

How a minor change had a major impact

Siegfried Fürst - SOFTWARE ENGINEERING GmbH



Like clockwork, every Thursday evening, an international supermarket conglomerate systematically “pushes” the latest weekly updates into Production.

These updates consist of approved program changes and functional upgrades for the data center of two different supermarket chains. During one of these routine cycles in mid July, an unintentional change was introduced which had far-reaching consequences.

Early Friday morning, reports of users experiencing problems began to trickle in: they were having issues with the login to the merchandising system—which runs under DB2 z/OS. The initial complaints were about delays in the login process. Later, complaints of poor performance in general were reported too.

A cursory check of the system revealed excessive CPU consumption. The initial conclusion was that the Thursday evening update had somehow led to the problems.

As would be expected under these circumstances, the initial response was to “rollback” the prior day’s release. The feeling among the IT staff was that a program “bug” had somehow been unintentionally introduced.

Unfortunately, even after the data had been restored, the problem persisted further. The login process and general usage of the system was still unacceptable to the user base. The regional branch managers were none too happy.

The IT team continued with their resolution endeavors and a further analysis was undertaken, with a complete review of the IBM Websphere Application Servers, and all of their Java components.

The assessment of all programs running at the time revealed one SQL statement in particular that really stood out. However, after further analysis, it was found that even though this SQL was related to the long wait times during user login, it was not involved in the Thursday evening release. The DBA staff were able to confirm that the execution of this SQL statement was also coincidental to the observed “spikes” in CPU consumption.



By pure coincidence, around this time, a performance project team had been looking into the issue of “nudging” access paths and their general stability. They were working together with the company SOFTWARE ENGINEERING GmbH (SE) to automate their complex requirements and benefit from SE’s long experience with large IBM DB2 z/OS customers around the globe.

SE had installed products that were using fast and powerful methods, including the new DB2 10 IFCID technology and access path control.

Management agreed to involve the performance project team in the problem diagnosis to help pinpoint the issues faster.



Enter SQL WorkloadExpert (aka WLX) and Bind ImpactExpert (BIX). Using these new, innovative solutions, the project team was able to very quickly confirm the “suspect” SQL as the culprit. Continuing with the tools, both teams were able to see that there had been a dramatic change to the access path. It was soon determined that the performance had degraded by a factor of 230!

As WLX captures all historical information on SQL and the associated access paths, users are able to view “point-in-time” or “point-to-point” scenarios.

In this case, the DBA’s were able to a) analyze the stored Dynamic Statement Cache (DSC) and confirm the offending SQL statement was one of the top consumers at the point in time, and b) correctly confirm that the offending statement was previously seen as operating within acceptable limits.

Armed with this information from WLX and BIX, the DBA team was then able to quickly pinpoint the reason for the dramatic degradation: On the Thursday evening in question, a RUNSTATS utility had been executed across the tables used by the offending SQL statement. Following this, new catalog statistics had been created, and the DB2 Optimizer had, unfortunately, selected a slightly different but dramatically worse access path than before.

Using the SE Tools, the team was easily able to compare “before” and “after” access paths side by side: Following the update of the catalog data on Thursday evening, the DB2 optimizer had chosen WW09_ORDER as the initial table, whereas previously WWS18_ORDER_STATUS had been used. (Please refer to the PLAN_Tables in Figure 2 and Figure 3 on pages 6 and 7 respectively for more details.)

The company employs more than 400,000 people in the German-speaking stores alone.

They run a three-tier 16+ way data-sharing environment with over 50,000 MIPS and DB2 z/OS on the mainframe.

Over 30 logistics centers deliver “Just in Time” to around 3,500 supermarkets. The supply chain is extremely important to the business for guaranteeing product freshness and availability, and it is also a key factor for competitive pricing.

In addition, BIX showed in the SQL Statement Cache after the Thursday evening update, that the SQL statement had a changed access path. Prior to this, the analysis showed that the statement was unchanged.

The small change with the large impact is highlighted in red in the PLAN_TABLE comparison below.

Access path OLD -----								! Access path NEW -----							
!								!							
TABLE	QB	PN	AC	MA	ME	IX	PR	TABLE	QB	PN	AC	MA	ME	IX	PR
INDEX			TY	CO	TH	ON	FT	INDEX			TY	CO	TH	ON	FT
RES	1	1	R	0	0	N	S	RES	1	1	R	0	0	N	S
	1	2		0	3	N			1	2		0	3	N	
WWS18_ORDER_STATUS	3	1	I	1	0	Y	S	WWS09_ORDER	3	1	I	1	0	N	S
WWS18X01								WWS09X01							
.								!	.						
.								!	.						
.								!	.						
Milliseconds:				25				! Milliseconds:					38		
Serviceunits:				1269				! Serviceunits:					1861		

The simple solution to the problem, was to create an additional Index for the table WWS18_ORDER_STATUS, in order to nudge the DB2 Optimizer into choosing that table as the initial table.

The new index was quickly created and RUNSTATED in Production in a matter of minutes, then everything starting running as normal. They were then able to restart all the WebSphere servers and all was well.



In Summary

The preceding (true!) story illustrates how just one small change can have dire consequences for the business.

Simply by following their routine procedures, unsuspecting teams were able to wreak havoc. This can happen anywhere.

In this case, the problems were 3-fold:

1. On Thursday night, RUNSTATS were run and the Optimizer chose a very slightly different, but nonetheless extremely poorly-performing access path.
2. Because historical SQL data wasn't readily available to the DBAs, the consequences of this were not immediately recognized.
3. As, we suspect, with most sites, time was wasted in a) rolling back the release, and b) attempting to remediate by running REORG and RUNSTATS utilities. This is not atypical, and in many cases may actually remediate similar issues, but not in this case.



The two-step solution was a) the implementation of an SQL Warehouse (WLX), allowing DBA's to very quickly assess offending SQL's in order to remediate and improve on them. b) To implement a methodology to quickly determine access path changes (BIX) and either approve them or disallow them because of access path degradation.

Now, not only does the performance team benefit from these solutions, but the DBAs are able to quickly respond to any future problems.

Traditional Tools

ISPF
Classical Drill Down
High CPU Consumption
Allows restricted Monitoring
High license costs

SQL WorkloadExpert

GUI-Interface
> 20 Uses Cases out-of-the-box
Low CPU consumption
Allows 24/7 IFCID Catching
Moderate license costs

WLX

- Catch all SQL running on a typical system, without the overhead of a trace or utilizing monitors.
- Analyze SQL Workload by aggregate, or in detail.
- Comparatively analyze access paths and SQL.
- Report and chart metrics via a standard Eclipse interface—an inexpensive and commonly-accepted desktop tool.

BIX

- Compares before and after access paths side-by-side.
- Provides an overview of historical data.

DB2 access path after the RUNSTATS on Thursday evening (As of 2013-07-13):
 Dynamic Statement Cache DB2/Group member: DBAP
 RunID : 1307200911 Created : 2013-07-13-09.11.12.969221 StmtID : 7322372
 Milliseconds: 38 Service Units: 1861 Cost Category: B


QBNO	QBTYPE	CREATOR	TABLE NAME		MTCH	IX	METH	PRNT	TABL	PRE	MXO
PLNO	TABNO	XCREATOR	INDEX NAME	ACTYP	COLS	ON	OD	QBLK	TYPE	FTCH	PSQ
1	SELECT	SBXXA00P	RES	R	0	N	0	0	W	S	0
1	2										
1	SELECT				0	N	3	0	-		0
2	0										
2	UNION				0		3	1	-		0
1	0										
3	NCOSUB	WCRKRK	WWS09_ORDER	I	1	N	0	2	T	S	0
1	6	WCRKRK	WWS09X01								
3	NCOSUB	WCRKRK	WWS06_SCM_MCCD	I	1	N	1	2	T		0
2	8	WCRKRK	WWS06X01								
3	NCOSUB	WCRKRK	WWS04_DELIVERER	I	1	N	1	2	T		0
3	1	WCRKRK	WWS04X0A								
3	NCOSUB	WCRKRK	WWS18_ORDER_STATUS	I	1	Y	1	2	T	S	0
4	1	WCRKRK	WWS18X01								
3	NCOSUB	WCRKRK	WWS10_WE_DEL_ORDER	I	6	N	1	2	T		0
5	2	WCRKRK	WWS10X01								
3	NCOSUB	WCRKRK	WWS11_WE_CAN_ORDER	I	6	N	1	2	T		0
6	4	WCRKRK	WWS11X01								
3	NCOSUB	WCRKRK	WWS18_ORDER_STATUS	I	7	Y	1	2	T	S	0
7	1	WCRKRK	WWS18X03								
9	NCOSUB	WCRKRK	WWS09_ORDER	I	1	N	0	2	T	S	0
1	1	WCRKRK	WWS09X01								
9	NCOSUB	WCRKRK	WWS18_ORDER_STATUS	I	1	Y	1	2	T	S	0
2	1	WCRKRK	WWS18X01								
9	NCOSUB	WCRKRK	WWS09_ORDER	I	4	N	1	2	T		0
3	1	WCRKRK	WWS09X01								
9	NCOSUB	WCRKRK	WWS10_WE_DEL_ORDER	I	6	N	1	2	T		0
4	1	WCRKRK	WWS10X01								
9	NCOSUB	WCRKRK	WWS06_SCM_MCCD	I	1	N	1	2	T		0
5	2	WCRKRK	WWS06X01								
9	NCOSUB	WCRKRK	WWS04_DELIVERER	I	1	N	1	2	T		0
6	2	WCRKRK	WWS04X0A								
9	NCOSUB	WCRKRK	WWS11_WE_DEL_ORDER	I	6	N	1	2	T		0
7	1	WCRKRK	WWS11X01								
9	NCOSUB	WCRKRK	WWS18_ORDER_STATUS	I	7	Y	1	2	T	S	0
8	2	WCRKRK	WWS18X03								

Figure 2

DB2 access path before the RUNSTATS on Thursday evening (As of 2013-07-11):
 Dynamic Statement Cache DB2/Group member: DBAP
 RunID : 1307111618 Created : 2013-07-11-16.18.12.544794 StmtID : 5837498
 Milliseconds: 25 Service Units: 1269 Cost Category: B

QBNO	QBTYPE	CREATOR	TABLE NAME	MTCH	IX	METH	PRNT	TABL	PRE	MXO	
PLNO	TABNO	XCREATOR	INDEX NAME	ACTYP	COLS	ON	OD	QBLK	TYPE	FTCH	PSQ
1	SELECT	SBXXA00P	RES	R	0	N	0	0	W	S	0
1	2										
1	SELECT				0	N	3	0	-		0
2	0										
2	UNION				0		3	1	-		0
1	0										
3	NCOSUB	WCRKRR	WWS18_ORDER_STATUS	I	1	Y	0	2	T	S	0
1	1	WCRKRR	WWS18X01								
3	NCOSUB	WCRKRR	WWS09_ORDER	I	6	N	1	2	T		0
2	6	WCRKRR	WWS09X01								
3	NCOSUB	WCRKRR	WWS10_WE_DEL_ORDER	I	6	N	1	2	T		0
3	2	WCRKRR	WWS10X01								
3	NCOSUB	WCRKRR	WWS11_WE_CAN_ORDER	I	6	N	1	2	T		0
4	4	WCRKRR	WWS11X01								
3	NCOSUB	WCRKRR	WWS06_SCM_MCCD	I	1	N	1	2	T		0
5	8	WCRKRR	WWS06X01								
3	NCOSUB	WCRKRR	WWS04_DELIVERER	I	1	N	1	2	T		0
6	1	WCRKRR	WWS04X0A								
3	NCOSUB	WCRKRR	WWS18_ORDER_STATUS	I	7	Y	1	2	T	S	0
7	1	WCRKRR	WWS18X03								
9	NCOSUB	WCRKRR	WWS18_ORDER_STATUS	I	1	Y	0	2	T	S	0
1	1	WCRKRR	WWS18X01								
9	NCOSUB	WCRKRR	WWS09_ORDER	I	6	N	1	2	T		0
2	1	WCRKRR	WWS09X01								
9	NCOSUB	WCRKRR	WWS18_ORDER_STATUS	I	7	Y	1	2	T	S	0
3	2	WCRKRR	WWS18X03								
9	NCOSUB	WCRKRR	WWS09_ORDER	I	4	N	1	2	T		0
4	1	WCRKRR	WWS09X01								
9	NCOSUB	WCRKRR	WWS10_WE_DEL_ORDER	I	6	N	1	2	T		0
5	1	WCRKRR	WWS10X01								
9	NCOSUB	WCRKRR	WWS11_WE_CAN_ORDER	I	6	N	1	2	T		0
6	1	WCRKRR	WWS11X01								
9	NCOSUB	WCRKRR	WWS06_ERP_MCCD	I	1	N	1	2	T		0
7	2	WCRKRR	WWS06X01								
9	NCOSUB	WCRKRR	WWS04_DELIVERER	I	1	N	1	2	T		0
8	2	WCRKRR	WWS04X0A								

Figure 3



Meet the SOFTWARE ENGINEERING GMBH “Expert” Tools:

▣ **SQL PerformanceExpert® for DB2 z/OS (SPX)**

- What-If Analysis
- SQL recommendations
- Statistics Check
- Standards-conformity (Extensible rules supported)

▣ **Bind ImpactExpert® for DB2 z/OS (BIX)**

- Compares access paths for static and dynamic SQL
- Allows only improved access paths for REBINDS and BINDS
- Helps to secure DB2 Version migration and PTFs/APARs application

▣ **SQL WorkloadExpert™ for DB2 z/OS (WLX)**

- Workload catcher
- Comparison, tuning and trending
- Out-of-the-box use cases

The Expert Performance Products are “stand-alone” or may be integrated as a complete solution package. They are sold in North America by SEGUS Inc.

For more information, please visit www.segus.com or www.seg.de.