DB2 Performance Tuning: Where do we start?

Jeff M. Sullivan
IBM Systems and Technology Group Lab Services

Thursday, August 11, 2011 1:30 PM-2:30 PM
Europe 6 (Walt Disney World Dolphin Resort)

So many places to look…

And adjust...

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- **The politics of performance**
  - The data base administrator role
  - Management expectations
  - DBA expectations
- **Performance Tuning**
  - Graphs and Trends
  - Playbooks
- **DB2 biggest performance issues**
- **Tooling**
- **Performance Touch Points with DB2 for z/OS**

The data base administrator role

- **Protection of the company's data assets**
- **Availability of the company's data assets**
- **Fast delivery of the company's data assets**
- **Provide analysis on problems as required**

Protection, Availability, and Speed are the goals.

Performance aspect – logging/recoverability

Protection of the company's data assets

- Adequate backup and recovery
- Retention of logging media
- Ensuring adequate storage and retention of logs and backups
- Performing disaster recovery scenarios

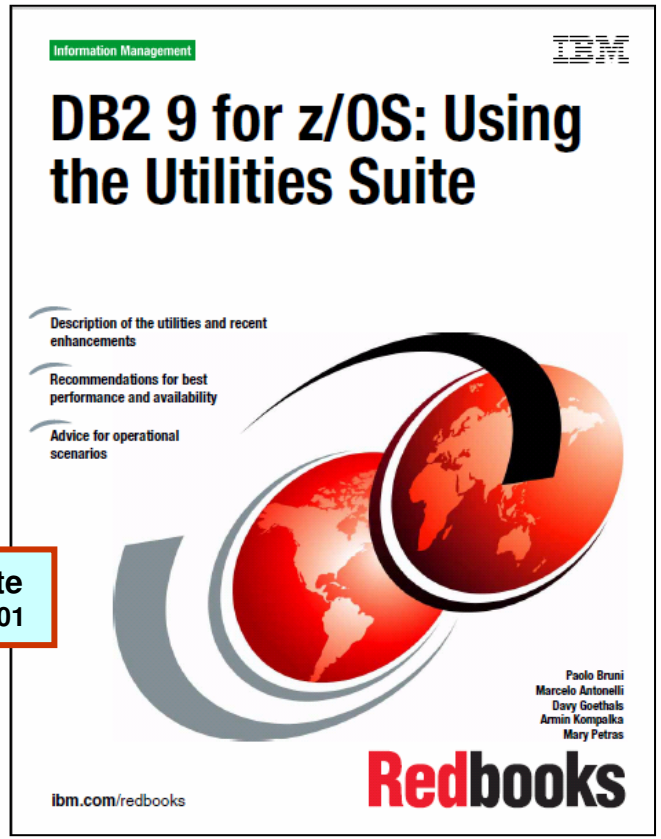

DB2 9 for z/OS: Using the Utilities Suite
February 2010 **SG24-6289-01**

Performance aspect – application and SQL, maintenance

Availability of the company's data assets

- Reviewing data paths into and out of the DBMS
  - Efficiency
  - Connections
  - Minimal server "hops“
  - Outage prevention
- Routine maintenance (data growth)

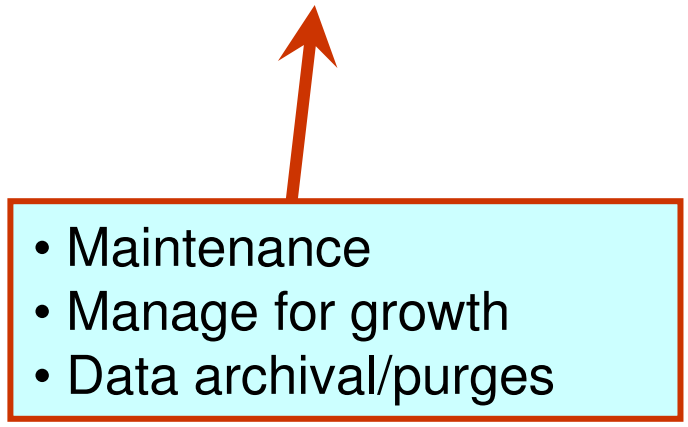

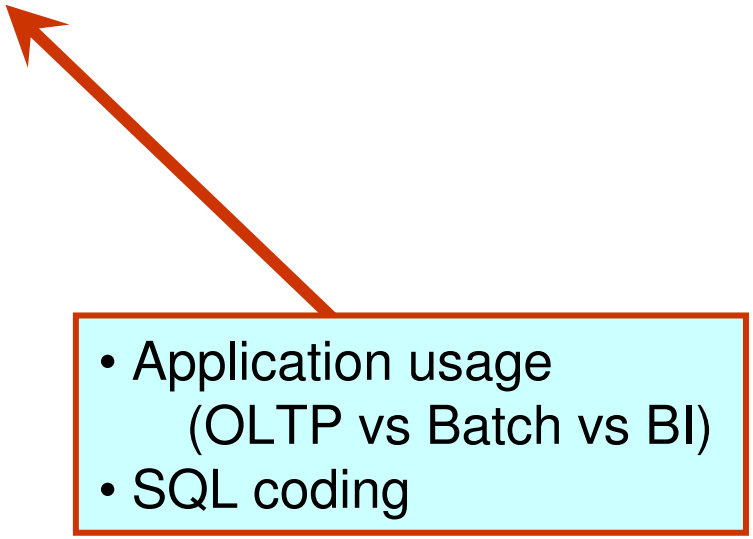

Performance aspect – SLA/SLO

Fast delivery of the company's data assets

- Creation of and meeting service level objectives
- Access paths and access strategies
- Ensure the data is delivered as quickly and efficiently as possible

- Performance data collection
- Measurable SLA/SLOs
- Object placement and I/O

Management's expectation & perspective

- ***Ability to look at the overall environment***
- ***Make projections on the growth***
- ***Plan for any future projects***

Motivation: To rein in spending and control costs

To an I.T manager perspective:

- ***Making a pretty good estimate on how much the overall environment will grow***
- ***Resolve outstanding issues that would enable postponing upgrades***
- ***Not being blindsided by an issue not caught by the database administration staff***

Data base administrator's perspective

- *Faster mean time to problem resolution*
- *Get to the problem faster than having the users complain to his manager*

Motivation: To not be blindsided by a problem

To a DBA this translates to:

- *Monitoring by rules*
- *Responding quickly to the problems when they happen*
- *Planning for that next growth issue*

A few words on performance “politics”

- Usually do not happen when applications are…
  - Unit tested
  - System tested or shakedown tested
  - Integration tested
- But when they happen, DBAs ask…
  - Is this a one-off situation?
  - Or is this a new “steady state”?
    - New workload
    - Existing work changes
    - DBA misstep

The DBA “unwritten” Code

(1) To minimize problem phone calls
(2) To make the on-call rotation a non-item task
(3) To never receive the problem from your boss

A few words on performance "myths"

"…If we tune the environment, we can reduce MIPS being used by the application…"

"…Performance tuning is easily resolved by finding the spikes and correcting for those spikes…"

"…We're getting performance problems in _____. We did not change anything…"

Better approach:

- Tune the application
- Put the "spikes" in context
  - New app(s)?
  - Change in app(s)?
  - Change in environment?
  - Continued growth?
- Consistent measurement – How do you know there is poor performance?

Agenda

- **The politics of performance**
  - The data base administrator role
  - Management expectations
  - DBA expectations
- **Performance Tuning**
  - Graphs and Trends
  - Playbooks
- **DB2 biggest performance issues**
- **Tooling**
- **Performance Touch Points with DB2 for z/OS**

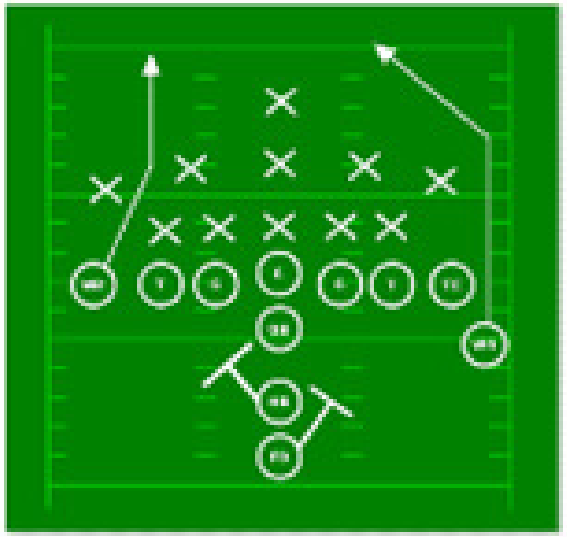

System-level vs. Local: Symptoms & Causes

- Performance problems come in two main flavors
  - Localized, impacting a subset of statements or applications
  - Broad-scope, impacting the entire system
- Many individual problems can combine to create a system-level symptom!

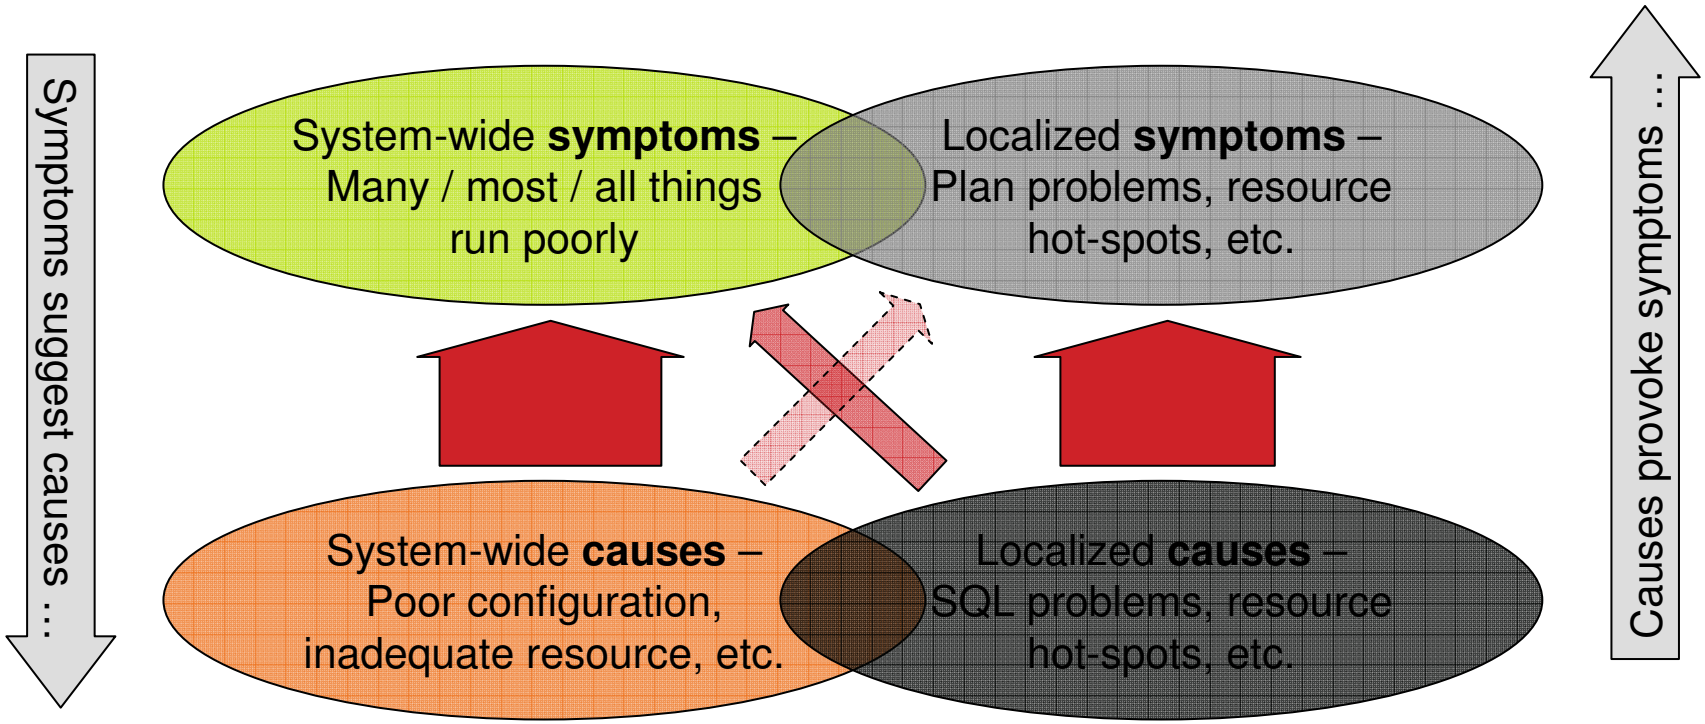

Usage and Growth

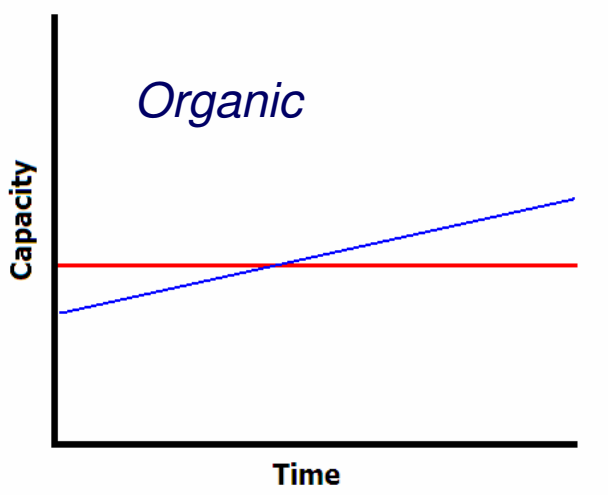

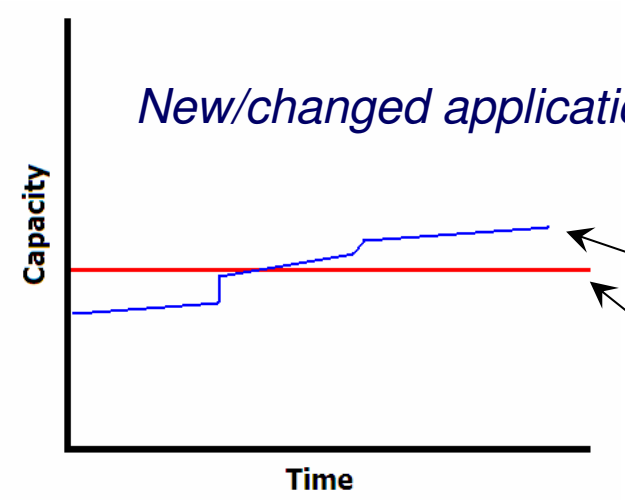

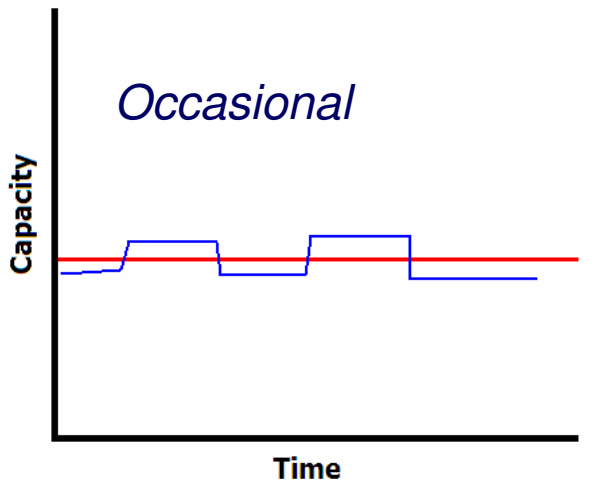

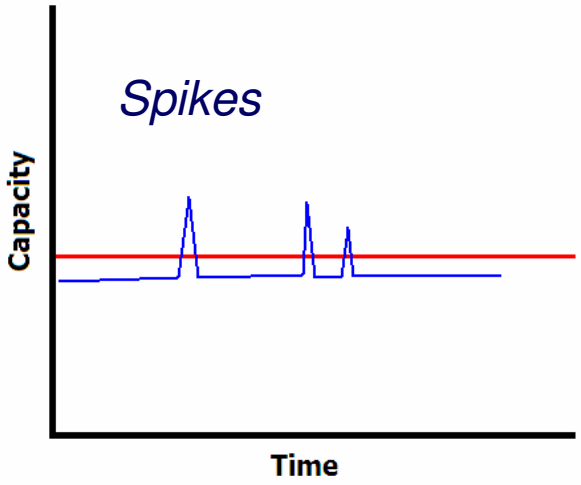

Usage and Growth

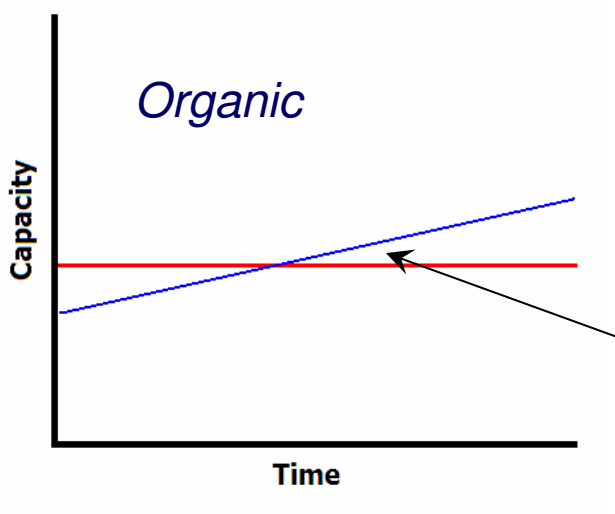

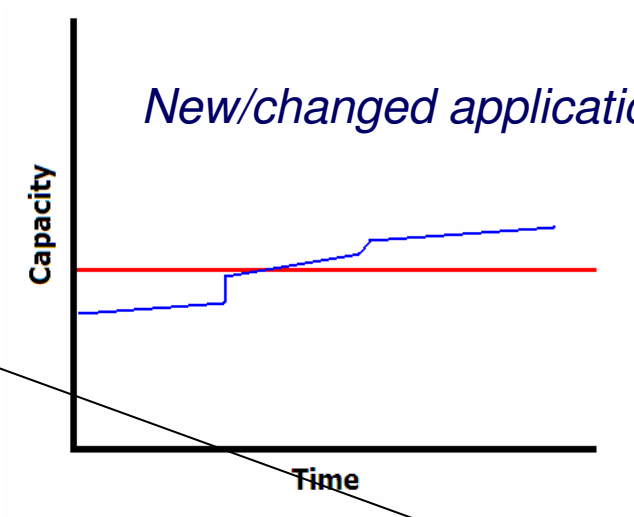

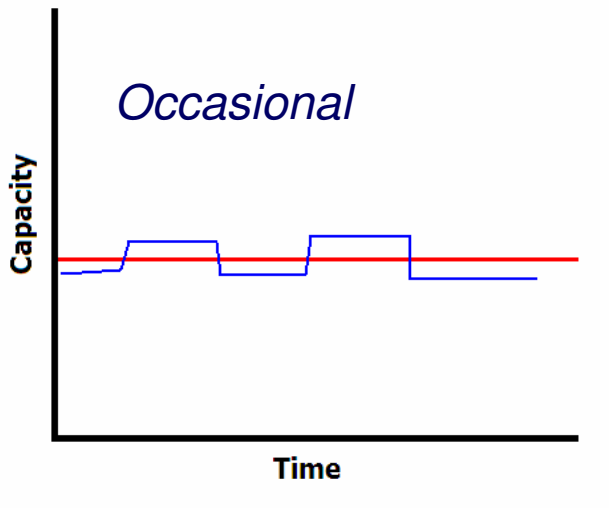

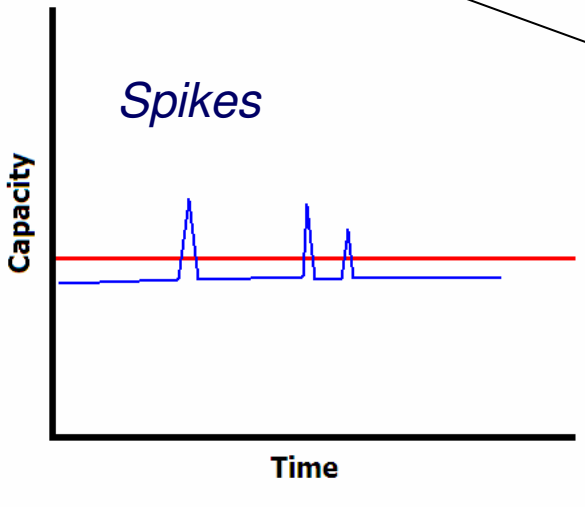

Data Growth:
- Old statistics
- I/O bottlenecks

Usage and Growth

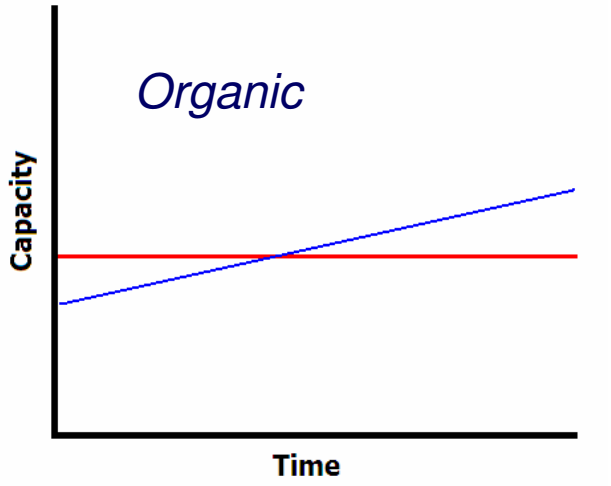

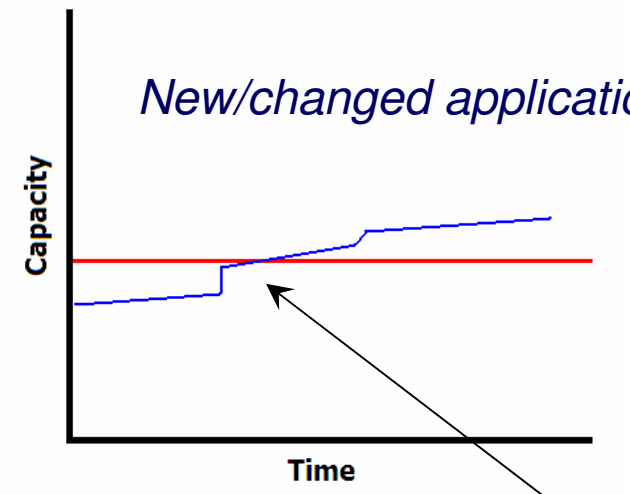

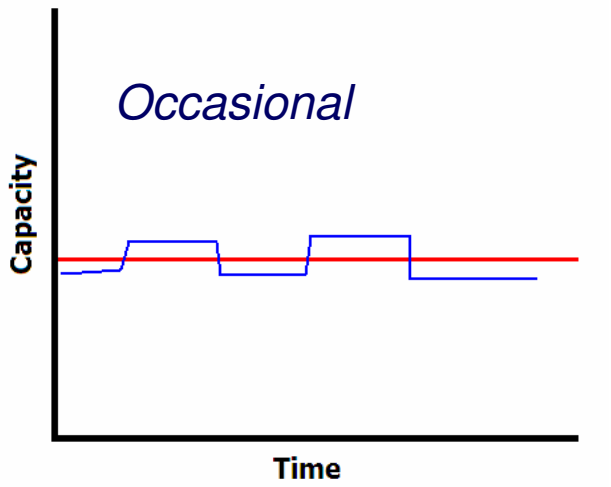

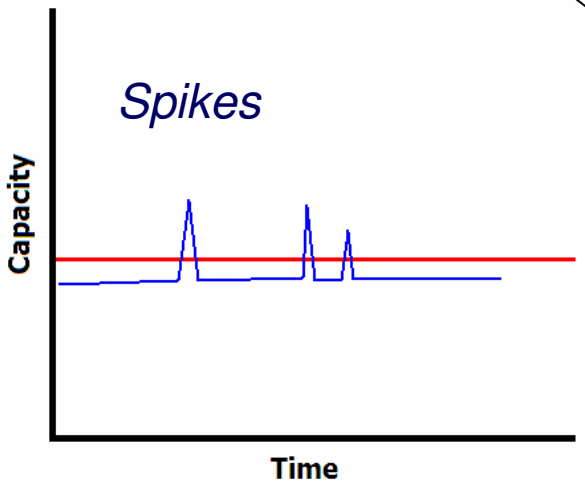

Installation success
But…
- Data grows with use
- Widespread adoption

Classic TIP-TOP

Usage and Growth

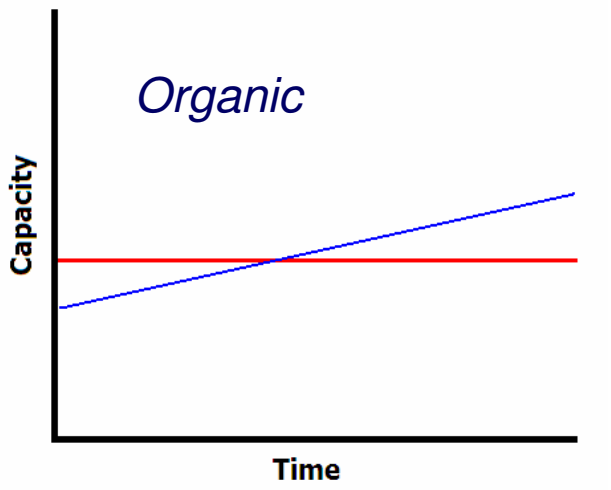

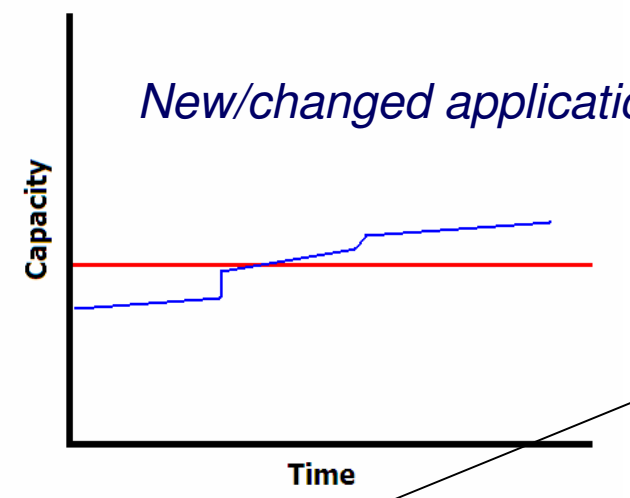

Was everything tested with this change?

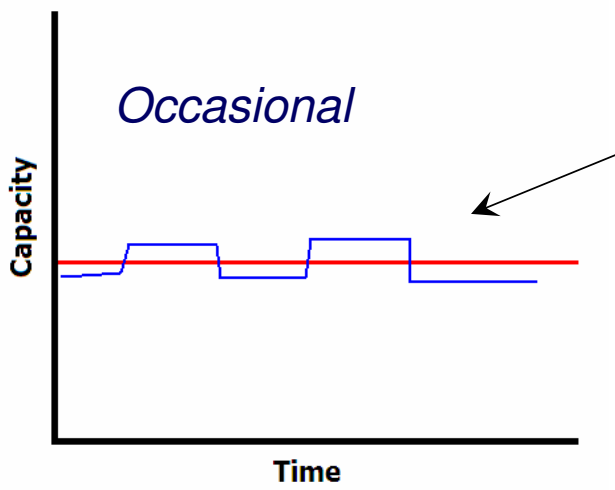

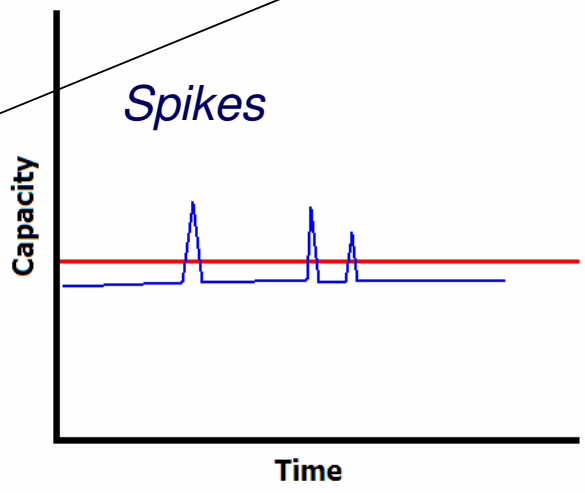

Usage and Growth

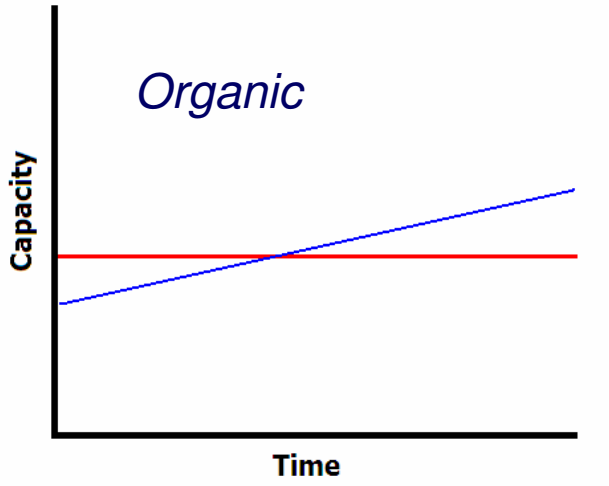

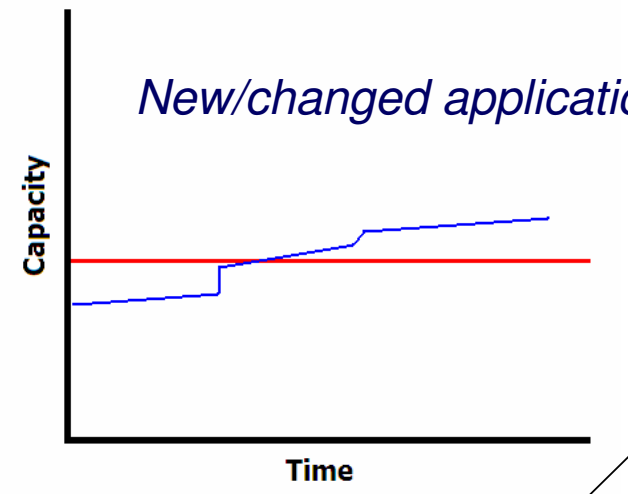

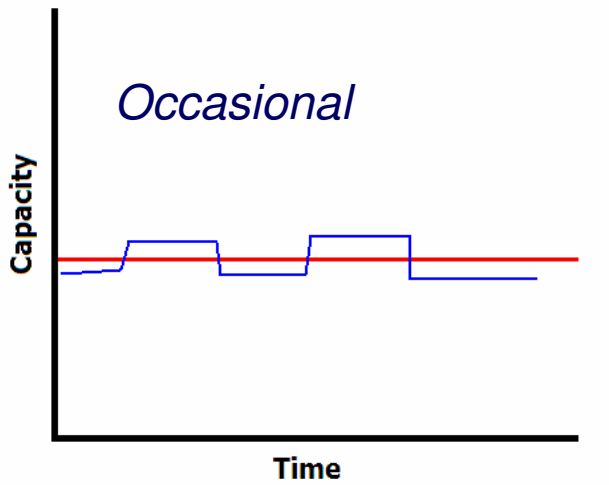

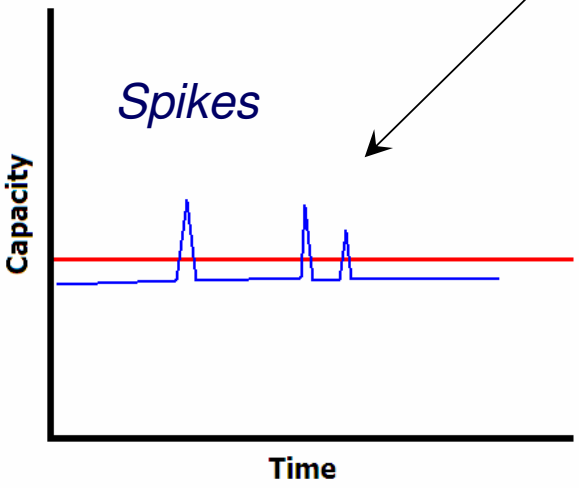

Most likely causes:
- Ad-hoc queries…
- …in PRODUCTION

((Yikes!))

Secondary cause:
- Competition for resources

See Next Chart!!!!

Spike analysis technique

- Look for the spikes in anything and do cross-comparisons
- Adjust one setting at a time to correct
  - SQL spike:
    Correct (tune) SQL and cross collaborate with table activity
  - I/O spike:
    Review type of activity, cross collaborate with bufferpool and SQL activity
  - Bufferpool spike:
    Review type of spike, cross collaborate with I/O and SQL activity
  - Memory spike:
    Cross compare with bufferpool, I/O and SQL activity
  - Locking issues:
    Review application for commit behavior, tune SQL, change config parameters.

On to the Playbooks…

Non-spike analysis technique

- Best used with I/O (SMF type 42)
- Look at calculated total response time and sort descending
  I/O AVERAGE RESPONSE TIME x TOTAL NUMBER of I/Os
- Top 1,000 highest times
- Find Low, medium, and high ranges
- Isolate as…
  - Steady
  - Occasional
  - Spikes

Next slide shows example

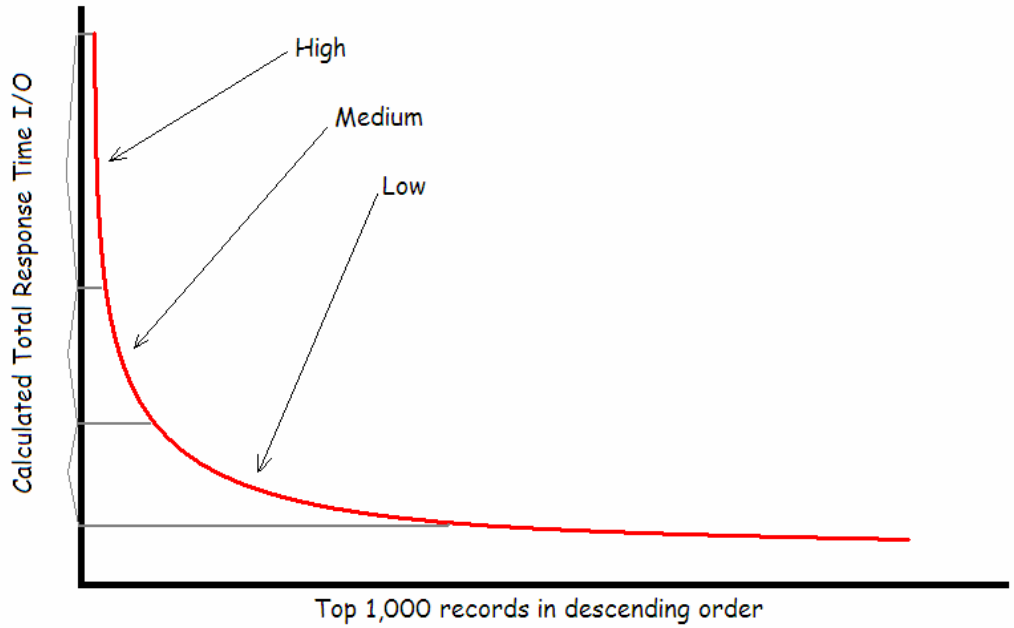

Where is the real bottleneck?

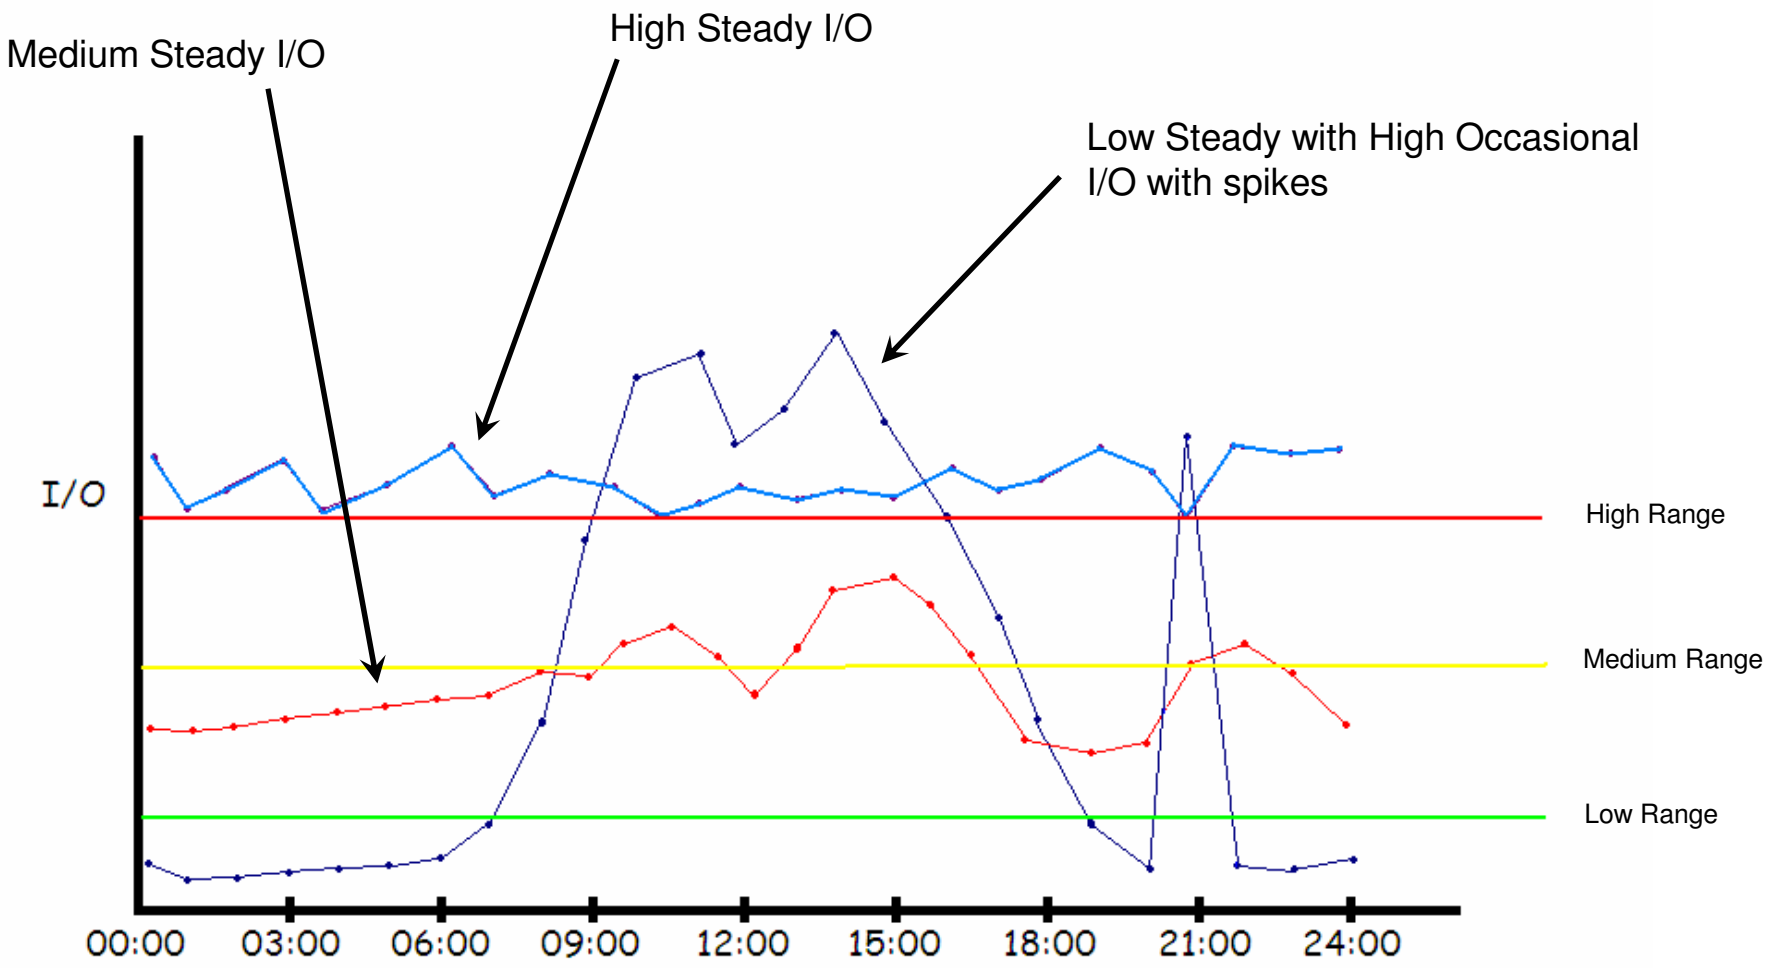

I/O usage analysis

DB	TSP/IDX	Extreme High I/O	Steady	Occasional	Spikes	Table Creator	Table Name	Indexspace	Index Creator	Index Name	Buffer Pool	Cardinality	STATS Date	Partition
EBKD3P03	TSKD0044	0	3	2	0	EBKDP03	ET_CPN				BP7	676,655,745	12/9/2010	40
EBKD3P03	TSKD0063	0	3	2	0	EBKDP03	ET_CPN_CRR_ROLE				BP7	980,436,569	10/25/2010	40
EBKD3P03	TSKD0042	0	3	0	0	EBKDP03	ET_CPN_ST				BP7	1,190,174,040	12/10/2010	100
EBKD3P03	IX1KD055	0	2	0	0	EBKDP03	ET_TKDC_FARE_CALTN	IX1KD055	EBKDP03	IX1KD055	BP6	202,351,254	10/26/2010	0
EBKD3P03	IX1KD037	0	2	3	0	EBKDP03	ET_TKT_DOC	IX1KD037	EBKDP03	IX1KD037	BP6	317,968,414	10/21/2010	0
EBKD3P03	TSKD0037	0	3	2	0	EBKDP03	ET_TKT_DOC				BP7	317,939,820	10/21/2010	40
EBKD3P03	IX1KD306	0	3	0	0	EBKDP03	FL_REV_ACCT_CTRL	IX1KD306	EBKDP03	IX1KD306	BP6	69,079,979	1/3/2011	0
EBKD3P03	IX1KD302	0	3	0	0	EBKDP03	FLT_LE					6,155	1/3/2011	0
EBKD3P03	IXCKD305	0	2	0	0	EBKDP03	FLT_LE					6,769	1/3/2011	0
EBKD3P03	IX1KD300	0	3	0	0	EBKDP03	LPX_FL					1,965	1/3/2011	0
EBKD3P03	IX4KD300	0	3	0	0	EBKDP03	LPX_FL					6,044	1/3/2011	0
EBKD3P03	TSKD0300	0	3	0	0	EBKDP03	LPX_FL					5,200	1/3/2011	40
EBKD3P03	IX1KD201	0	2.5	0	0	EBKDP03	MT_CPN_ATP_SVC_FEE	IX1KD201	EBKDP03	IX1KD201	BP6	275,107,100	7/13/2010	0
EBKD3P03	IX1KD202	0	3	0	0	EBKDP03	MT_CPN_ATPSF_DTL	IX1KD202	EBKDP03	IX1KD202	BP6	275,699,343	7/12/2010	0
EBKD3P03	TSKD0275	0	3	2	0	EBKDP03	MT_CPN_ST				BP7	2,668,108,230	8/31/2010	500
EBKD3P03	IX1KD172	1	1.5	1	1	EBKDP03	MT_DOTRP_TKT_DOC	IX1KD172	EBKDP03	IX1KD172	BP6	16,391,141	7/8/2010	0
EBKD3P03	IX1KD119	0	3	0	0	EBKDP03	MT_TKDC_BKG_RLOCR	IX1KD119	EBKDP03	IX1KD119	BP6	523,230,843	2/24/2010	0
EBKD3P03	IX1KD134	0	3	0	0	EBKDP03	MT_TKDC_PMT_COLTN	IX1KD134	EBKDP03	IX1KD134	BP6	384,716,014	7/13/2010	0
EBKD3P03	TSKD0372	0	1.5	0	0	EBKDP03	MT_TKT_DOC				BP7	507,245,026	10/15/2010	200
EBKD3P03	IX1KD372	0	2.5	0	0	EBKDP03	MT_TKT_DOC	IX1KD372	EBKDP03	IX1KD372	BP6	507,433,509	10/15/2010	0
EBKD3P03	IX3KD372	0	2.5	0	0	EBKDP03	MT_TKT_DOC	IX3KD372	EBKDP03	IX3KD372	BP6	390,832,771	10/15/2010	0
EBKD3P03	TSKD0804	0	2	0	0	EBKDP03	MT_TKT_DOC_CHRG				BP7	1,864,701,252	7/25/2009	200
EBKD3P03	IX1KD804	0	2.5	0	0	EBKDP03	MT_TKT_DOC_CHRG	IX1KD804	EBKDP03	IX1KD804	BP6	2,347,664,567	9/30/2010	0
EBKD3P03	IX1KD121	0	3	0	0	EBKDP03	MT_TKT_DOC_CMSN	IX1KD121	EBKDP03	IX1KD121	BP6	239,497,063	9/1/2010	0
EBKD3P03	IX2KD832	0	3	0	0	EBKDP03	MT_TKT_DOC_MSG	IX2KD832	EBKDP03	IX2KD832	BP6	96,127,992	12/8/2010	0
EBKD3P03	IX1KD117	0	3	0	0	EBKDP03	MT_TKT_DOC_PMT	IX1KD117	EBKDP03	IX1KD117	BP6	374,840,724	7/13/2010	0
EBKD3P03	IX2KD117	0	3	2	0	EBKDP03	MT_TKT_DOC_PMT	IX2KD117	EBKDP03	IX2KD117	BP6	73,874,562	7/13/2010	0
EBKD3P03	IX3KD117	0	3	0	0	EBKDP03	MT_TKT_DOC_PMT	IX3KD117	EBKDP03	IX3KD117	BP6	156,695	7/13/2010	0
EBKD3P03	TSKD0117	0	3	0	0	EBKDP03	MT_TKT_DOC_PMT				BP7	374,816,927	7/13/2010	40
EBKD3P03	IX1KD105	0	2.5	2	1	EBKDP03	MT_TKT_DOC_POS_SRC	IX1KD105	EBKDP03	IX1KD105	BP6	963,907,962	9/29/2010	0
EBKD3P03	TSKD0105	1	1.5	1	1	EBKDP03	MT_TKT_DOC_POS_SRC				BP7	729,212,405	9/3/2009	40
EBKD3P03	IX1KD071	0	2.5	0	0	EBKDP03	MT_TKT_DOC_TPBK	IX1KD071	EBKDP03	IX1KD071	BP6	1,234,623,831	9/29/2010	0
EBKD3P03	TSKD0071	1	1	0	1	EBKDP03	MT_TKT_DOC_TPBK				BP7	952,462,782	9/3/2009	40
EBKZ1P01	TSKZ0802	0	3	0	0	EBKDP03	RPT_LPX_FLRCTRL				BP7	42,661,347	1/18/2011	40

Buffer pools used:
Overuse found on BP6 and BP7

High-usage / spiked repository datasets

Heavy-usage tables

Performance tuning playbooks DB2 for z/OS

New application on a new subsystem:

- zParm settings
- Above/below the line memory allocations
- WLM settings for service class pertaining to DB2
- SQL
  - If less than 5 SQLs running slow, tune the SQL
  - If more than 5 SQLs running slow, look at the I/O and bufferpool
- Locking behavior

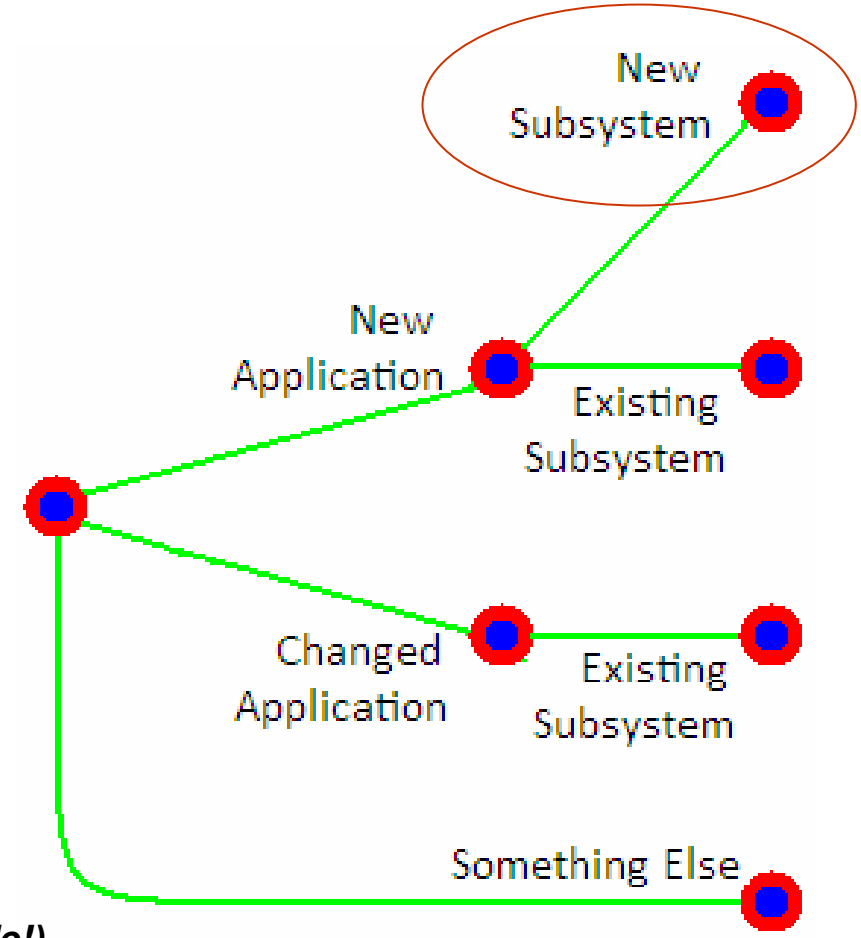

Jeffy's Rule-of-10 (next slide!)

Sidebar: Why 5 SQL statements?

- Jeffy's rule of "10"
  - Most I/O is caused by no more than 10 tables/indexes
  - No more than 10% of all I/O data show high usage sync activity
  - Most applications have a top 10 list of poor performing SQL
- Of the top 10 poor performing SQL…
  - There are 2 to 3 versions running concurrently
  - Therefore, average of 5 SQL statements

Back to the Playbooks…

Performance tuning playbooks DB2 for z/OS

New application on an existing subsystem hosting applications:

- Activity in the buffer pools and I/O
- SQL
  - If less than 5 SQLs running slow, tune the SQL
  - If more than 5 SQLs running slow, look at the I/O and bufferpool
- Locking behavior
- Capacity like tablespace growth

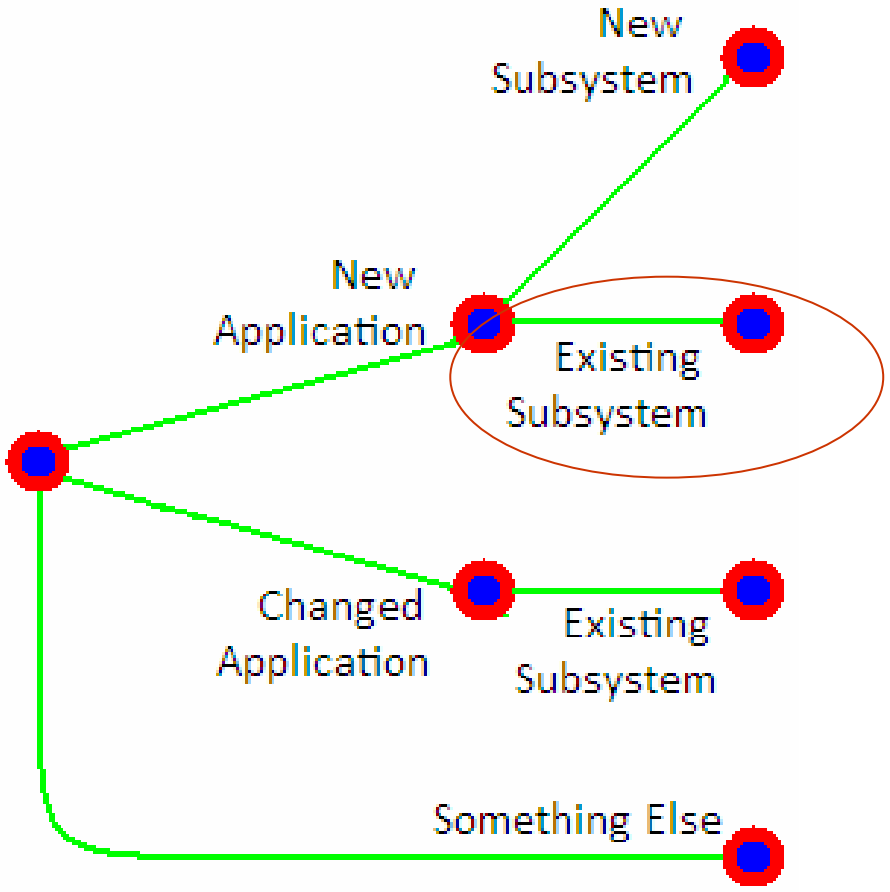

Performance tuning playbooks DB2 for z/OS

Change to an existing application:

- SQL
  - If less than 5 SQLs running slow, tune the SQL
  - If more than 5 SQLs running slow, look at the I/O and bufferpool
- Locking behavior
- Activity on the subsystem

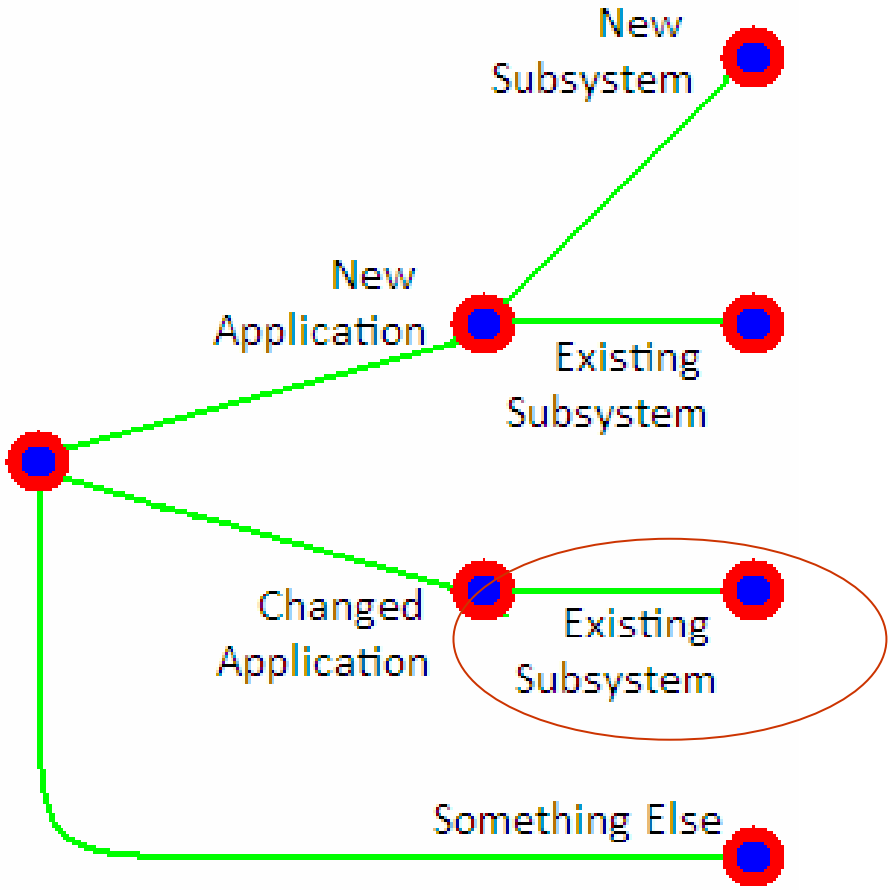

Performance tuning playbooks DB2 for z/OS

No perceptible change whatsoever (application and environment)

- Organic growth
- Something unexpected is running

How did “you” hear about the problem?

- Tool showed problem
- Identified by an end-user or application programmer

What should be researched?

- Check for something unexpected additional (Heavy utilities running)
- SQL to find the long running queries, then
- RUNSTATS on tables identified in the long running queries
- Followed by a health check across-the-board
- Trend analysis and capacity planning like table space growth

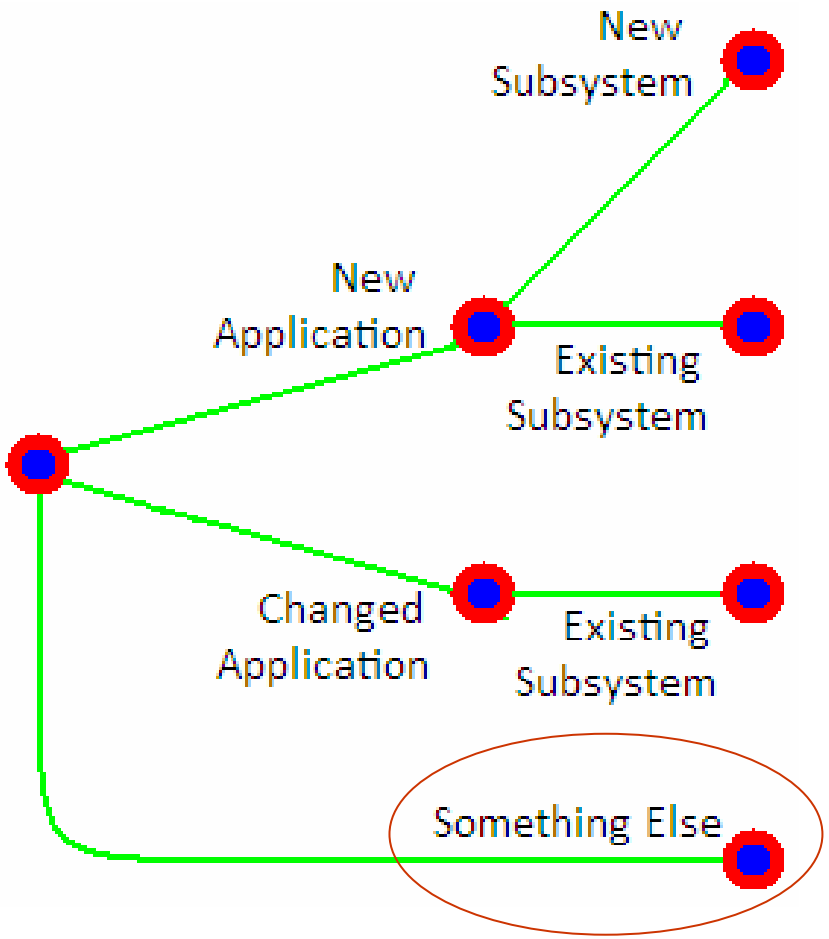

A Thoughtful, Enlightened Strategy

Structured, methodical, closed-loop approach

- Be prepared! Understand how the system works when things are well
- Look at high-level performance symptoms with tooling
  - Optim Performance Manager
  - Omegamon XE
  - Iostat / db2look
  - Resource Management Facility
  - Other tools
- "Divide and conquer" the problem
  - What causes do the symptoms indicate?
  - What do they rule out?
- Make one or more hypotheses
- Important: change one thing at a time!

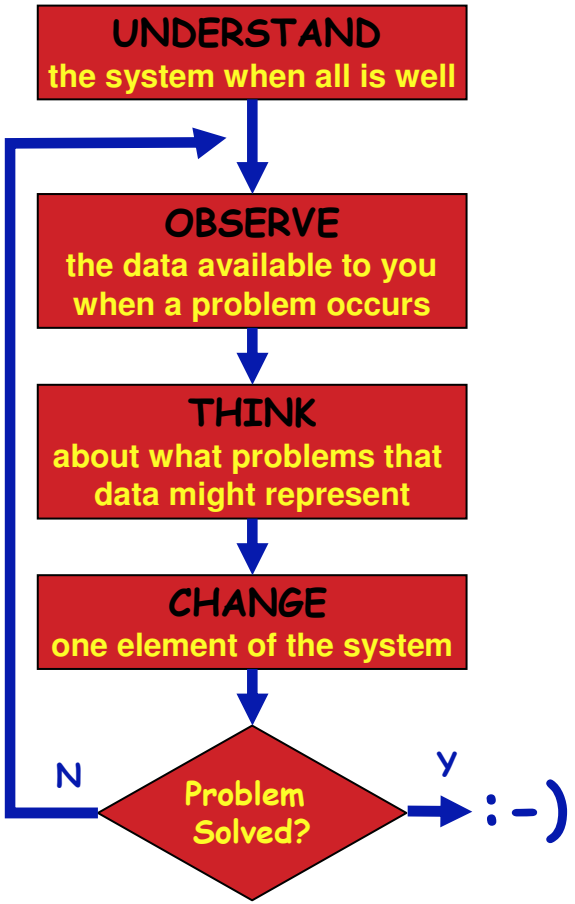

Agenda

- **The politics of performance**
  - The data base administrator role
  - Management expectations
  - DBA expectations
- **Performance Tuning**
  - Graphs and Trends
  - Playbooks
- **DB2 biggest performance issues**
- **Tooling**
- **Performance Touch Points with DB2 for z/OS**

DB2 biggest performance issues

Underestimating the effects of …

- the DB2 setup
- the I/O (DB2 for z/OS)
- poorly written SQL
- the workload

DB2 setup - Buffer pool strategy

- All I/O is buffered, no direct reads
  - DB2 does many operations autonomously (Predictable)
  - Unexpected things in the buffer pool (Unpredictable)
- DB2 has to “drop what it is doing” to perform a read
- Contention
  - Buffer pool too small
  - Competition within the pool
  - Threshold settings not reflective of usage

DB2 9 for z/OS: Buffer Pool Monitoring and Tuning
http://www.redbooks.ibm.com/redpieces/abstracts/redp4604.html

DB2 setup - Connections

- Know thy traffic
  - Application traffic
  - Server-to-server traffic and latency
  - Workload and usage
- Contention
  - Number of hops
  - Latency
  - Number of connections
  - Settings not reflective of usage

Underestimating the I/O – High usage objects/Logical Control Unit

- Highly used DB2 objects
  - Table spaces
  - Index spaces
- Highly used disk controllers
- Combination of both

DB2 9 for z/OS and Storage Management (SG24-7823-00)
http://www.redbooks.ibm.com/abstracts/sg247823.html?Open

Poorly written SQL and the workload

- Dynamic SQL challenge
- Need to find the “bad” SQL
- How do you know the “good” SQL from the “bad” SQL
- Situation of execution
  - One-off inconsistent
  - Same time / same day-of-week but different day
  - Did it just start happening without change
- Dynamic SQL – Finding the application

Sequence to the Hunt for Bad SQL

- Capturing the SQL
  - On-line monitors
  - If that doesn't work, set a trace
- Analyzing/Tuning the SQL
- Getting the SQL changed
  - Hopefully there is a feedback/code review process
  - Regression test "sandbox"

Omegamon XE or equivalent

Scanning the dynamic statement cache:
- Data Studio Standalone
- Optim Query Tuner
- Query Workload Tuner

Optim Query Tuner for DB2 for z/OS or equivalent

Data Studio Standalone:
- Query formatter
- Access plan graph
- Statistics advisor

Optim Performance Manager

pureQuery

- Static binding plus WLM

You have the SQL, now what?

1. Minimize I/O
   - Find in the buffer pool
2. Minimize synchronous activity
   - Have DB2 predictably prefetch into the buffer pool
3. Combine SQL operations to minimize CPU instructions
   - Find an SQL guru
   - Use tooling
     - OQT - Tune SQL pre-production while costs and impact are low
     - OQWT - Optimize workload for peak performance

Agenda

- **The politics of performance**
  - The data base administrator role
  - Management expectations
  - DBA expectations
- **Performance Tuning**
  - Graphs and Trends
  - Playbooks
- **DB2 biggest performance issues**
- **Tooling**
- **Performance Touch Points with DB2 for z/OS**

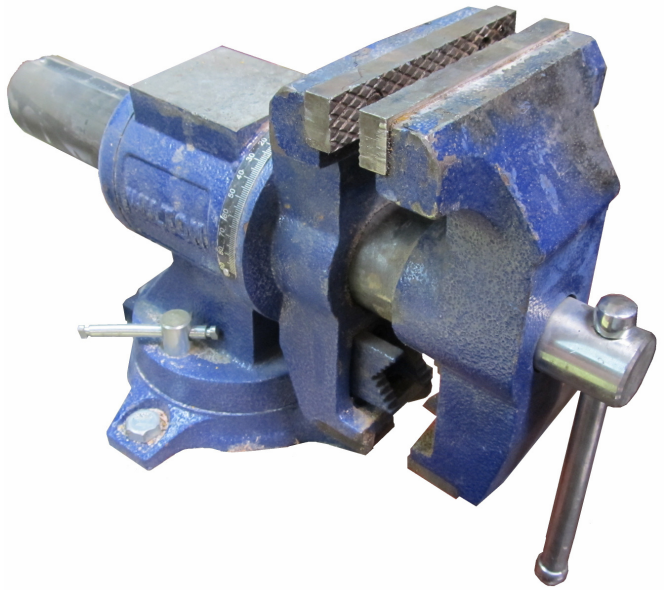

Tooling in Performance Management

DB2 for z/OS

OMEGAMON XE DB2 Performance Expert
OMEGAMON XE DB2 Performance Monitor
DB2 Buffer Pool Analyzer
Optim Query [Workload] Tuner
DB2 Query Monitor
DB2 Performance Toolkit SAP Edition

DB2 for LUW

Optim Performance Manager
Optim Query Tuner

IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows
http://www.redbooks.ibm.com/redpieces/abstracts/sg247925.html?Open

DB2 Developer Workbench vs. Data Studio

before *now*

IBM DB2 Developer Workbench V9.1

- SQL Query Editor
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

Data Studio is a full replacement of DB2 Developer Workbench plus much more

- ***DB2 for Linux, Unix, Windows v8.x, v9.x***
- ***DB2 for z/OS v7, v8, v9***
- ***DB2 for i5/OS v5r2, v5r3, v5r4***
- ***Informix Dynamic Server (IDS) v9.x, v10.x, v11***

IBM Data Studio V2.2

No-charge

- Integrated Query Editor – SQL + XQuery
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

- ER Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support

Data Studio V2.2

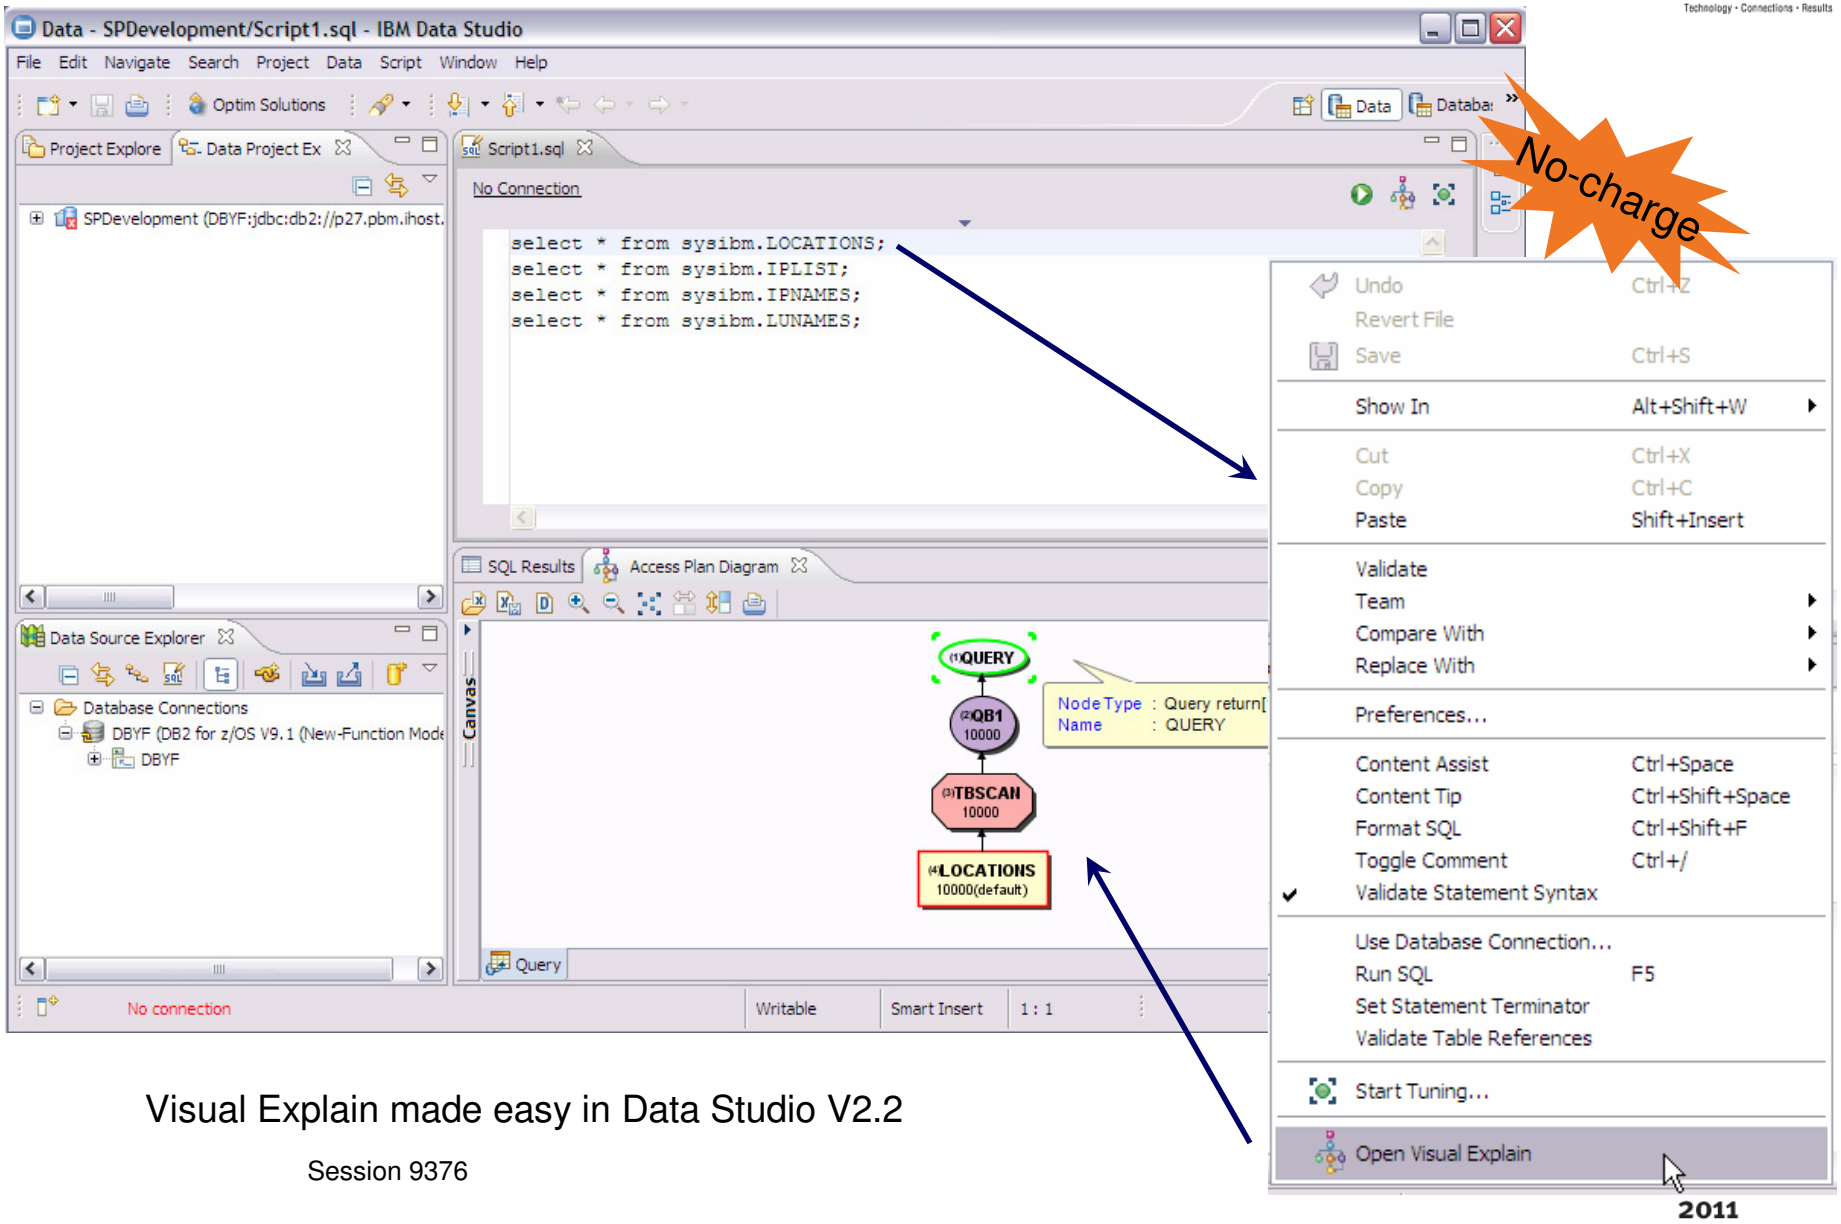

Visual Explain made easy in Data Studio V2.2

Session 9376

Other tooling
DB2 for z/OS high usage objects

- Methods
  - RMF I/O activity by VOLSER reports
  - SMF type 42 records (SMS: subtype 6)
- Observation: There are, on average, 10 tables accessed exponentially above all other tables

SMF Type 42 Parser for zOS
https://www.ibm.com/developerworks/mydeveloperworks/files/app?lang=en#/person/100000P902/file/402e49fe-e5bf-496f-8e4f-0bbcde625f34

Agenda

- **The politics of performance**
  - The data base administrator role
  - Management expectations
  - DBA expectations
- **Performance Tuning**
  - Graphs and Trends
  - Playbooks
- **DB2 biggest performance issues**
- **Tooling**
- **Performance Touch Points with DB2 for z/OS**

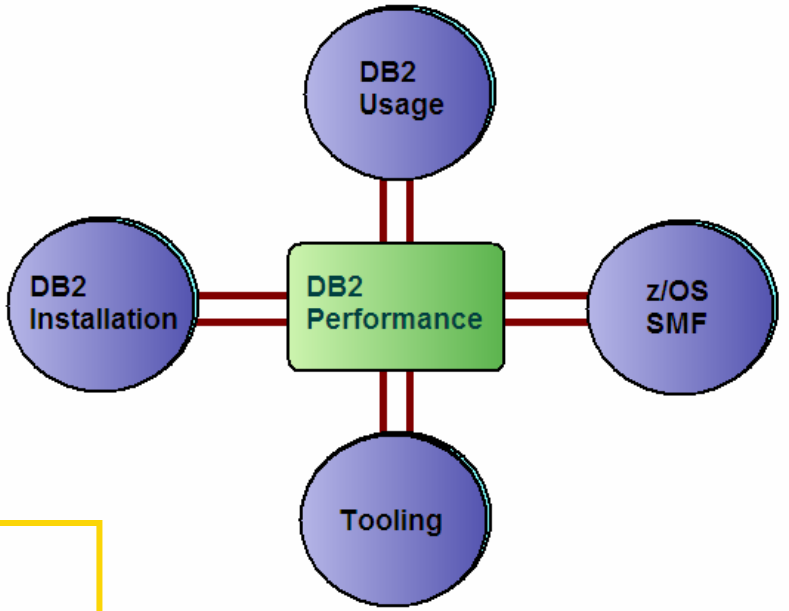

Performance Touch Points with DB2 for z/OS

DB2 Installation

- SMFACCT=(1,2,3) and SMFSTAT=(1,3,4)
- Log sizes - dual for onsite/offsite

DB2 Usage

- Table spaces sized for manageability
- Buffer pool strategy
- Understand the usage for any new feature
  - PGFIX
  - Use sliding scale if not constrained by DASD

z/OS

- SMF Records
  - Use appropriate interval
  - Type 100:102
  - Type 42
  - Need type 30 records with intervals

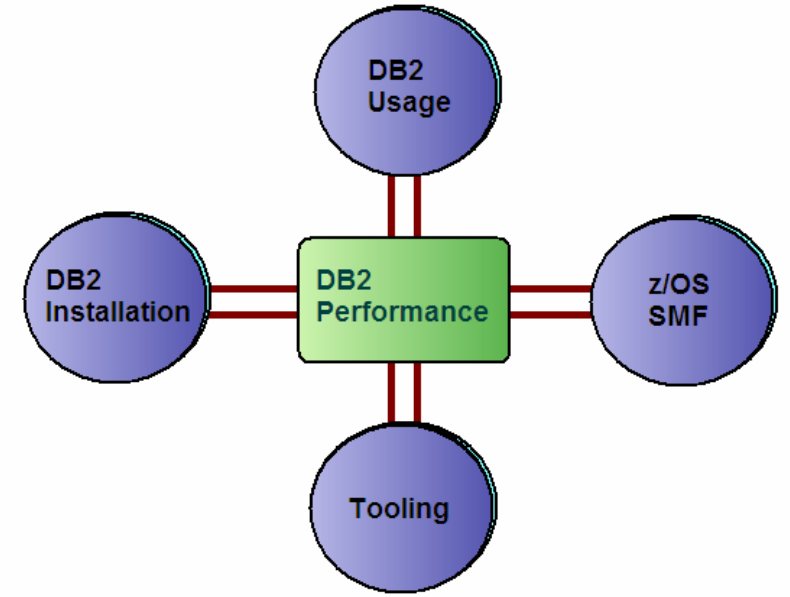

Performance Touch Points with DB2 for z/OS

Tooling

- Omegamon XE for DB2 or equivalent
- DB2PE or equivalent
- Data Studio (OSC replacement)
- Optim Query Tuner
- Optim Workload Query Tuner

Optional Tooling

- RMF or equivalent
  - RMF Spreadsheet Reporter
- Optim Performance Manager (DB2 for LUW on zLinux)

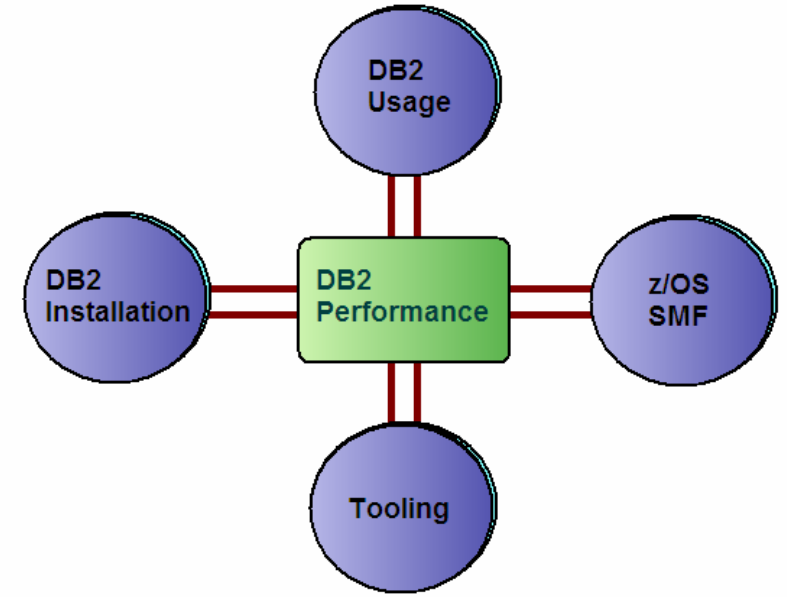

Performance Touch Points with DB2 for z/OS – Links to tools

IBM Data Studio and pureQuery

http://www-01.ibm.com/software/data/optim/

DB2 Accessories Suite for z/OS

http://www-01.ibm.com/software/data/db2imstools/db2tools/accessories-suite/

SMF Type 42 Parser for z/OS

https://www.ibm.com/developerworks/mydeveloperworks/files/app?lang=en#/person/100000P902/file/402e49fe-e5bf-496f-8e4f-0bbcde625f34

IBM Tivoli Monitoring zOMEGAMON and Related Products Best Practices Informational links

http://www.ibm.com/developerworks/wikis/download/attachments/141165182/Best+Practices+Links.pdf?version=1

Performance Touch Points with DB2 for z/OS - Training

- CV960 - DB2 9 for z/OS Application Performance and Tuning
- CV950 - DB2 9 for z/OS System Performance Analysis and Tuning
- ES545 - Basic z/OS Tuning Using the Workload Manager (WLM)

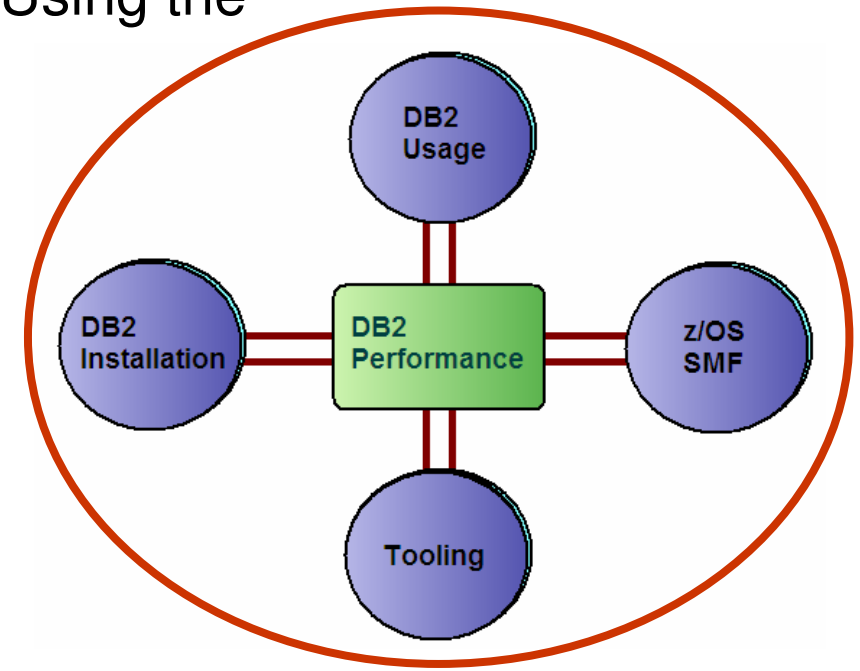

In Summary

- It is possible to performance tune ad infinitum, ad nauseam...Don't!
- Tune to the level of what you, your end-users, and the company can live with
- Meet your service level agreements and service level objectives
- Remember these 3 guidelines:
  - Bottlenecks can be either Memory, I/O, or Processor but in most cases it is a combination of factors
  - The limit of any machine is ultimately a bottleneck by definition
  - Workload and performance management is the art of juggling what you can live with

Need a jump-start?

Skill gap?

Need it yesterday?

IBM Systems Lab Services and Training

www.ibm.com/systems/services/labservices

stgls@us.ibm.com ● Systemz@us.ibm.com

Thank You!!!

Jeff M. Sullivan
jeffsull@us.ibm.com

Session 9376

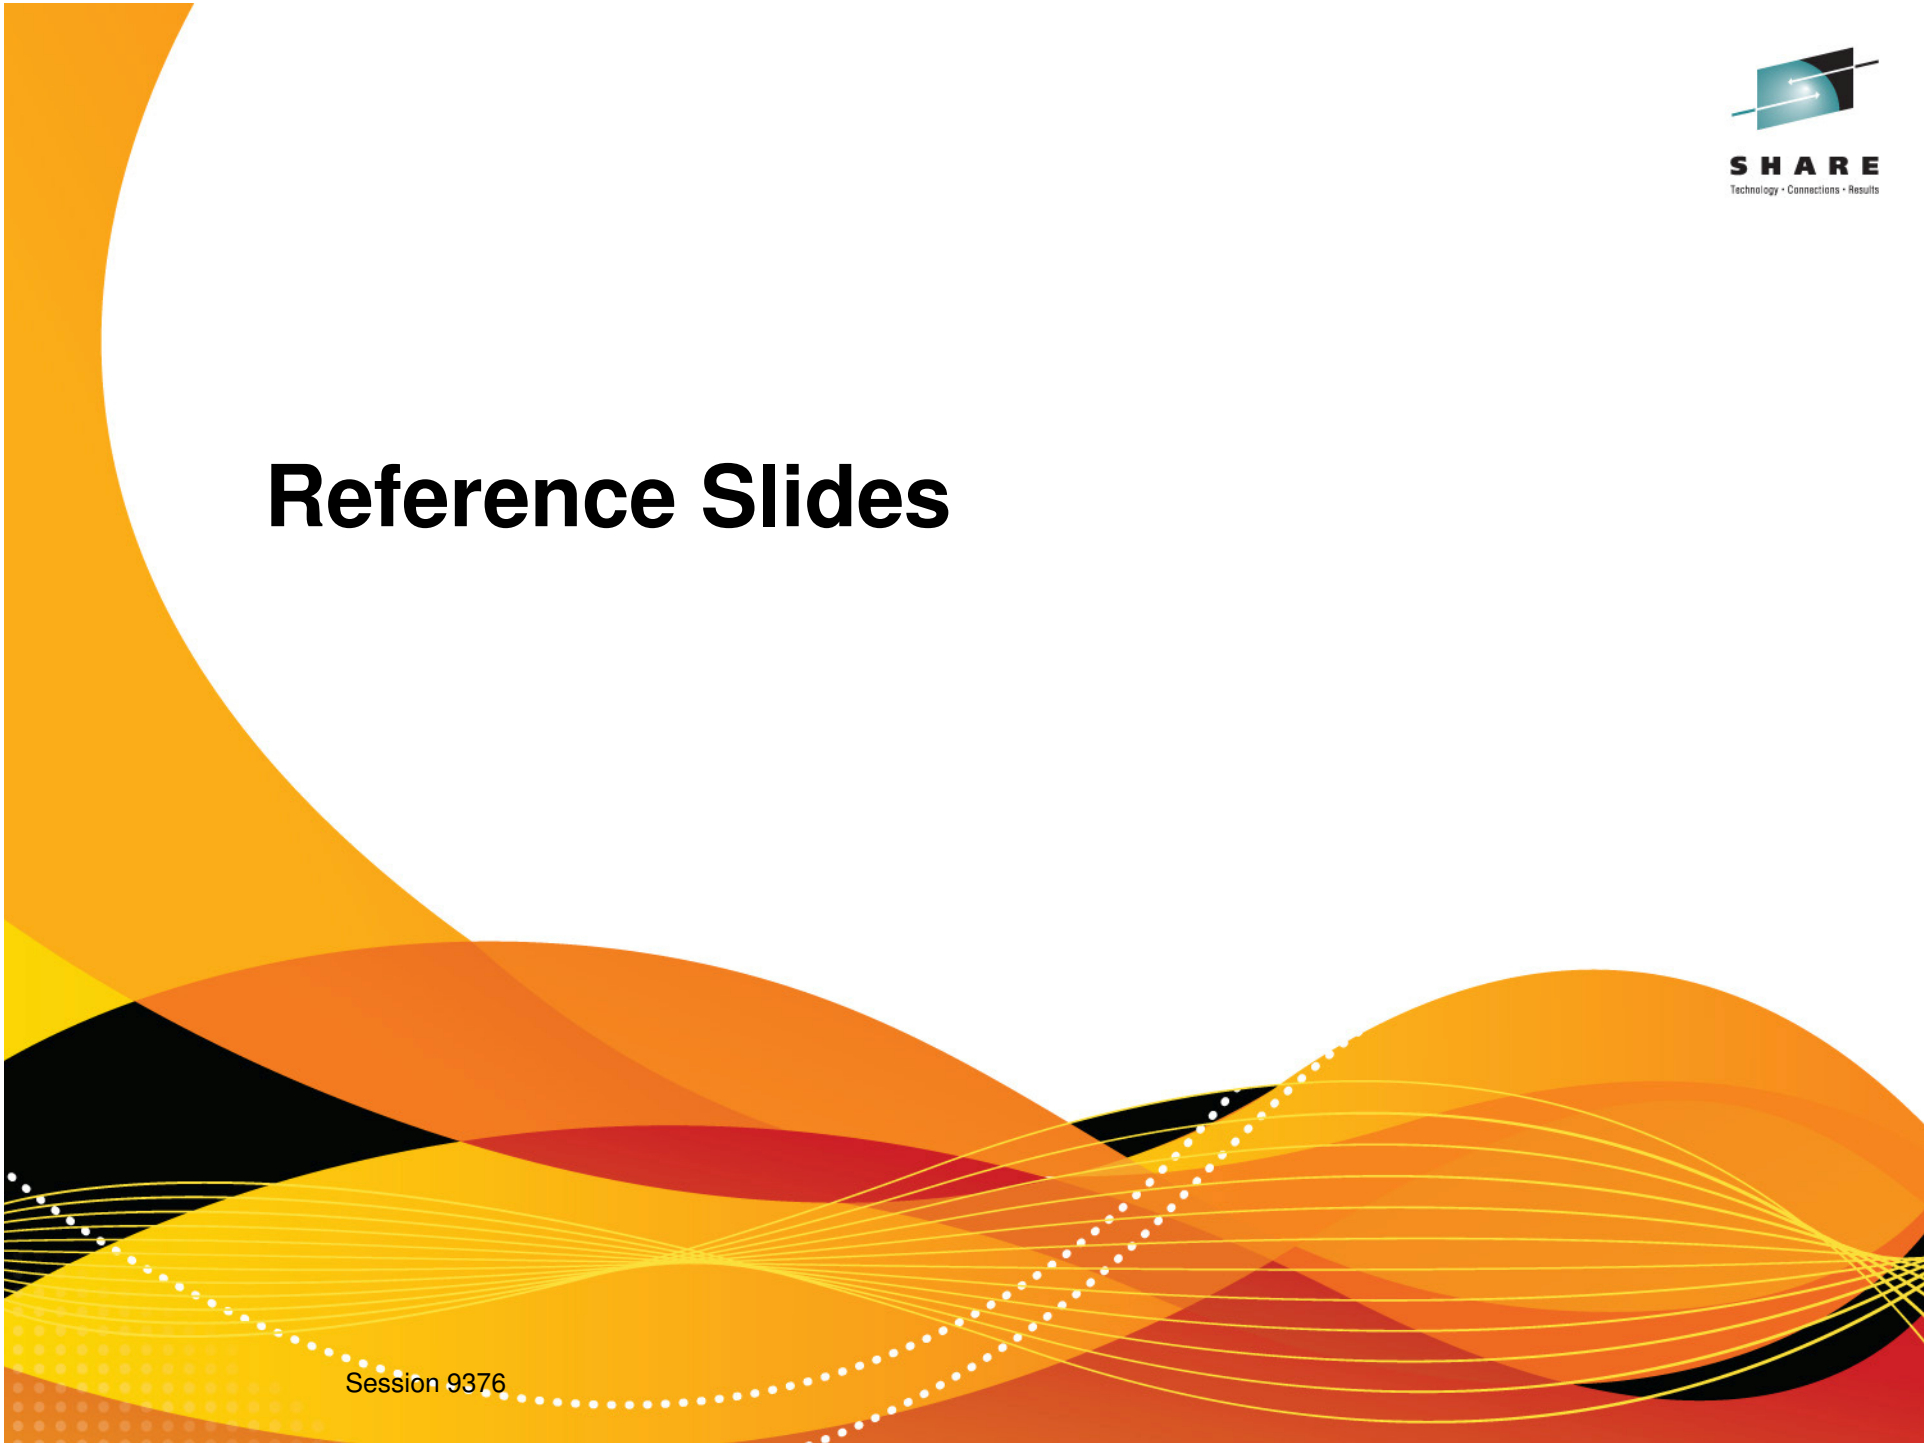

Reference Slides

The data base administrator - Job description

Protection, Availability, and Speed are the goals.

- **Protection of the company's data assets**
  - Adequate backup and recovery
  - Retention of logging media
  - Ensuring adequate storage and retention of logs and backups
  - Performing disaster recovery scenarios
  - Making sure any recovery performed is done efficiently with little data loss and little outage
  - Jealous data protection and being involved with any change in the system
- **Availability of the company's data assets on an as-required basis**
  - Ensuring the correct security roles and responsibilities are delivered to the correct organizations
  - Reviewing that all required paths into and out of the DBMS are efficient with such things as connections, minimal server "hops", prevention of any outage (including failover of a server in the chain)
  - Making sure all software is up-to-date
  - Performing proactive routine maintenance including REORGS, RUNSTATS, space usage (organic growth), and review of DBMS usage for memory and I/O
- **Fast delivery of the company's data assets to those authorized**
  - Involvement with creation of and meeting service level objectives
  - Reviewing (and creating) table, view, access paths and access strategies with application personal
  - Ensure that the DBMS, the DBMS servers, and the data is delivered as quickly and efficiently as possible
- **Provide analysis on problems as required**

DB2 for z/OS
Buffer Pool Recommendations

- A buffer pool strategy should…
  - Separate, at a minimum, by tablespaces and indexes
  - Separate highly accessed tablespaces and indexes from the less active tablespaces and indexes
  - Separate by random (online, OLTP) versus sequential access
  - Separate by size
- A good subsystem buffer pool strategy should…
  - Put work tablespaces into their own pool
  - Put temporary tablespaces into their own pool
  - DB2 Catalog into their own pool

DB2 9 for z/OS: Buffer Pool Monitoring and Tuning
http://www.redbooks.ibm.com/redpieces/abstracts/redp4604.html

Recommended DB2 buffer pool Strategy

- 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