	
Data Sharing	UDFs		✓	✓	✓	✓	
	Global locks total	✓	✓				
	Parent L-locks			✓	✓	✓	
	Child L-locks			✓	✓	✓	
	Other L-locks			✓	✓	✓	
	Page set/Partition P-locks			✓	✓	✓	
	Page P-locks			✓	✓	✓	
	Other P-locks			✓	✓	✓	
	Sending Notify messages	✓	✓	✓	✓	✓	
	Asynchronous coupling facility request completion		✓	✓	✓	✓	



# Db2 Instrumentation Data

## Statistics Records

### • Db2 Statistics

#### – ZPARM SMFSTAT

- Yes (default) starts the trace for the default classes (1,3,4,5,6)
- CPU overhead of all Db2 Statistic Traces is negligible
- Recommendation to start all of the statistic traces by specifying SMFSTAT=\*
- Db2 Statistic records are written to SMF 100 records
- IFCIDS 2, 202, 217, 225 and 230 are written in fixed 1 minute intervals
- ZPARM STATIME applies to IFCID 105, 106, 199 and 365, default is 1 minute

<b>Class</b>	<b>Data Collected</b>	<b>IFCIDs Activated <sup>(1)</sup></b>
1	Statistics data	1, 2, 105,106, 202, 225
2	Installation-defined statistics record	152
3	Deadlock, lock escalation, group buffer pool, data set extension information, indications of long-running URs, and active log space shortages	172, 196, 250, 258, 261, 262, 313, 330, 335, 337
4	Db2 exceptional conditions	173,191-195, 203-210, 235, 236, 238, 267, 268, 343, <b>402</b>
5	Db2 data sharing statistics record	230 , 254
6	Storage usage details	225
<u>7</u>	<b><u>DRDA location statistics</u></b>	<b><u>365</u></b>
8	Data set I/O statistics	199
<u>9</u>	<b><u>Aggregated CPU and wait time by connection type</u></b>	<b><u>369</u></b>

# Troubleshooting Methodology

## *Establishing a Baseline*

- **The establishment of a performance baseline is a principal foundation component for effective troubleshooting.**

– What input data should be used

### Db2 Accounting Records

✓ IFCID 3

✓ Class 1, Class 2, Class 3

✓ IFCID 239

✓ Class 7, Class 8, Class 10

### Db2 Statistics Records

✓ For additional analysis

✓ Baseline for comparison

– What correlation id's should be captured

- Top x OLTP transactions
  - ✓ Different definitions per organization, typically CICS or DDF transactions
- Business critical batch programs

– When ... should the data be captured

- Time of day
- Day of week
- Day of month

– Capturing Approach

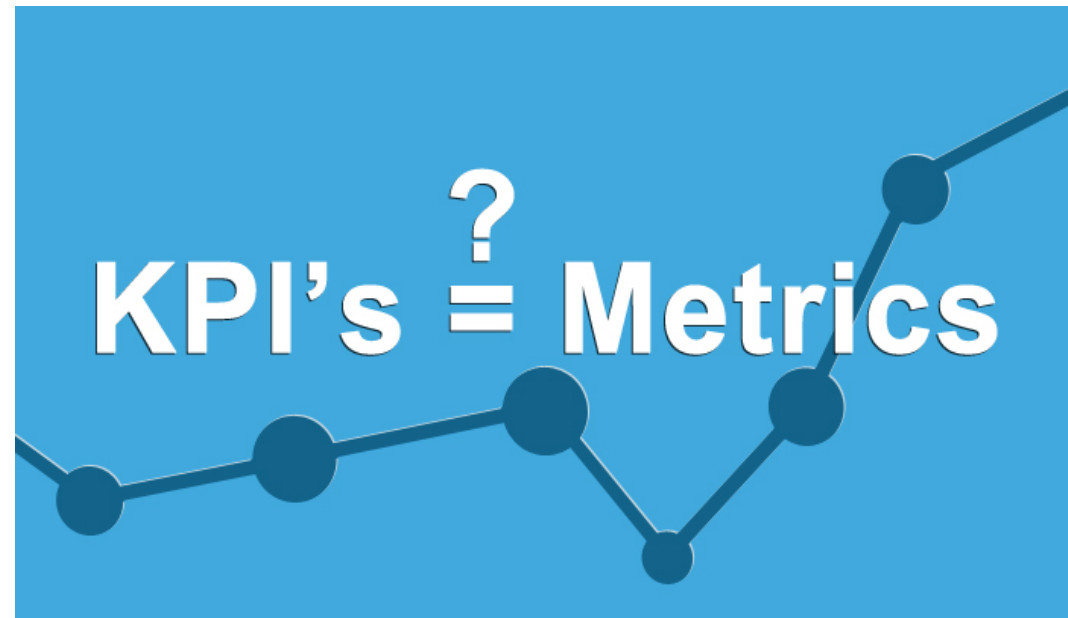
- Performance warehouse
- Ad-hoc baseline reporting

# Troubleshooting Methodology

## Metrics

- **Key Performance Indicators (KPI's)**

- Number of occurrences
  - Transactions/Second
  - Commits
- Elapsed Time
  - Elapsed/Transaction
  - Elapsed/Commit
- CPU Time
  - CPU/Transaction
  - CPU/Commit
- Additional Indicators
  - DML/Transaction and type – Select, Insert, Update, Delete
  - Getpages/Transaction
  - Average Sync I/O/Transaction
  - Number of select, Insert, update and deletes



# Troubleshooting Methodology

## *Additional Analysis*

- **Additional Analysis might be needed**
  - Class 1, Class 2 and Class 3 analysis can lead to other Db2 components to research
- **Subsystem Analysis**
  - Problem is more widespread than application transactions
  - Statistic Reports
  - System Monitors
    - Warning and Alerting
    - Capture exceptions for additional research
- **Package Analysis**
  - IFCID 239
    - Package level detail
    - Accounting Class 7, Class 8 and Class 10
  - Db2 Accounting batch reporting
  - System Monitors
- **SQL Analysis**
  - Monitors
  - Dynamic Statement Cache

# Omegamon XE for Db2

## Batch Reporting Methods

```
GLOBAL
  INTERVAL (1)
  INCLUDE (CONNECT (CICSXA2*), CORRNAME (TRN*))
ACCOUNTING
  REDUCE
    INTERVAL (1)
  REPORT
    LAYOUT (LONG)
    ORDER (INTERVAL-CORRNAME)
    DDNAME (IVRPTDD)
EXEC
```

### Omegamon XE for Db2 sample control statements

- Accounting Long Report
- Filtering on CICS Region and Transaction Id
- 1 Minute intervals
- Order by Interval (1 minute) and Transaction Id (Correlation)
- Short reporting interval
- Accessing SMF Dump Datasets to reduce latency

# Omegamon XE for Db2 Accounting Report

## Formatted Report

1 LOCATION: USSSDBXG OMEGAMON XE FOR DB2 PERFORMANCE MONITOR ( V5 )				ACCOUNTING REPORT - LONG		PAGE: 1-1	
GROUP: DBXG.				REQUESTED FROM: NOT SPECIFIED		TO: NOT SPECIFIED	
MEMBER: DBX3				ORDER: INTERVAL		INTERVAL FROM: 09/05/17 10:23:00.00	
SUBSYSTEM: DBX3				SCOPE: MEMBER		TO: 09/05/17 10:40:00.00	
DB2 VERSION: V11							
INTERVAL: 09/05 10:23 - 09/05 10:24							
ELAPSED TIME DISTRIBUTION				CLASS 2 TIME DISTRIBUTION			
APPL	=====> 14%			CPU	> 1%		
DB2	=====> 8%			SECPU	> 1%		
SUSP	=====> 77%			NOTACC	=====> 8%		
SUSP				SUSP	=====> 90%		
AVERAGE	APPL(CL.1)	DB2 (CL.2)	IFI (CL.5)	CLASS 3 SUSPENSIONS	AVERAGE TIME	AV.EVENT	HIGHLIGHTS
ELAPSED TIME	0.090611	0.077615	N/P	LOCK/LATCH(DB2+IRLM)	0.009877	0.42	#OCCURRENCES : 29475
NONNESTED	0.048225	0.037520	N/A	IRLM LOCK+LATCH	0.006485	0.27	#ALLIEDS : 0
STORED PROC	0.042386	0.040095	N/A	DB2 LATCH	0.003392	0.15	#ALLIEDS DISTRIB: 0
UDF	0.000000	0.000000	N/A	SYNCHRON. I/O	0.053502	8.70	#DBATS : 10575
TRIGGER	0.000000	0.000000	N/A	DATABASE I/O	0.052823	8.58	#DBATS DISTRIB. : 18900
CP CPU TIME	0.000780	0.000717	N/P	LOG WRITE I/O	0.000679	0.12	#NO PROGRAM DATA: 0
AGENT	0.000780	0.000717	N/A	OTHER READ I/O	0.006368	0.80	#NORMAL TERMINAT: 10575
NONNESTED	0.000458	0.000408	N/P	OTHER WRTE I/O	0.000000	0.00	#DDFRRSAF. ROLLUP: 189
STORED PRC	0.000322	0.000309	N/A	SER.TASK SWTCH	0.000010	0.00	#ABNORMAL TERMIN: 0
UDF	0.000000	0.000000	N/A	UPDATE COMMIT	0.000000	0.00	#CP/X PARALLEL. : 0
TRIGGER	0.000000	0.000000	N/A	OPEN/CLOSE	0.000010	0.00	#UTIL PARALLEL. : 0
PAR.TASKS	0.000000	0.000000	N/A	SYSLGRNG REC	0.000000	0.00	#IO PARALLELISM : 0
SE CPU TIME	0.000892	0.000830	N/A	EXT/DEL/DEF	0.000000	0.00	#PCA RUP COUNT : 0
NONNESTED	0.000578	0.000516	N/A	OTHER SERVICE	0.000000	0.00	#RUP AUTONOM. PR: 0
STORED PROC	0.000314	0.000314	N/A	ARC.LOG(QUIES)	0.000000	0.00	#AUTONOMOUS PR : 0
UDF	0.000000	0.000000	N/A	LOG READ	0.000000	0.00	#INCREMENT. BIND: 0
TRIGGER	0.000000	0.000000	N/A	DRAIN LOCK	0.000000	0.00	#COMMITTS : 32694
PAR.TASKS	0.000000	0.000000	N/A	CLAIM RELEASE	0.000000	0.00	#ROLLBACKS : 0
SUSPEND TIME	0.001956	0.069992	N/A	PAGE LATCH	0.000013	0.00	#SVPT REQUESTS : 0
AGENT	N/A	0.069992	N/A	NOTIFY MSGS	0.000000	0.00	#SVPT RELEASE : 0
PAR.TASKS	N/A	0.000000	N/A	GLOBAL CONTENTION	0.000013	0.00	#SVPT ROLLBACK : 0
STORED PROC	0.001956	N/A	N/A	COMMIT PH1 WRITE I/O	0.000000	0.00	MAX SQL CASC LVL: 1
UDF	0.000000	N/A	N/A	ASYNCH CF REQUESTS	0.000209	0.13	UPDATE/COMMIT : 0.11
NOT ACCOUNT.	N/A	0.006075	N/A	TCP/IP LOB XML	0.000000	0.00	SYNCH I/O AVG. : 0.006151
DB2 ENT/EXIT	N/A	4.59	N/A	ACCELERATOR	0.000000	0.00	
EN/EX-STPROC	N/A	1.63	N/A	AUTONOMOUS PROCEDURE	0.000000	0.00	
EN/EX-UDF	N/A	0.00	N/A	PQ SYNCHRONIZATION	0.000000	0.00	
DCAPT.DESCR.	N/A	N/A	N/P	TOTAL CLASS 3	0.069992	10.05	
LOG EXTRACT.	N/A	N/A	N/P				

# Omegamon XE for Db2 Accounting Report

- Accounting Long Report
- Each reporting section is broken up into 1 minutes time intervals

**DB2 for z/OS – Ultimate Database for Cloud, Analytics and Mobile**


Industry-leading performance, security, scale and reliability



## Trouble Shooting Examples...

Utilizing Db2 Accounting  
Class 1, Class 2 and Class 3 records

## High Accounting Class 1 CPU and/or Elapsed Time



- Application Logic inefficiency
- Network problems
- User think time
- Class 2 not active all the time
- Potential Db2 pitfalls
  - Thread reuse and long time between consecutive transactions (allied threads)
  - Long running DDF threads when CMTSTAT=ACTIVE
- Potential non-Db2 causes
  - Waiting on a 2-phased commit
    - CICS
    - Application
  - Network related issues
  - Application Server issues
  - CPU Spin



## Example #1 - High Accounting Class 1 Elapsed Time Con't

SSID	Date	Start-Time	End-Tim	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upsds/commt	T3-LLatch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs	
DB21	9/17/17	2:49:00 PM	2:50:00 PM	TRN-1	203	0.316884	0.012724	0.011820	0.001341	0.000929	N/C	0.002996	17.34	0.31	0.000018	0.011278	0.011244	0.000034	0.000312	32.11	
DB21	9/17/17	2:50:00 PM	2:51:00 PM	TRN-1	177	0.402622	0.015676	0.014680	0.001456	0.001022	N/C	0.003568	18.40	0.53	0.000012	0.014159	0.014120	0.000039	0.000358	35.80	
DB21	9/17/17	2:51:00 PM	2:52:00 PM	TRN-1	230	0.301680	0.015201	0.014218	0.001450	0.001010	N/C	0.003330	17.67	0.46	0.000009	0.013482	0.013434	0.000047	0.000481	41.10	
DB21	9/17/17	2:52:00 PM	2:53:00 PM	TRN-1	180	0.353893	0.013670	0.012816	0.001265	0.000879	N/C	0.003177	15.60	0.32	0.000006	0.012062	0.012028	0.000034	0.000716	32.30	
DB21	9/17/17	2:53:00 PM	2:54:00 PM	TRN-1	195	1.484222	0.014163	0.013098	0.001561	0.001084	N/C	0.002877	20.36	0.59	0.000007	0.012596	0.012533	0.000044	0.000440	39.70	
DB21	9/17/17	2:54:00 PM	2:55:00 PM	TRN-1	222	0.234721	0.015722	0.014824	0.001321	0.000924	N/C	0.003388	15.65	0.25	0.000014	0.013767	0.013724	0.000043	0.000839	31.56	
DB21	9/17/17	2:55:00 PM	2:56:00 PM	TRN-1	166	0.439619	0.016010	0.015007	0.001472	0.001028	N/C	0.003202	17.03	0.32	0.000019	0.014438	0.014398	0.000040	0.000438	39.33	
Avg					196	0.341570	0.014834	0.013894	0.001384	0.000965	N/C	0.003277	16.948333	0.365000	0.000013	0.013198	0.013158	0.000040	0.000524	35.99	
Max					230	1.484222	0.024714	0.023505	0.001933	0.001375	N/C	0.004985	21.870000	0.590000	0.000096	0.022976	0.022929	0.000370	0.001890	41.10	

### Increased Db2 Accounting Class 1 elapsed time

- Db2 is typically innocent
- Elongated transaction elapsed time is due to something outside of Db2 (Class 1).
  - After the initial Db2 thread was created

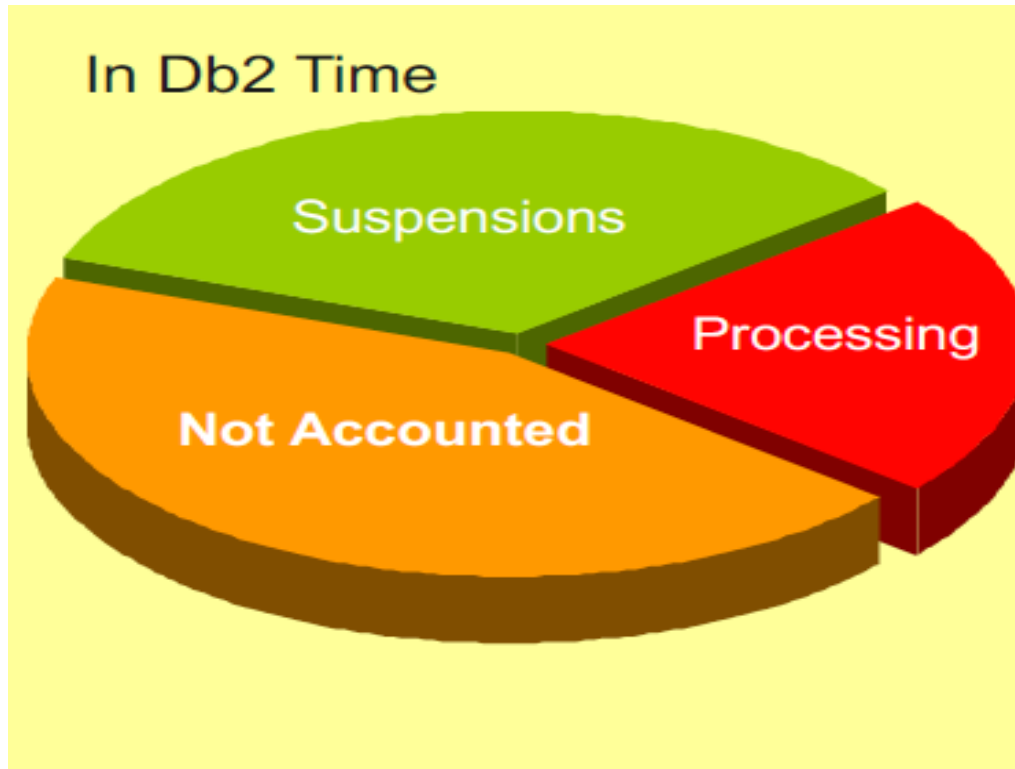
## Example #2 - High Accounting Class 1 CPU Time

SSID	Date	Start-Time	End-Time	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upds/commt	T3-Llatch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs
DB21	8/21/17	10:49:00 PM	10:50:00 PM	TRN-2	162	0.316884	0.012724	0.011820	0.001341	0.000929	N/C	0.002996	17.34	0.31	0.000018	0.011278	0.011244	0.000034	0.000312	35.43
DB21	8/21/17	10:50:00 PM	10:51:00 PM	TRN-2	141	0.402622	0.015676	0.014680	0.001456	0.001022	N/C	0.003568	18.40	0.53	0.000012	0.014159	0.014120	0.000039	0.000358	34.15
DB21	8/21/17	10:51:00 PM	10:52:00 PM	TRN-2	184	0.301680	0.015201	0.014218	0.001450	0.001010	N/C	0.003330	17.67	0.46	0.000009	0.013482	0.013434	0.000047	0.000481	38.27
DB21	8/21/17	10:52:00 PM	10:53:00 PM	TRN-2	144	0.353893	0.013670	0.012816	0.001265	0.000879	N/C	0.003177	15.60	0.32	0.000006	0.012062	0.012028	0.000034	0.000716	33.70
DB21	8/21/17	10:53:00 PM	10:54:00 PM	TRN-2	173	0.414222	0.014163	0.013098	0.001561	0.001084	N/C	0.002877	20.36	0.59	0.000007	0.012596	0.012533	0.000044	0.000440	40.90
DB21	8/21/17	10:54:00 PM	10:55:00 PM	TRN-2	178	0.234721	0.015722	0.014824	0.001321	0.000924	N/C	0.003388	15.65	0.25	0.000014	0.013767	0.013724	0.000043	0.000839	33.74
DB21	8/21/17	10:55:00 PM	10:56:00 PM	TRN-2	133	0.659619	0.016010	0.015007	0.021472	0.001028	N/C	0.003202	17.03	0.32	0.000019	0.014438	0.014398	0.000040	0.000438	37.33
Avg					159	0.337337	0.014526	0.013576	0.001399	0.000975	N/C	0.003223	17.503333	0.410000	0.000011	0.012891	0.012847	0.000040	0.000524	36.22
Max					184	0.659619	0.016010	0.015007	0.021472	0.001084	0.000000	0.003568	20.360000	0.590000	0.000019	0.014438	0.014398	0.000047	0.000839	40.90

### Increased Db2 Accounting Class 1 CPU time

- Db2 is typically innocent
- Increased z/OS CPU time is due to a z/OS (Class 1) but not attributed with a Db2 process.

## High Accounting Class 2 Not Accounted for Time



- GCP or zIIP processing wait
  - Wrong dispatching priorities
    - ✓ No Db2 work should run in discretionary
  - Overloaded GCP or zIIP
    - ✓ Use RMF for details
    - ✓ General tuning
- Paging
  - Overcommitted real storage
    - ✓ Use RMF for details
    - ✓ Check Buffer pools, EDM pools, Sort, number of threads, MAXKEEPD too large...
- Excessive detailed online tracing with vendor tools


## Example #3 - High Accounting Class 2 Elapsed Time Not Accounted for Time

SSID	Date	Start-Time	End-Time	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upds/commt	T3-LLatch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs	
DB21	5/23/17	9:11:00 AM	9:12:00 AM	TRN-3	162	1.082613	0.778453	0.011820	0.001341	0.000929	0.765729	0.002996	17.34	0.31	0.000018	0.011278	0.011244	0.000034	0.000312	32.11	
DB21	5/23/17	9:12:00 AM	9:13:00 AM	TRN-3	141	1.356568	0.969622	0.014680	0.001456	0.001022	0.953946	0.003568	18.40	0.53	0.000012	0.014159	0.014120	0.000039	0.000358	35.80	
DB21	5/23/17	9:13:00 AM	9:14:00 AM	TRN-3	184	1.384857	1.098378	0.014218	0.001450	0.001010	1.083177	0.003330	17.67	0.46	0.000009	0.013482	0.013434	0.000047	0.000481	41.10	
DB21	5/23/17	9:14:00 AM	9:15:00 AM	TRN-3	144	1.696874	1.356651	0.012816	0.001265	0.000879	1.342981	0.003177	15.60	0.32	0.000006	0.012062	0.012028	0.000034	0.000716	32.30	
DB21	5/23/17	9:15:00 AM	9:16:00 AM	TRN-3	158	2.057514	1.587455	0.013098	0.001561	0.001084	1.573292	0.002877	20.36	0.59	0.000007	0.012596	0.012533	0.000044	0.000440	39.70	
DB21	5/23/17	9:16:00 AM	9:17:00 AM	TRN-3	178	1.394030	1.175031	0.014824	0.001321	0.000924	1.159309	0.003388	15.65	0.25	0.000014	0.013767	0.013724	0.000043	0.000839	31.56	
DB21	5/23/17	9:17:00 AM	9:18:00 AM	TRN-3	133	1.165117	0.741508	0.015007	0.001472	0.001028	0.725498	0.003202	17.03	0.32	0.000019	0.014438	0.014398	0.000040	0.000438	39.33	
Avg					157	1.346677	1.019941	0.013894	0.001384	0.000965	1.086276	0.003277	16.948333	0.365000	0.000013	0.013198	0.013158	0.000040	0.000524	35.99	
Max					184	2.057514	1.587455	0.023505	0.001933	0.001375	1.573292	0.004985	21.870000	0.590000	0.000096	0.022976	0.022929	0.000370	0.001890	41.10	

### Increased Db2 Accounting Class 2 elapsed time – Not Accounted for Time

- Increased Class 2 elapsed time is due to a non accounted for Db2 is unable to classify where the time is being spent and the time is placed in the Not Accounted for bucket.
- Increased Not Accounted for Time typically points to an event or configuration outside of Db2

## High Accounting Class 3 Average Sync I/O



- High Average Sync I/O
  - Make sure there are enough I/O resources
  - Use RMF for additional analysis
    - DASD contention
    - High IOSQ time – use PAV
    - Control Unit cache misses
    - CPU contention
    - Inappropriate I/O priority
  - Tune I/O subsystem
  - User faster I/O devices (e.g. SSD)
  - For detailed analysis, use performance class 4

## Example #4 - High Accounting Class 3 Elapsed Time


### Average Sync I/O

SSID	Date	Start-Time	End-Time	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upsds/commt	T3-Llatch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs	
DB23	11/13/17	18:19	18:20	QTRN	210	0.013842	0.007802	0.006866	0.000515	0.000475	N/C	0.001046	4.28	0.24	0.000044	0.006529	0.006091	0.000438	0.000264	43.13	
DB23	11/13/17	18:20	18:21	QTRN	230	0.012103	0.007197	0.006334	0.000502	0.000463	N/C	0.000984	4.21	0.25	0.000041	0.006068	0.005706	0.000362	0.000198	43.41	
DB23	11/13/17	18:21	18:22	QTRN	241	0.047339	0.041338	0.039546	0.000501	0.000461	0.000744	0.005731	3.93	0.24	0.001924	0.031720	0.031360	0.000361	0.005828	45.18	
DB23	11/13/17	18:22	18:23	QTRN	270	0.012454	0.007669	0.006841	0.000486	0.000448	N/C	0.001078	3.87	0.27	0.000043	0.006509	0.006181	0.000327	0.000259	44.02	
DB23	11/13/17	18:23	18:24	QTRN	233	0.014783	0.009512	0.008522	0.000511	0.000472	N/C	0.001234	4.02	0.26	0.000069	0.007868	0.007535	0.000333	0.000550	44.23	
DB23	11/13/17	18:24	18:25	QTRN	221	0.055783	0.052411	0.051110	0.000460	0.000423	0.000323	0.007320	3.52	0.27	0.001432	0.038744	0.038302	0.000442	0.010889	43.16	
DB23	11/13/17	18:25	18:26	QTRN	246	0.011822	0.008428	0.007588	0.000473	0.000435	N/C	0.001204	3.87	0.27	0.000048	0.007230	0.006672	0.000559	0.000274	42.72	
Avg						0.013001	0.008122	0.007230	0.000497	0.000459	0.000534	0.001109	4.05	0.26	0.000049	0.006841	0.006437	0.000404	0.000309	43.50	
Max						0.055783	0.052411	0.051110	0.000522	0.000481	0.000744	0.007320	4.28	0.27	0.001924	0.038744	0.038302	0.000559	0.010889	45.18	

## Increased Db2 Accounting Class 2 and Class 3 elapsed time

- Increased Class 3 elapsed times does not pinpoint the root cause, further research is needed
- The captured QTRN transaction below is a skinny distributed native SQL stored procedure
- Db2 is typically a victim and innocent when the average sync i/o increases
- The identified customer root cause below was incorrect configuration of new Parallel Access Volumes(PAV) storage volumes resulting in a spike in average sync I/O

## High Accounting Class 2 CPU & Elapsed



- What's the actual number
  - Using a percentage can be misleading
- Is this a Db2 intensive process anyway?
  - Compare to baseline
- Need to analyze time distribution in Db2
  - Where is the time actually spent in Db2?
  - Start with KPI's

## High Accounting Class 2 CPU & Elapsed Con't

### *Breaking down time spent in Db2*

- Burning CPU cycles (Class 2 CPU time)
  - Non-nested or nested
  - ON a general CP or zIIP
- In Db2 but don't know what it is we are doing
  - Mostly things happen outside of Db2's control, such as lack of CPU or paging
  - Should be a small percent
- Any known changes?
  - New Code, DML behavior changes
  - Index changes
  - Rebinds/new access path
  - Data related?



## Example #5 - High Accounting Class 2 CPU & Elapsed


SSID	Date	Start-Time	End-Time	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upds/commt	T3-Latch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs	
DB21	9/6/17	9:07:00 PM	9:08:00 PM	TRN-5	162	0.316884	0.012724	0.011820	0.001341	0.000929	N/C	0.002996	17.34	0.31	0.000018	0.011278	0.011244	0.000034	0.000312	32.11	
DB21	9/6/17	9:08:00 PM	9:09:00 PM	TRN-5	141	0.402622	0.015676	0.014680	0.001456	0.001022	N/C	0.003568	18.40	0.53	0.000012	0.014159	0.014120	0.000039	0.000358	35.80	
DB21	9/6/17	9:09:00 PM	9:10:00 PM	TRN-5	184	0.301680	0.015201	0.014218	0.001450	0.001010	N/C	0.003330	17.67	0.46	0.000009	0.013482	0.013434	0.000047	0.000481	41.10	
DB21	9/6/17	9:10:00 PM	9:11:00 PM	TRN-5	144	0.356777	0.016554	0.012816	0.004149	0.003763	N/C	0.003177	15.60	0.32	0.000006	0.012062	0.012028	0.000034	0.000716	138.30	
DB21	9/6/17	9:11:00 PM	9:12:00 PM	TRN-5	173	0.427662	0.017603	0.013098	0.005001	0.004524	N/C	0.002877	20.36	0.59	0.000007	0.012596	0.012533	0.000044	0.000440	165.70	
DB21	9/6/17	9:12:00 PM	9:13:00 PM	TRN-5	178	0.234721	0.015722	0.014824	0.001321	0.000924	N/C	0.003388	15.65	0.25	0.000014	0.013767	0.013724	0.000043	0.000839	31.56	
Avg					163	0.342051	0.015315	0.013894	0.001865	0.001446	N/C	0.003277	16.95	0.365000	0.000013	0.013198	0.013158	0.000040	0.000524	69.13	
Max					184	0.427662	0.017603	0.014824	0.005001	0.004524	N/C	0.003568	20.36	0.590000	0.000018	0.014159	0.014120	0.000047	0.000839	165.70	

### Increased Db2 Accounting Class 2 CPU time

- Db2 is typically guilty when Class 2 CPU increases
- In the scenario below the average number of get pages per DML increases substantially. This could be a data driven event where Db2 needs to access more data for the desired results.

# High Accounting Class 3 Elapsed Time

## *Db2 Locking/Latching Contention*



- Time per suspension large?
- Lock Suspensions
  - Non-committing transactions
  - Concurrent DDL
    - `RELEASE(DEALLOCATE)`
  - For detailed analysis, use performance trace classes 6 & 7
- Db2 latch suspension
  - Designed to be short in duration
  - Db2 10...
    - More efficient latch resume logic
    - Less chance to get resumed and suspended again without ever obtaining the latch in between
  - Can be a symptom of another problem
  - For detailed analysis, use performance trace class 11

## Example #6 - High Accounting Class 3 Elapsed Time

### *Db2 Locking/Latching Contention*

SSID	Date	Start-Time	End-Tim	Tran-id	Tran/Sec	Avg-Cl1-Elap	Avg-Cl2-Elap	Avg-Cl3-Elap	Avg-Cl1-Cpu	Avg-Cl2-Cpu	Not-Acct-Time	Avg-Syncio	Avg-DML	Upsds/commt	T3-Llatch	T3-Syncio	T3-DB-io	T3-Log-Wio	T3-OthRead	Total-GPs	
DB21	6/3/17	1:45:00 AM	1:46:00 AM	TRN-6	283	0.544504	0.026046	0.024780	0.001816	0.001271	0.000000	0.004777	20.16	0.45	0.000000	0.023787	0.023736	0.000051	0.000917	41.46	
DB21	6/3/17	1:46:00 AM	1:47:00 AM	TRN-6	485	0.223610	0.016890	0.015851	0.001504	0.001040	0.000000	0.004060	16.93	0.31	0.000017	0.015290	0.015256	0.000034	0.000501	35.24	
DB21	6/3/17	1:47:00 AM	1:48:00 AM	TRN-6	176	1.811042	1.441796	1.440717	0.001547	0.001076	0.000003	0.003700	16.98	0.40	1.424867	0.015022	0.014981	0.000041	0.000774	37.88	
DB21	6/3/17	1:48:00 AM	1:49:00 AM	TRN-6	183	1.962918	1.691741	1.690665	0.001558	0.001082	0.000000	0.003846	18.40	0.35	1.673630	0.016366	0.016332	0.000034	0.000597	37.43	
DB21	6/3/17	1:49:00 AM	1:50:00 AM	TRN-6	251	1.095249	0.818311	0.817220	0.001575	0.001090	0.000001	0.003561	17.61	0.37	0.802740	0.013811	0.013775	0.000036	0.000619	37.28	
DB21	6/3/17	1:50:00 AM	1:51:00 AM	TRN-6	431	0.301384	0.014809	0.013327	0.001414	0.000978	0.000004	0.003520	15.77	0.21	0.000006	0.013332	0.013300	0.000033	0.000455	32.81	
DB21	6/3/17	1:51:00 AM	1:52:00 AM	TRN-7	425	0.580979	0.023776	0.022606	0.001666	0.001167	0.000004	0.004651	17.69	0.39	0.000002	0.021946	0.021891	0.000055	0.000589	34.21	
Avg					319	0.931384	0.576196	0.575095	0.001583	0.001101	0.000002	0.004016	17.65	0.35	0.557323	0.017079	0.017039	0.000041	0.000636	36.62	
Max					485	1.962918	1.691741	1.690665	0.001816	0.001271	0.000004	0.004777	20.16	0.45	1.673630	0.023787	0.023736	0.000055	0.000917	41.46	

### Increased Db2 Accounting Class 2 and Class 3 elapsed time

- Increased Class 3 elapsed times does not pinpoint the root cause, further research is needed
- Increase in Class 3 Db2 latch elapsed time is due to Db2 lock/latch contention
- Db2 is typically guilty
- Additional research is needed to identify the root cause

## Post-Outage Research

### *Root Cause Analysis*

- **While an event is occurring, the primary objective is restore an application, infrastructure or system as expeditiously (or quickly) as possible:**
  - Minimize SLA penalties
  - Minimize customer impact
- **After an event occurs, additional time is typically spent to identify the root cause:**
  - Detailed data analysis and research on sequence of events
- **Determinations to be made as part of the root cause analysis:**
  - What was the root cause(s) of the event?
  - What was the impact of the outage (Failed SLAs?)
  - How did the problem occur?
  - Why did the problem occur?
  - Are we exposed to experiencing the problem again?
  - What can be done to eliminate this problem from happening in the future?
  - What alerting can be added to proactively notify us prior to an outage?

## Post-Outage Research

### *Pinpointing the Root Cause*

- **Change is a necessity for progress and success and is also a potential cause of problems, outages and negatively impacting events**
- **When there is a negatively impacting event, organizations must ask the fundamental question...”Have any documented or non-documented changes been made prior to the event?”**
  - Db2 Changes
    - Maintenance and/or version upgrades
    - Configuration changes (ZPARMs), bufferpool changes (group or local)
    - Object changes, ie, altered tables, new or modified indexes
    - REORGs, RUNSTATS or REBINDs
  - Non-Db2 Related Changes
    - z/OS component changes , z/OS, CICS, MQ Series, Network, Firewall
    - Distributed Infrastructure changes
      - ✓ New or modified distributed servers
      - ✓ Non-z/OS network/firewall changes
    - Application Changes
      - ✓ Code changes, new or modified SQL
      - ✓ Data access changes, day of week, month, year

## Summary

- **Availability and SLAs will continue to be a top objective for organizations**
- **Customer impacting events are inevitable**
- **Without a proper troubleshooting methodology z/OS and Db2 will continue to be an an easy initial culprit**
- **Identifying KPIs, baseline data and an utilizing Db2 Accounting records can change the paradigm from a target to a go-to metric**
- **Root cause analysis hints and tips**

# Questions

