


IBM Data Server Manager for z/OS

Saghi Amirsoleymani,
WW Principal Solution Advisor , IBM Champion



IBM Data Server Manager

Help / Welcome

Home

Administer >

Run SQL >

Monitor >

Jobs >

Optimize >

Settings >

Help >

Welcome to Data Server Manager

Set Up and Configure

- Manage users and privileges
- Set up a repository database
- Add database connections

Monitor Databases

- Add database connections
- Edit monitoring profiles
- View monitoring data

Administer and Configure Databases

- Add database connections
- Manage databases
- Create and manage database objects
- Track changes

Execute and Optimize SQL Statements

- Add database connections
- Edit and run SQL statements
- Tune SQL statements
- Discover object dependencies

IBM Data Server Manager

User name

Password

Log In

IBM Knowledge Center
Find answers quickly with IBM product documentation

Redbooks
Complimentary, step-by-step guides for download and mobile access

Community & Forum
Explore technical topics, find trial software, and join the community

Support portal
Get help today for the IBM services and software you own


Video
Start your learning journey with best practices and new product features

Twitter
Contact us by Twitter

What Customers Keep Telling Us



Too many tools, too many installs, too many repositories



Make it easier to use; provide workflow, diagnostic smarts, actionable recommendations



Provide better performance and scalability



Manage databases across the enterprise



Be competitive with other database vendors that provide one stop shop tooling

IBM Data Server Manager




Deliver a Simplified User Experience

- Single installer and integrated repository




Common integrated web console

- Provides enterprise view of your environment
- Guided workflow and analysis



Deliver familiar capabilities from the Database Tools

- Database Administration, Query Tuning and Configuration Management as extensible services



Data Server Manger – Where DBAs Spend Time



Administration

- helps you manage, and maintain complex database environments for increased productivity and optimized use of system resources



Performance Tuning

- helps you develop and implement a performance strategy including providing expert recommendations to improve query workload performance



Identifying Environment changes

- offers centralized management of database and client configuration



Troubleshooting

- capture production application workloads then compare capture and enforce configuration settings

Welcome to Data Server Manager

Set Up and Configure



- Manage DSM users and privileges
- Set up a repository database
- Add database connections

Monitor Databases



- Add database connections
- Edit monitoring profiles
- View monitoring data

Administer and Configure Databases



- Add database connections
- Manage databases
- Create and manage database objects
- Track changes

Execute and Optimize SQL Statements



- Add database connections
- Edit and run SQL statements
- Tune SQL statements
- Discover object dependencies

Features of the Data Server Manager z/OS Based Tools - At A Glance

NO CHARGE

Data Server Manager Base

- Connect to DB2 for z/OS V10/ V11/V12
- Database object navigation, viewing object detail, and linking to related objects.
- Database object dependency display.
- Data browsing and editing.
- Basic database object operations, such as creation of tables, indexes, constraints, and tablespaces; dropping of tables, indexes and constraints; altering tables.
- Show system privilege from the perspective of Group/User, Role, or SQL object. Choose:
 - "Group/User" to see the role and the relative object privilege for a user account;
 - "Role" to see the role a user account belongs to and its relative object privilege;
 - "SQL object" to see a specific object and users or roles that have the relative authority.
- Single query tuning
 - Statistics Advisor
 - Query Environment Capture
 - Access Path Graph

DB2 Performance Solution Pack

IBM Query Workload Tuner

- Launch of visual explain and tune query on the SQL editor
- Tuning wizard to capture SQL statements from multiple sources
- Tuning advisors provide recommendations for:
 - Statistics Advisor
 - Index Advisor
 - IDAA Advisor
- Problem analysis of query or workload
 - Access plan graph
 - Query formatting and annotation
 - Tuning Report
 - Test Candidate Index
 - Access Plan Comparison
 - Index Impact Analysis
 - Query and Workload Environment Capture
 - Selectivity Override

IBM DB2 Query Monitor

- Launching of DSM from Query Monitor Web UI for end to end performance analysis
- Host variable collection

OMEGAMON XE for DB2 PE

- Key Performance Indicators (KPIs) displayed in Data Server Manager on the Subsystems

DB2 Admin Solution Pack

Configuration Manager for z/OS V5.1

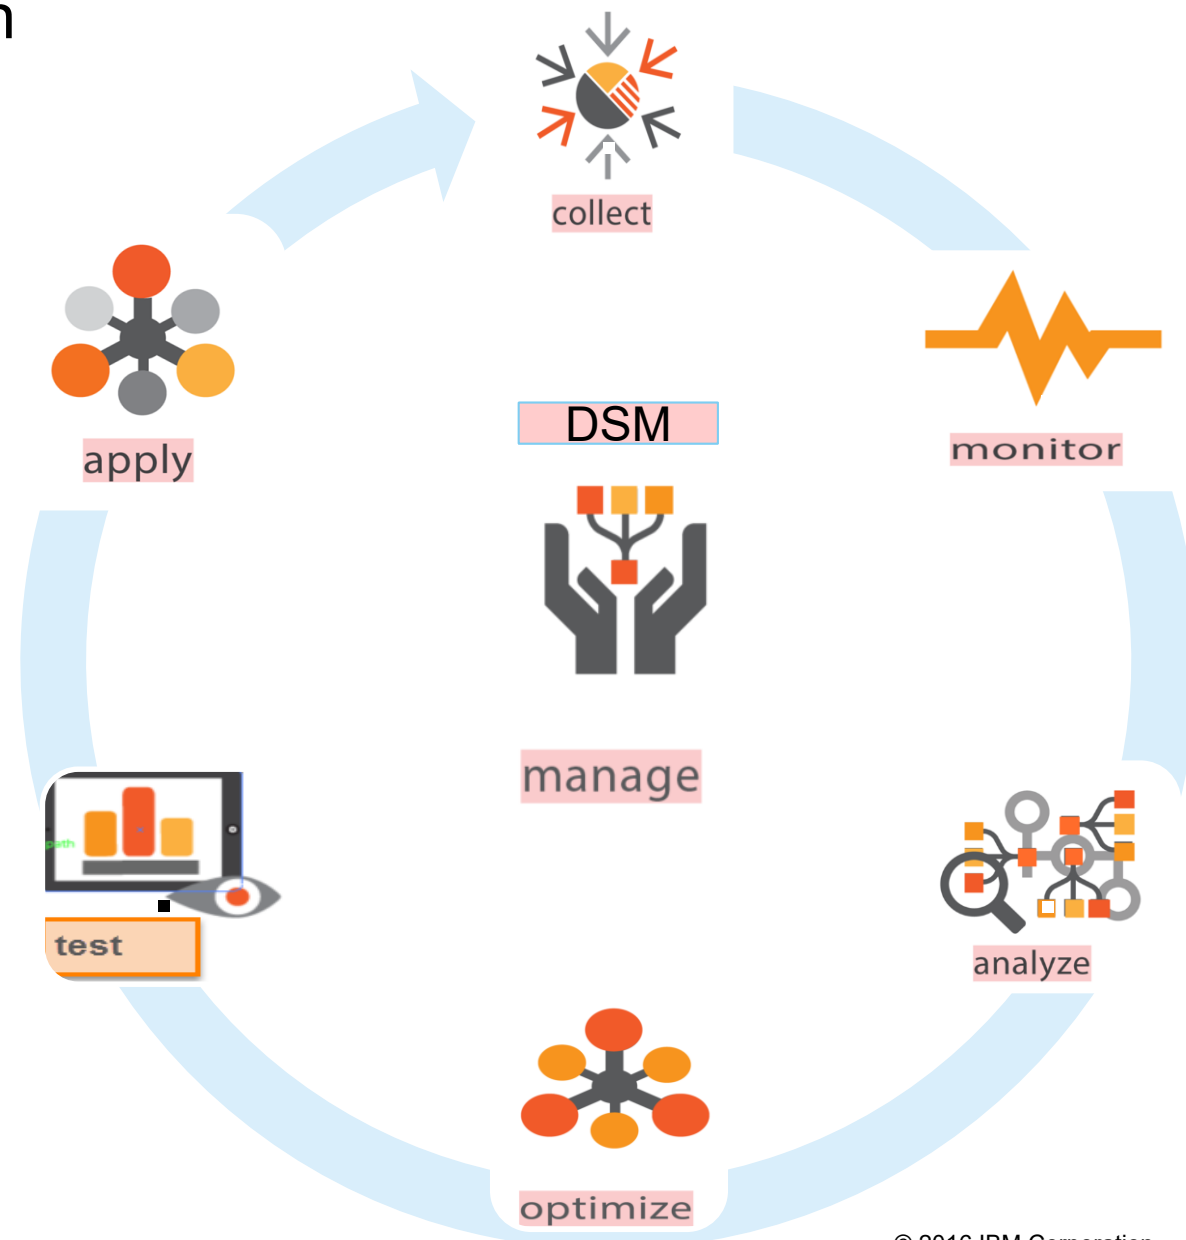
- Track configuration changes
- Configure zParm
- Compare and clone configurations
- Manage application profile
- Manage alias
- Manage and control clients

DB2 Utility Solution Pack V4.1

- Customizable profiles for performing conditional object evaluations and generating actions mapped to resolving utilities (reorg, copy, runstats, etc)
- Ability to control prioritization of objects, evaluation conditions and generated resolving actions.
- Ability to define maintenance windows for enabling autonomics, allowing DB2 to self manage utility runs
- Graphical trend analysis of historical RTS
- Capture of utility history, recording utility output, time, duration, etc.

IBM DB2 Performance Management Solution provides:

- ✓ Fast identification with automated alerts, proactive notification and 24x7 monitoring
- ✓ Tuning of queries and workloads proactively
- ✓ Expert advice with built-in advisors
- ✓ Diverse set of capabilities managed via Data Server Manager (DSM)
 - ✓ Easy-to-use integrated view of overall DB2 performance management
 - ✓ Seamless navigation and movement via functional capabilities versus individual products



Where do I start ?..... Data Server Manager



<http://ibm.biz/IWANTDSM>

Simple 3-step setup

Step 1 of 3

Accept the terms of the license agreement to continue

International License Agreement for Easy
Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, OR OTHERWISE USING THE SOFTWARE, YOU AGREE TO THE TERMS OF THIS AGREEMENT. YOU HAVE FULL AUTHORITY TO BIND LICENSING.

* DO NOT DOWNLOAD, INSTALL, COPY, OR OTHERWISE USE THE SOFTWARE WITHOUT THE EXPRESS WRITTEN PERMISSION OF IBM CORPORATION.

* PROMPTLY RETURN THE UNUSED SOFTWARE TO IBM CORPORATION.

1. Definitions

"Authorized Use" - the specified level of service units ("MSUs"), Processor Value Units ("PVUs"), or other units of measurement.

"Early Release" - a release of a Program potentially unreliable.

"IBM" - International Business Machine Corporation.

☒ I accept the terms of the license agreement.

Step 2 of 3

Specify credentials and port numbers

Specify a user name

User name:

Password:

Re-enter password

Change port

HTTP port for access

HTTPS port for access

Port used internally

Step 3 of 3

Setup is complete

The server is registered as a Windows service

You can open the IBM Data Server Manager web console by using the following URLs:

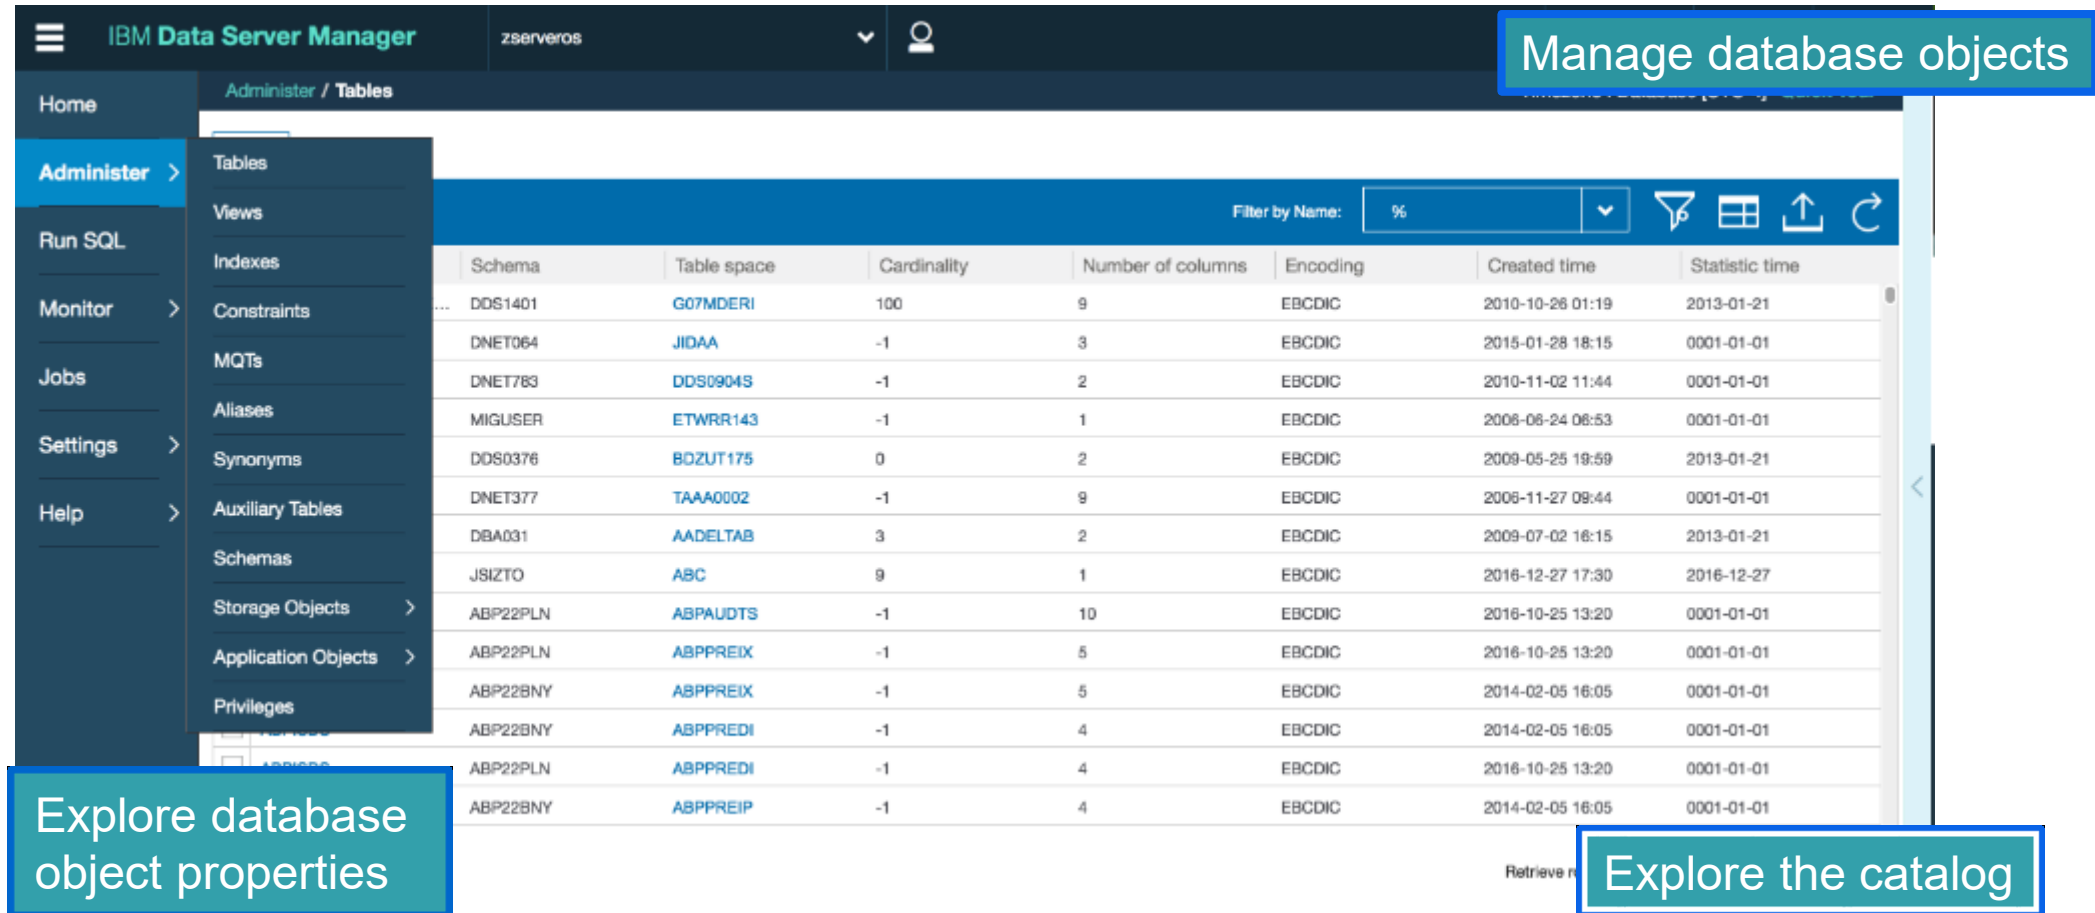
<http://localhost:11080>

<https://localhost:11081> (secure)

Record the URLs of the web console so that you can use them to log in later.

<http://ibm.biz/IWANTDSM>

Manage Databases Using the Database Explorer



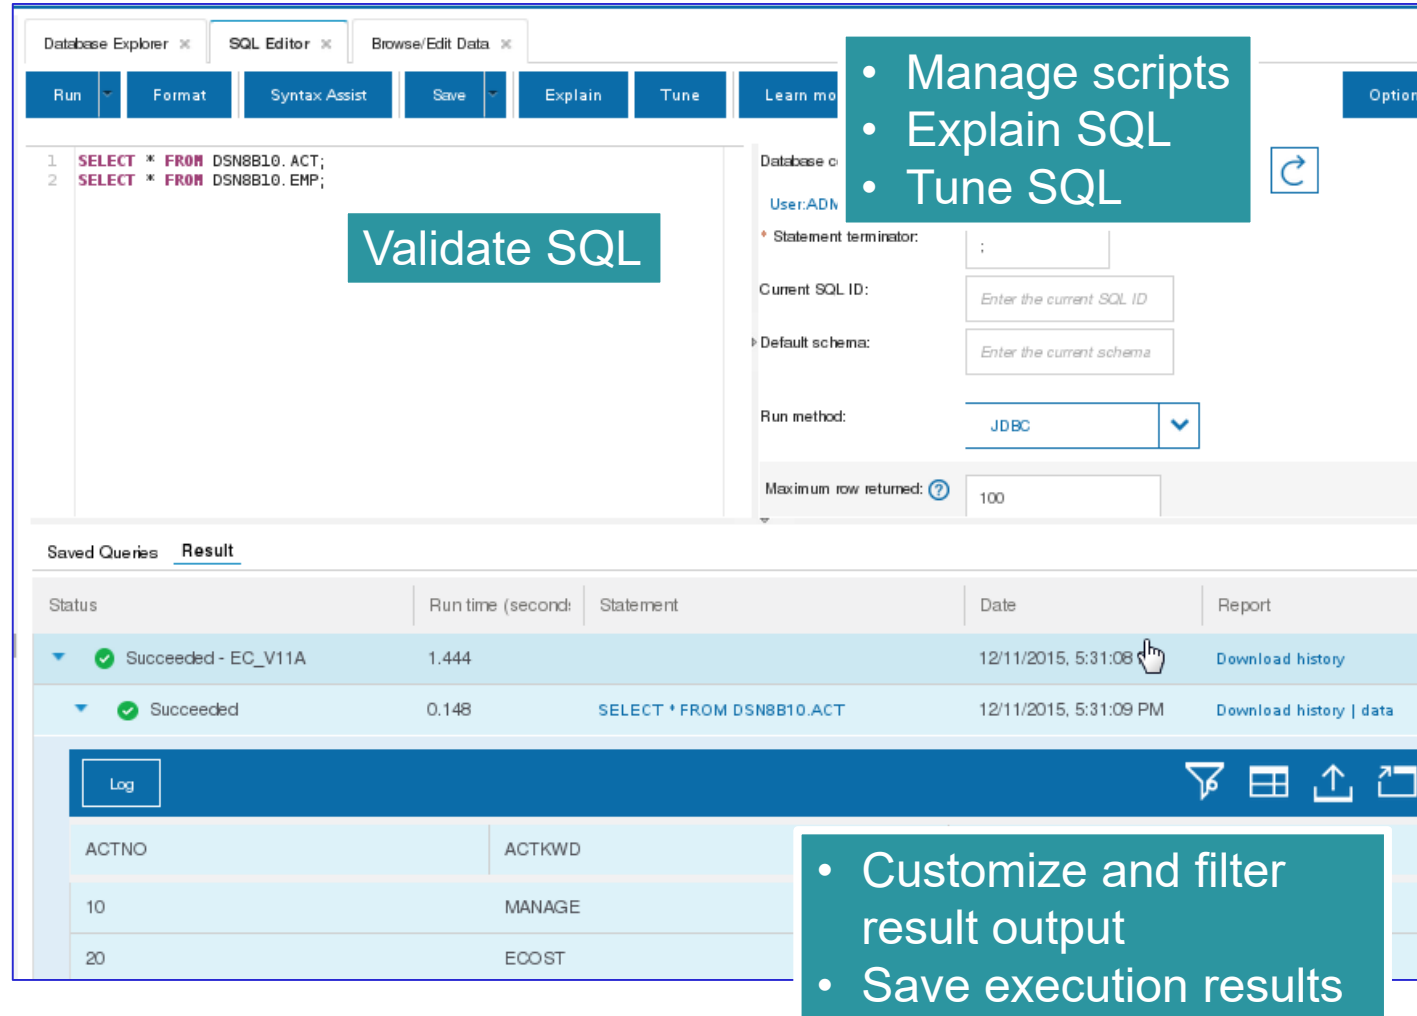
Manage database objects

Explore database object properties

Explore the catalog

Schema	Table space	Cardinality	Number of columns	Encoding	Created time	Statistic time
DDS1401	G07MDERI	100	9	EBCDIC	2010-10-26 01:19	2013-01-21
DNET064	JIDAA	-1	3	EBCDIC	2015-01-28 18:15	0001-01-01
DNET783	DDS0904S	-1	2	EBCDIC	2010-11-02 11:44	0001-01-01
MIGUSER	ETWRR143	-1	1	EBCDIC	2006-08-24 08:53	0001-01-01
DDS0376	BDZUT175	0	2	EBCDIC	2009-05-25 19:59	2013-01-21
DNET377	TAAA0002	-1	9	EBCDIC	2006-11-27 09:44	0001-01-01
DBA031	AADELTAB	3	2	EBCDIC	2009-07-02 16:15	2013-01-21
JSIZTO	ABC	9	1	EBCDIC	2016-12-27 17:30	2016-12-27
ABP22PLN	ABPAUDTS	-1	10	EBCDIC	2016-10-25 13:20	0001-01-01
ABP22PLN	ABPPREIX	-1	5	EBCDIC	2016-10-25 13:20	0001-01-01
ABP22BNY	ABPPREIX	-1	5	EBCDIC	2014-02-05 16:05	0001-01-01
ABP22BNY	ABPPREDI	-1	4	EBCDIC	2014-02-05 16:05	0001-01-01
ABP22PLN	ABPPREDI	-1	4	EBCDIC	2016-10-25 13:20	0001-01-01
ABP22BNY	ABPPREIP	-1	4	EBCDIC	2014-02-05 16:05	0001-01-01

Develop and Run SQL Scripts



The screenshot displays the IBM Analytics SQL Editor interface. At the top, there are tabs for 'Database Explorer', 'SQL Editor', and 'Browse/Edit Data'. Below these are buttons for 'Run', 'Format', 'Syntax Assist', 'Save', 'Explain', 'Tune', and 'Learn more'. The main area shows a SQL script with two lines: 'SELECT * FROM DSN8B10. ACT;' and 'SELECT * FROM DSN8B10. EMP;'. A teal box labeled 'Validate SQL' is overlaid on the script. To the right, there are input fields for 'Statement terminator', 'Current SQL ID', 'Default schema', 'Run method' (set to 'JDBC'), and 'Maximum row returned' (set to '100'). A teal box with a list of actions is overlaid on the right side: 'Manage scripts', 'Explain SQL', and 'Tune SQL'. Below the script, there is a 'Saved Queries' section and a 'Result' table. The 'Result' table has columns for 'Status', 'Run time (second)', 'Statement', 'Date', and 'Report'. It shows two rows: 'Succeeded - EC_V11A' with a run time of 1.444 and 'Succeeded' with a run time of 0.148. A teal box with a list of actions is overlaid on the bottom right: 'Customize and filter result output' and 'Save execution results'.

- Manage scripts
- Explain SQL
- Tune SQL

Validate SQL

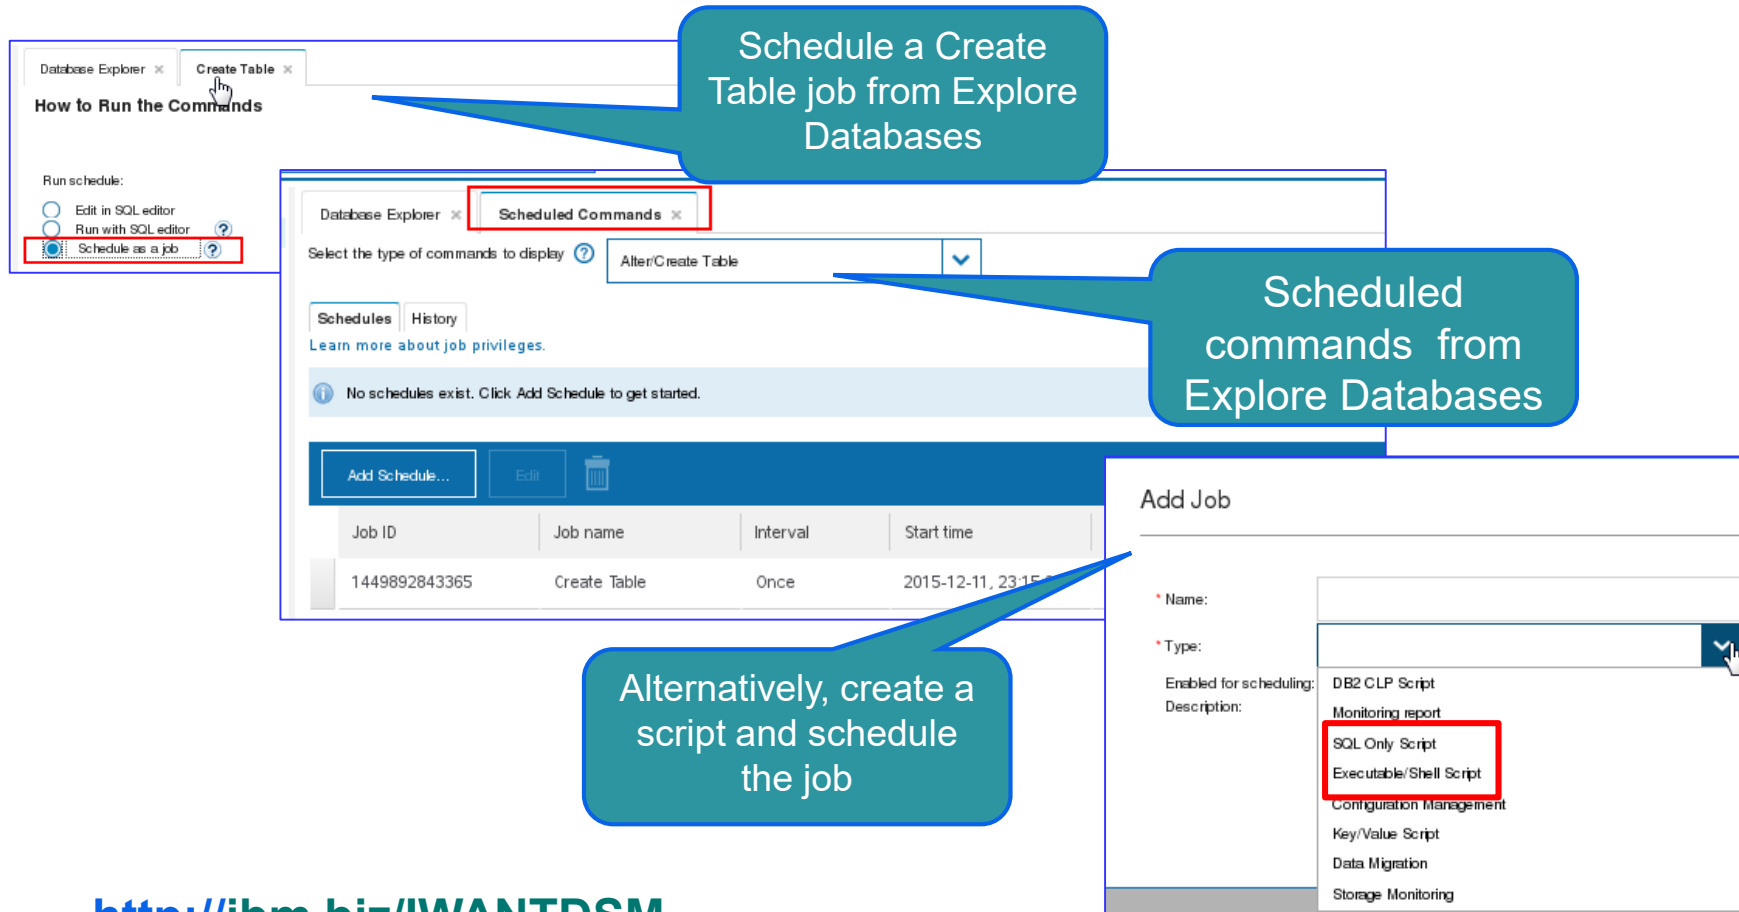
Log

Status	Run time (second)	Statement	Date	Report
✓ Succeeded - EC_V11A	1.444		12/11/2015, 5:31:08	Download history
✓ Succeeded	0.148	SELECT * FROM DSN8B10.ACT	12/11/2015, 5:31:09 PM	Download history data

ACTNO	ACTKWD
10	MANAGE
20	ECOST

- Customize and filter result output
- Save execution results

Create and schedule jobs

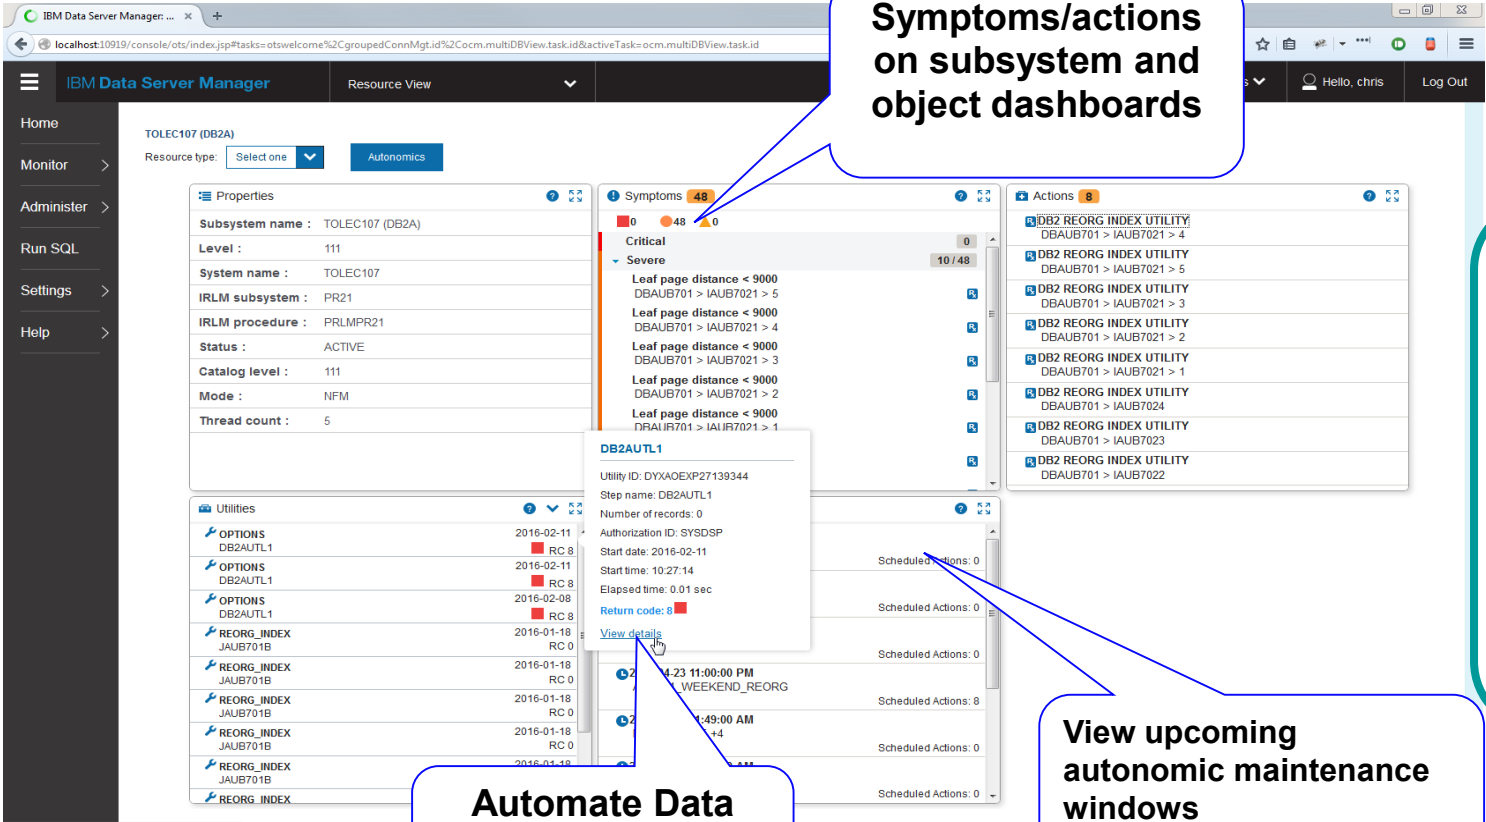


The screenshot illustrates the process of creating and scheduling jobs in IBM Analytics. It features three main panels with callout boxes explaining the steps:

- How to Run the Commands:** A panel on the left showing three options: "Edit in SQL editor", "Run with SQL editor", and "Schedule as a job". The "Schedule as a job" option is highlighted with a red box and a callout box stating: "Schedule a Create Table job from Explore Databases".
- Scheduled Commands from Explore Databases:** A central panel showing the "Scheduled Commands" tab. It displays a table with one entry: "Create Table". A callout box points to the "Add Schedule..." button, stating: "Alternatively, create a script and schedule the job".
- Add Job:** A panel on the right showing the "Add Job" form. The "Type" dropdown menu is open, showing options: "DB2 CLP Script", "Monitoring report", "SQL Only Script", "Executable/Shell Script", "Configuration Management", "Key/Value Script", "Data Migration", and "Storage Monitoring". The "Executable/Shell Script" option is highlighted with a red box.

<http://ibm.biz/IWANTDSM>

DB2 Utilities Solution Pack 2.2



Symptoms/actions on subsystem and object dashboards

“More integration, greater value”

Optimize, control manage & automate

Components:

- DB2 Automation Tool*
- DB2 High Performance Unload for z/OS*
- DB2 Sort for z/OS*
- DB2 Utilities Enhancement Tool*
- Autonomics support*
- Data Server Manager*

DB2 Utilities Solution

Automate Data Collection Utility History

View upcoming autonomic maintenance windows with scheduled actions

The screenshot displays the IBM Data Server Manager interface for a DB2 subsystem named TOLEC107 (DB2A). The interface is divided into several sections: Properties, Symptoms, Actions, and Utilities. The Properties section shows details about the subsystem, including its level, system name, IRLM subsystem, IRLM procedure, status, catalog level, mode, and thread count. The Symptoms section displays a list of critical and severe symptoms, such as 'Leaf page distance < 9000'. The Actions section shows a list of scheduled actions, including 'DB2 REORG INDEX UTILITY'. The Utilities section displays a table of utility execution history, including options, reorg index, and reorg index utilities. A callout box provides details for a specific utility, DB2AUTL1, showing its utility ID, step name, number of records, authorization ID, start date, start time, elapsed time, and return code. Another callout box shows the 'View details' link for a specific utility execution.

Reactive vs. Proactive performance management

- Problems addressed after performance impact
- Measuring flashing light indicators
- Noticing either the very good or the very bad
- Takes longer to react to bad performance because of measuring lagging indicators



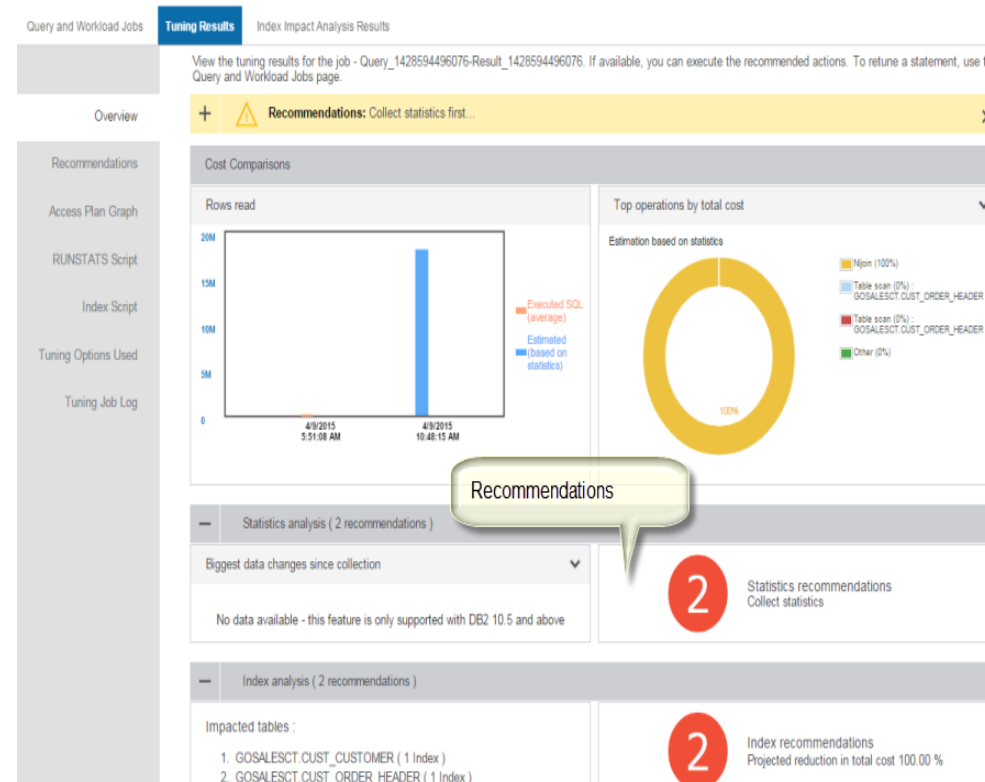
- Understanding what behavior is desired
- Measuring leading indicators
- Capture best practices and procedures
- Team responsible for creating the measurements understands the what and why

Traditional Reactive Tuning



Expert-Enabled Query Tuning Out of the Box

- **Identify query candidates from numerous sources**
 - DB2 catalog
 - Dynamic statement cache
 - User defined repository
- **Facilitate analysis**
 - Query formatting
 - Access path visualization and annotation
- **Get expert tuning advice**
 - Improve statistics quality



Capture from Statement Cache & customize collection

IBM Data Server Manager

LABEC504

Optimize / Start Tuning

Home

Administer >

Run SQL

Monitor >

Jobs

Optimize >

Settings >

Help >

Source

Scope

Input SQL

Files

Statement cache

Packages and plans

User defined repository

Select an existing filter or template to customize:

Manage Filters

More options

Default SQLID: SYSADM

Using the SYSPROC.OPT_RUNSQL: ☐

Exclude system statements: ☒

Capture from data sharing: ☒

Select data sharing members from the following available connections. [Manage Database Connections](#)

<input type="checkbox"/>	NAME	LOCATION	HOST NAME
<input type="checkbox"/>	LABEC504Admf002	STLEC1	9.30.112.62
<input type="checkbox"/>	LOCD821	LOCD821	9.125.72.130

Capture & Select from Data Sharing members

- Accelerate analysis, reduce downtime

Manage filters

Add

Edit

Copy

Delete

Apply

<input type="checkbox"/>	NAME	SHARED	OWNER	DATE MODIFIED
<input type="checkbox"/>	STABILIZED=Y	<input type="checkbox"/>	admin	2017-06-11 19:20:39
<input type="checkbox"/>	EP_RATIO_DESC	<input checked="" type="checkbox"/>	SYSTEM	2017-05-23 19:43:26
<input type="checkbox"/>	PCT_CPU_DESC	<input checked="" type="checkbox"/>	SYSTEM	2017-05-23 19:43:15
<input type="checkbox"/>	PCT_ELAP_DESC	<input checked="" type="checkbox"/>	SYSTEM	2017-05-23 19:43:08
<input type="checkbox"/>	JASON_GETPAGES	<input checked="" type="checkbox"/>	admin	2017-05-16 11:00:02
<input type="checkbox"/>	aa-copied2	<input type="checkbox"/>	admin	2017-04-27 16:04:19
<input type="checkbox"/>	aa-copied1	<input type="checkbox"/>	admin	
<input type="checkbox"/>	cc	<input type="checkbox"/>	admin	
<input type="checkbox"/>	aa	<input type="checkbox"/>	admin	
<input type="checkbox"/>	samename	<input type="checkbox"/>	hjq	
<input type="checkbox"/>	samename-copied2	<input type="checkbox"/>	zhw	
<input type="checkbox"/>	samename-copied1	<input checked="" type="checkbox"/>	zhw	
<input type="checkbox"/>	defaultFilter468-copied1	<input checked="" type="checkbox"/>	zhw	
<input type="checkbox"/>	defaultFilter468	<input type="checkbox"/>	admin	
<input type="checkbox"/>	ACQUM_CPU_DESC-copied2	<input checked="" type="checkbox"/>	zhw	
<input type="checkbox"/>	samename	<input checked="" type="checkbox"/>	admin	
<input type="checkbox"/>	ACQUM_CPU_DESC-copied1	<input checked="" type="checkbox"/>	zhw	
<input type="checkbox"/>	defaultFilter70	<input checked="" type="checkbox"/>	zhw	
<input type="checkbox"/>	ACQUM_CPU_DESC	<input checked="" type="checkbox"/>	SYSTEM	
<input type="checkbox"/>	ACQUM_ELAP_DESC	<input checked="" type="checkbox"/>	SYSTEM	

Share & Save Filters to speed up analysis

Manage filters

Filter name:

defaultFilter897

Shared

☐

Select filters to capture statements from the dynamic statement cache on the subsystem.

PRIMAUTH

=

Select filters to capture statements if statement cache trace is enabled. [Learn More](#)

STAT_EXECB

=

Select column(s) to sort captured result.

ASC

Maximum statements: 100

Create new filter either shared or private

Home

Administer >

Run SQL

Monitor >

Jobs

Optimize >

Settings >

Help >

IBM Data Server Manager

REGECS80V12NFM

Select SQL statement source

Choose a source:

Statement Cache

Maximum number of statements to capture:

100

SQID

SYSADM

☐ Use the SYSPROC.OPT_RUNSQL stored procedure to run the EXPLAIN statement

Select filter for capturing SQL statements from the dynamic statement cache on the subsystem.

PRIMAUTH

PRIMAUTH

STMT_TOKEN

STMT_TEXT

STABILIZED

=

filter value

+

-

PRIMAUTH

PRIMAUTH

STMT_TOKEN

STMT_TEXT

STABILIZED

=

filter value

+

-

Additional filters if statement cache trace is enabled.

AVG_STAT_GPAG

AVG_STAT_GPAG

STAT_ELAP

STAT_CPU

STAT_EXECB

STAT_RSCNB

=

filter value

+

-

This stored procedure will allow the user to Explain the statement. Click the checkbox to activate this capability.

Add or remove additional filter criteria

19

19

© 2016 IBM Corporation

Home
Administer >
Run SQL
Monitor >
Jobs
Optimize >
Settings >
Help >

Review captured statements

View Full Statement View Runtime Access Plan Graph

Sort columns Export

STMT_ID	STMT_TEXT	SCHEMA	PRMAUTH	STMT_TOKEN	STAT_ELAP	S
80	SELECT COUNT(*) FROM SYSIBM.SYSTABLES WHERE TYPE = 'V'	SYSADM	SYSADM		0	
21	SELECT CURRENT SCHEMA AS SCHEMA FROM SYSIBM.SYSDUMMY1	SYSADM	SYSADM		0.00445	
28	SELECT CURRENT SCHEMA AS SCHEMA FROM SYSIBM.SYSDUMMY1	TPCDS	SYSADM		0.000416	
102	SELECT CURRENT SCHEMA AS SCHEMA FROM SYSIBM.SYSDUMMY1	SYSADM	SYSADM		0.000252	
66	/* IBM_DSM */ DELETE FROM SYSADM.DSN_VIRTUAL_INDEXES	TPCDS	SYSADM		0	
132	/* IBM_DSM */ DELETE FROM SYSADM.DSN_VIRTUAL_INDEXES	SYSADM	SYSADM		0.000427	

Total: 100 Selected: 1

< 1 2 3 ... 10 >

10 | 25 | 50 +

Previous Tune Highlighted Statement Tune All Statements

Create a workload tuning job to tune all the captured statements

Create a single-query tuning job to tune the highlighted statement

Analyze Access Plans

New Recommendation Integration

- **Visualize access path**
 - See flow of query processing
 - See indexes and operations
 - See optimizer rationale
- **Assess access path stability to reduce risk of performance regression**
 - Is the optimizer able to apply the filtering early?
 - Are there indexes that support an efficient path?
 - Do statistics allow distinction between the choices?

SQL Statement Warnings Environment & Explain Options Open in New Window Graph view

Description **Recommendations**

Recommended RUNSTATS control statement available below.

Repair

This version of the recommendations repairs problems only. Run this version if you want to save time and CPU resources by not collecting statistics.

Show RUNSTATS Control Statements

Complete

Index recommendations found.

Recommended indexes are listed below. You can view the DDL for the recommendations.

Show DDL

QWT0001V.CUST_ORDER_HEADER

Recommended RUNSTATS commands

```

RUNSTATS TABLESPACE "QWT0001DB"."TSQWT202"
TABLE("QWT0001V"."CUST_ORDER_HEADER") TABLESAMPLE SYSTEM 40
COLUMN("CUST_TOTAL")
COLGROUP("CUST_TOTAL") FREQUENCY COUNT 10
SORTDEVT SYSDA SORTNUM 4
INDEX("QWT0001V"."CUST_ORDER_HEADER_FK" HISTOGRAM NUMCOLS 1
NUMQUANTILES 20)
SERLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
  
```

QUERY

QBI

NLJOIN

F0

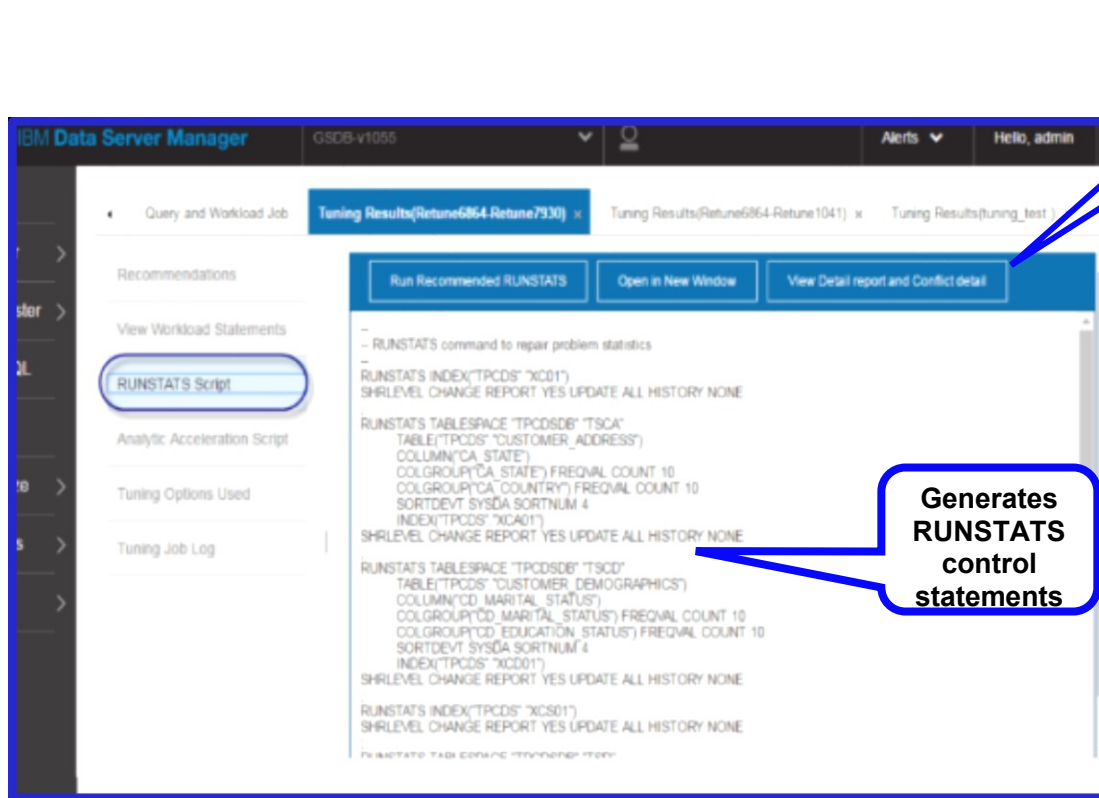
CUST_ORDER_HEADER 39,389,000

IXSCAN 1

IDX_CUST_CUSTOMER 5

CUST_CUSTOMER 30

<http://ibm.biz/IWANTDSM>



Conflicting statistics explanation

Generates RUNSTATS control statements

- Provides advice on
 - Missing statistics
 - Conflicting statistics
 - Out-of-date statistics

- Results

- Accurate estimated costs
- Better query performance
- Less CPU consumption
- Improved maintenance window throughput

“80 % of access path PMRs could be resolved by statistics advisor before calling IBM support.” – IBM Support

What to do next Performance Tuning Using Data Server Manager : Query Workload Tuner 5.1




<http://ibm.biz/IWANTDSM>

Why is workload tuning important?


Workload: Multiple SQL statements defined by user

- The effort for tuning the whole application with good performance by evaluating every statement is overwhelming. Optimization decisions are based on trade offs:
 - Statistics – CPU costs vs. query savings
 - Indexing – query speed vs resource and transaction
- Sometimes performance improvement for one statement in an application may regress other statements in the application.
- When your application data grows, allows you to do proactive application health check periodically to find potential problems earlier before costly application outages
- Workload tuning speeds up analysis
 - Analyzes multiple queries at once
- Workload tuning consolidates and optimizes recommendation for overall workload
 - Statistics recommendations
 - Index recommendation

Where is the most time spent?



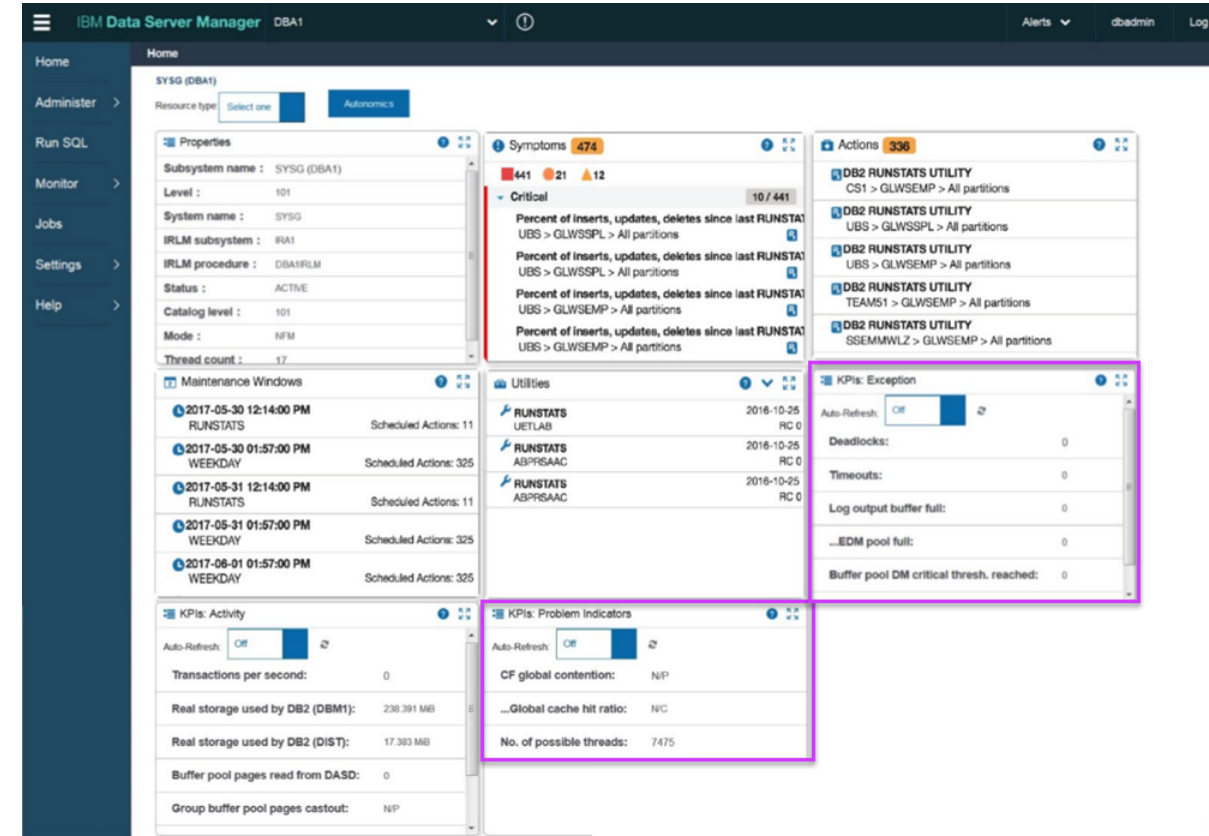
Need to analyze time distribution:
where is the time really spent?



- Application logic inefficiency
 - often combined with
Class 2 CPU << Class 1 CPU
- Network problems
- Class 2 not active all the time

Improve Performance and Reduce Costs


- Improve end-user experience of performance
 - Monitor KPIs that better reflect end-user experience
 - i.e., transaction response time
 - Get early warning of degrading performance before users are affected
 - Isolate problems to correct area for fast response
 - Get expert advice for improving query and workload performance
- Reduce costs
 - Improve performance and govern system utilization to defer upgrades
 - Save hours of staff time and stress
 - Isolate problems to the right layer of the application stack, database component, even the line of code
 - Enable developers and novice DBAs to tune like an expert




Accelerate analysis
and reduce downtime
for urgent situations



Where is the time spent within DB2?



- Turn off expensive traces
- Inefficient access paths
 - Explain



- What is the largest contributor
 - Class 3 and 8 analysis

What is your query tuning objective
reduced CPU usage or reduced
elapsed time ?



Common Scenarios & Collection

- Most expensive SQL statement in your DB2 subsystem
- Most expensive SQL within a PLAN
- All of the PLANS where a specific package is used
- All of the “exceptional SQL for a given plan”
- All of the objects accessed by a specific package
- All of the SQL which access a specific object
- Unnecessary negative SQLCODES





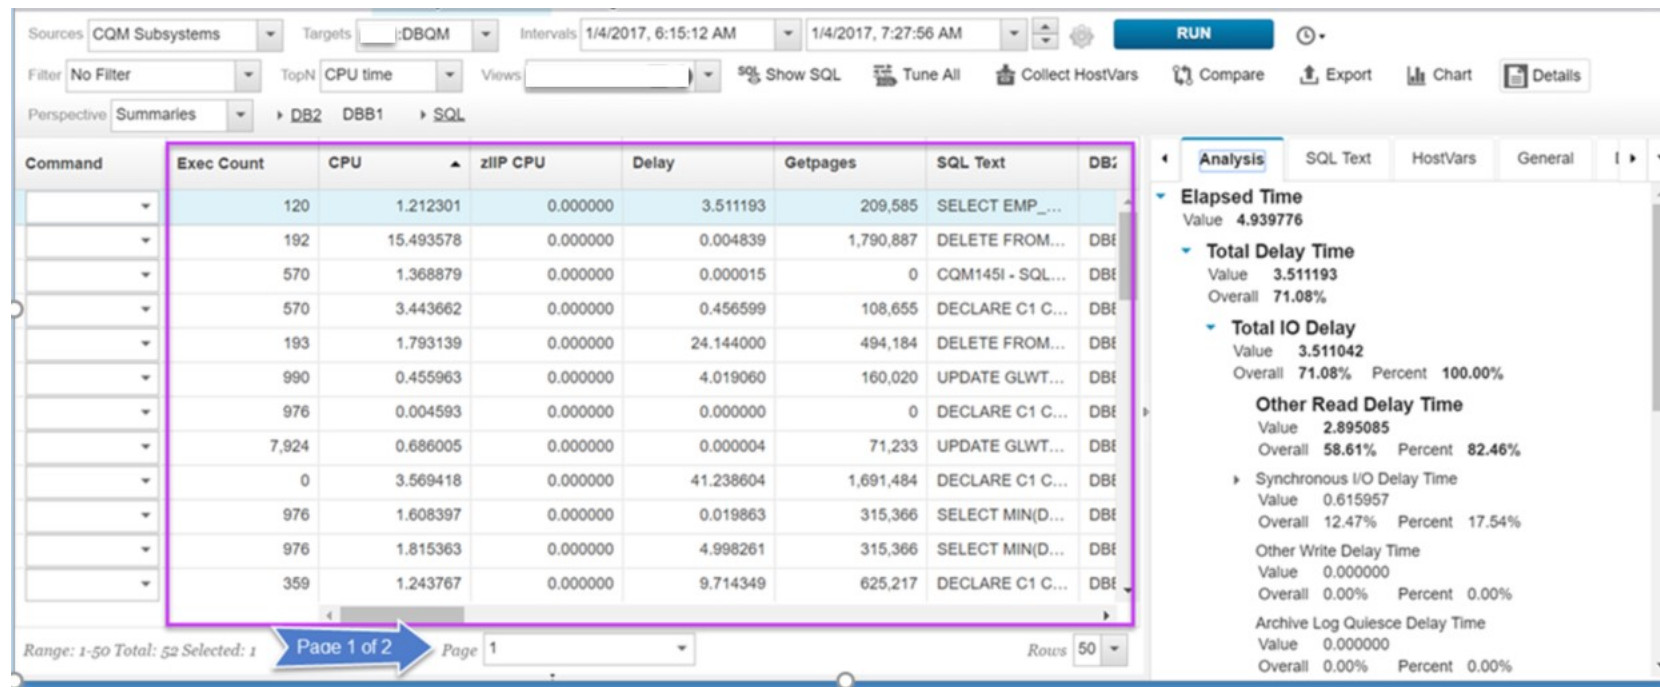
Determine what data needs to be collected

- Data available to be collected or viewed
 - SQL metrics
 - DB2 object access
 - SQL text and host variables
 - DB2 commands
 - Negative SQLCODES
 - Expanded and grouped Information about exceptions
 - Buffer Pool Statistics
 - Delays
- Three types of data
 - **Summary** – data summarized for each unique SQL statement executed in a particular interval of time
 - Plan + Program + Section + Statement # + Statement type
 - SQL Codes are not collected by default
 - **Exceptions** – individual SQL calls that have exceeded user defined thresholds
 - **Alerts** – events that require immediate attention; can be classified as exceptions



2 Identify workload for proactive tuning

- Identify the topN (50 , 200 , 250) most expensive queries
- Drill down into results
- Save workload and/or start tuning



The screenshot shows the IBM Analytics workload analysis interface. The top navigation bar includes 'Sources' (CQM Subsystems), 'Targets' (DBQM), and 'Intervals' (1/4/2017, 6:15:12 AM to 1/4/2017, 7:27:56 AM). A 'RUN' button is present. Below the navigation bar, there are filters for 'Filter' (No Filter), 'TopN' (CPU time), and 'Views'. The 'Perspective' is set to 'Summaries', and the 'DB2' and 'DBB1' perspectives are selected. The main table displays a list of queries with columns: Command, Exec Count, CPU, zIIP CPU, Delay, Getpages, SQL Text, and DB2. The table is sorted by CPU time, with the top row showing a query with an Exec Count of 120 and a CPU time of 1.212301. The right-hand pane shows the 'Analysis' tab, which includes a summary of the workload analysis. The 'Elapsed Time' is 4.939776, and the 'Total Delay Time' is 3.511193. The 'Total IO Delay' is 3.511042, with an overall percentage of 71.08%. The 'Other Read Delay Time' is 2.895085, with an overall percentage of 58.61%. The 'Synchronous I/O Delay Time' is 0.615957, with an overall percentage of 12.47%. The 'Other Write Delay Time' is 0.000000, with an overall percentage of 0.00%. The 'Archive Log Quiesce Delay Time' is 0.000000, with an overall percentage of 0.00%.

Command	Exec Count	CPU	zIIP CPU	Delay	Getpages	SQL Text	DB2
	120	1.212301	0.000000	3.511193	209,585	SELECT EMP_...	
	192	15.493578	0.000000	0.004839	1,790,887	DELETE FROM...	DBE
	570	1.368879	0.000000	0.000015	0	CQM145I - SQL...	DBE
	570	3.443662	0.000000	0.456599	108,655	DECLARE C1 C...	DBE
	193	1.793139	0.000000	24.144000	494,184	DELETE FROM...	DBE
	990	0.455963	0.000000	4.019060	160,020	UPDATE GLWT...	DBE
	976	0.004593	0.000000	0.000000	0	DECLARE C1 C...	DBE
	7,924	0.686005	0.000000	0.000004	71,233	UPDATE GLWT...	DBE
	0	3.569418	0.000000	41.238604	1,691,484	DECLARE C1 C...	DBE
	976	1.608397	0.000000	0.019863	315,366	SELECT MIN(D...	DBE
	976	1.815363	0.000000	4.998261	315,366	SELECT MIN(D...	DBE
	359	1.243767	0.000000	9.714349	625,217	DECLARE C1 C...	DBE

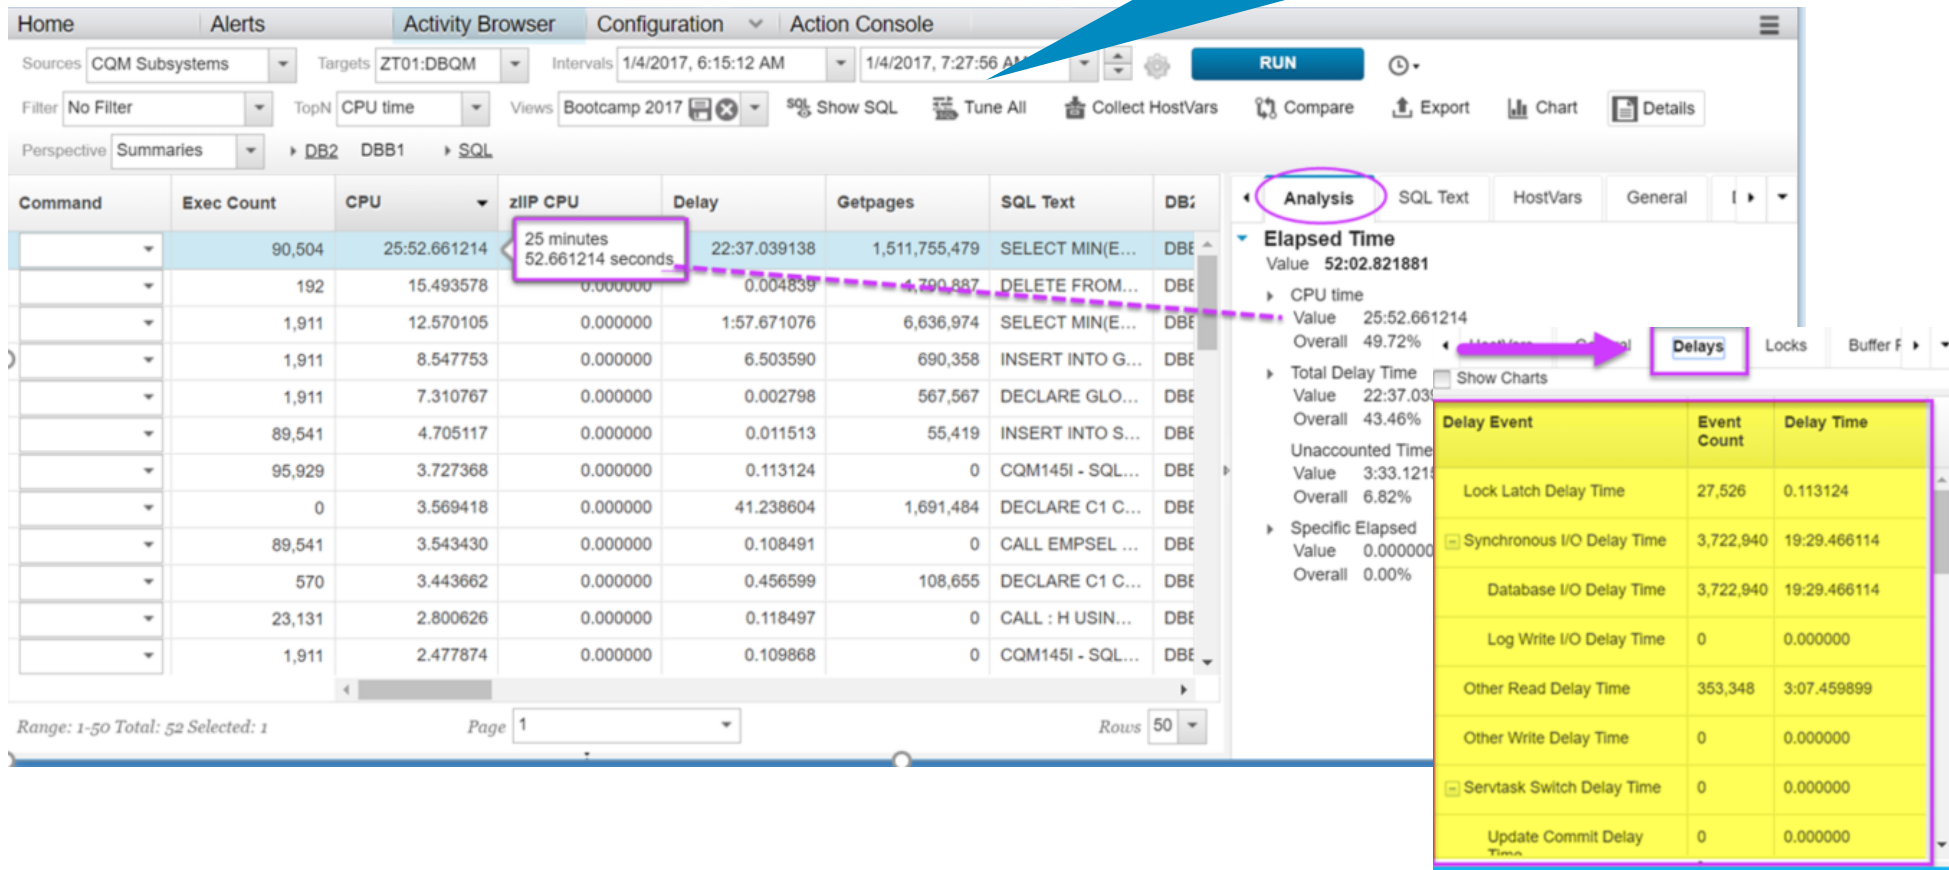
Range: 1-50 Total: 52 Selected: 1 Page 1 of 2 Page 1 Rows 50



2

Drill down into results

Tune or Tune all



The screenshot displays the IBM Analytics Activity Browser interface. The main table shows execution details for various SQL commands. A callout box highlights the CPU time for the first row, showing 25 minutes and 52.661214 seconds. A blue callout bubble points to the 'Tune or Tune all' button. The right-hand panel shows the 'Analysis' tab, which includes a table of delay events.

Command	Exec Count	CPU	zIIP CPU	Delay	Getpages	SQL Text	DB2
	90,504	25:52.661214	25 minutes 52.661214 seconds	22:37.039138	1,511,755,479	SELECT MIN(E...	DBE
	192	15.493578	0.000000	0.004839	4,790,887	DELETE FROM...	DBE
	1,911	12.570105	0.000000	1:57.671076	6,636,974	SELECT MIN(E...	DBE
	1,911	8.547753	0.000000	6.503590	690,358	INSERT INTO G...	DBE
	1,911	7.310767	0.000000	0.002798	567,567	DECLARE GLO...	DBE
	89,541	4.705117	0.000000	0.011513	55,419	INSERT INTO S...	DBE
	95,929	3.727368	0.000000	0.113124	0	CQM145I - SQL...	DBE
	0	3.569418	0.000000	41.238604	1,691,484	DECLARE C1 C...	DBE
	89,541	3.543430	0.000000	0.108491	0	CALL EMPSEL ...	DBE
	570	3.443662	0.000000	0.456599	108,655	DECLARE C1 C...	DBE
	23,131	2.800626	0.000000	0.118497	0	CALL : H USIN...	DBE
	1,911	2.477874	0.000000	0.109868	0	CQM145I - SQL...	DBE

Analysis Panel:

Elapsed Time
Value 52:02.821881

- CPU time
Value 25:52.661214
Overall 49.72%
- Total Delay Time
Value 22:37.039138
Overall 43.46%
- Unaccounted Time
Value 3:33.121881
Overall 6.82%
- Specific Elapsed
Value 0.000000
Overall 0.00%

Delays

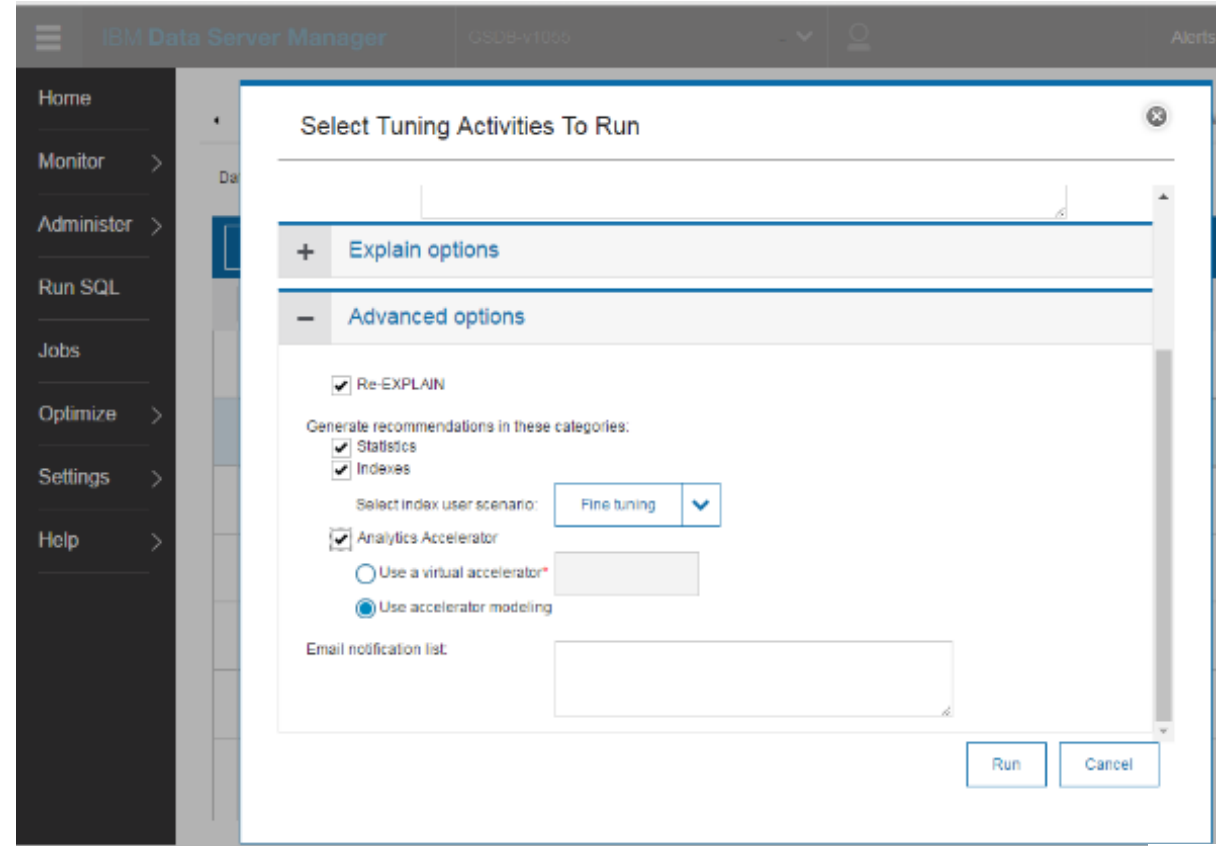
Delay Event	Event Count	Delay Time
Lock Latch Delay Time	27,526	0.113124
Synchronous I/O Delay Time	3,722,940	19:29.466114
Database I/O Delay Time	3,722,940	19:29.466114
Log Write I/O Delay Time	0	0.000000
Other Read Delay Time	353,348	3:07.459899
Other Write Delay Time	0	0.000000
Servtask Switch Delay Time	0	0.000000
Update Commit Delay	0	0.000000



analyze

3 Execute Advisors

- Statistics
 - Get recommendations on the best statistics to capture to influence access path selection
- Index
 - Get recommendations on indexes changes that can reduce database scans
- Analytics Accelerator
 - Get recommendations on optimizing and managing accelerated analytic queries and applications



Improve statistics quality and collection

Query and Workload Job

Tuning Results(tpch60GB) x

Tuning Results(Query_1449642937354-Result_1449642937354) x

Recommendations

View Workload Statements

RUNSTATS Script

Index Script

Analytic Acceleration Script

Tuning Options Used

Tuning Job Log

Run Recommended RUNSTATS

Open in New Window

View Detail report and Conflict detail

```
--
-- RUNSTATS command to repair problem statistics
--
RUNSTATS TABLESPACE "TPCH060DB"."TSCUS"
TABLE("TPCH60"."CUSTOMER")
COLUMN("C_ACCTBAL")
COLGROUP("C_ACCTBAL") HISTOGRAM NUMQUANTILES 20
COLGROUP("C_MKTSEGMENT") FREQVAL COUNT 10
SORTDEVT SYSDA SORTNUM 4
INDEX("TPCH60"."IXCUS01")
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
;
RUNSTATS TABLESPACE "TPCH060DB"."TSLI"
TABLE("TPCH60"."LINEITEM")
COLUMN("L_ORDERKEY")
COLGROUP("L_QUANTITY") FREQVAL COUNT 10 HISTOGRAM
NUMQUANTILES 20
COLGROUP("L_ORDERKEY","L_SUPPKEY")
COLGROUP("L_RETURNFLAG") FREQVAL COUNT 10
COLGROUP("L_PARTKEY","L_SUPPKEY")
COLGROUP("L_SHIPINSTRUCT") FREQVAL COUNT 10
SORTDEVT SYSDA SORTNUM 4
INDEX("TPCH60"."IXLI01",
"TPCH60"."LINEITEM_VIRT_IDX_144910527477340314")
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
;
RUNSTATS TABLESPACE "TPCH060DB"."TSNTN"
TABLE("TPCH60"."NATION")
COLUMN("N_NAME")
COLGROUP("N_NAME") FREQVAL COUNT 10
SORTDEVT SYSDA SORTNUM 4
INDEX("TPCH60"."IXN01")
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
;
```

Statistics Advisor Detail Report

Statistics Advisor Detail Report
 Analysis start time: 2015-12-02 18:31:48.504682
 Analysis end time: 2015-12-02 18:31:50.455052

TABLE TPCH60.CUSTOMER
 Table type: Table
 Cardinality: 60000.0
 Collection time: 2015-11-17 01:30:26.468977
 Statistics status: OK

INDEXES:

TPCH60.IXCUS01 (C_CUSTKEY)
 First key cardinality: 60000.0
 Full key cardinality: 60000.0
 Data repetition factor: 1170.0
 Collection time: 2015-11-17 01:30:26.468977
 Statistics status: conflicting

Interesting columns:

C_MKTSEGMENT
 Cardinality: 5.0
 Uniform statistics collection time: 2015-11-17 01:30:26.468977
 Uniform statistics status: OK
 Frequency statistics collection time: null
 Frequency statistics status: missing
 Histogram statistics collection time: null
 Histogram statistics status: missing
 Possibly point skewed: Yes
 Symptom: Columns with low COLCARD (the number of distinct values in a column)
 Possibly range skewed: No

Conflicting
statistics
explanation

Results

- Accurate estimated costs
- Better query performance
- Less CPU consumption
- Improved maintenance window throughput



optimize

“80 % of access path PMRs could be resolved by statistics advisor before calling IBM support.” – IBM Support

Indexing advice to improve database design

Workload Index Impact Analysis



analyze

- Indexes are decided at design stage
 - Lot of effort is spent making SQL to use the provided indexes
 - But what if the SQL is "right" and it's the indexes that are "wrong"
 - Cost resources to maintain
 - *How do you simply test your hypotheses without impacting production?*
- Removing obsolete indexes simplify use
 - Consolidate indexes and provide a single recommendation
 - Enables what-if analysis
 - Provides DDL to create indexes
 - Run immediately or save
- Test before deployment
 - Use virtual index capabilities built into the DB2 engine



% Performance gain on Statement/ Package level

Index	Table	Key Columns	Estimated Performance Gain(%)	Statement Impact	Action
Original Index Recommendations					
CUST_CREDIT_CARD_VIRT_IDX_145215367825162295	CUST_CREDIT_CARD	CUST_CODE(ASC), CREDIT_METHOD_CODE(ASC)	99.81	4	View Statements
CUST_ORDER_DETAIL_VIRT_IDX_145215367820745351	CUST_ORDER_DETAIL	CUST_ORDER_NUMBER(ASC)	99.85	6	View Statements
CUST_ORDER_HEADER_VIRT_IDX_145215367822725941	CUST_ORDER_HEADER	CUST_CODE(ASC), CUST_CODE(ASC)	98.63	1	View Statements
CUST_ORDER_HEADER_VIRT_IDX_145215367821540097	CUST_ORDER_HEADER	CUST_CODE(ASC), CREDIT_METHOD_CODE(ASC)	99.99	2	View Statements



test

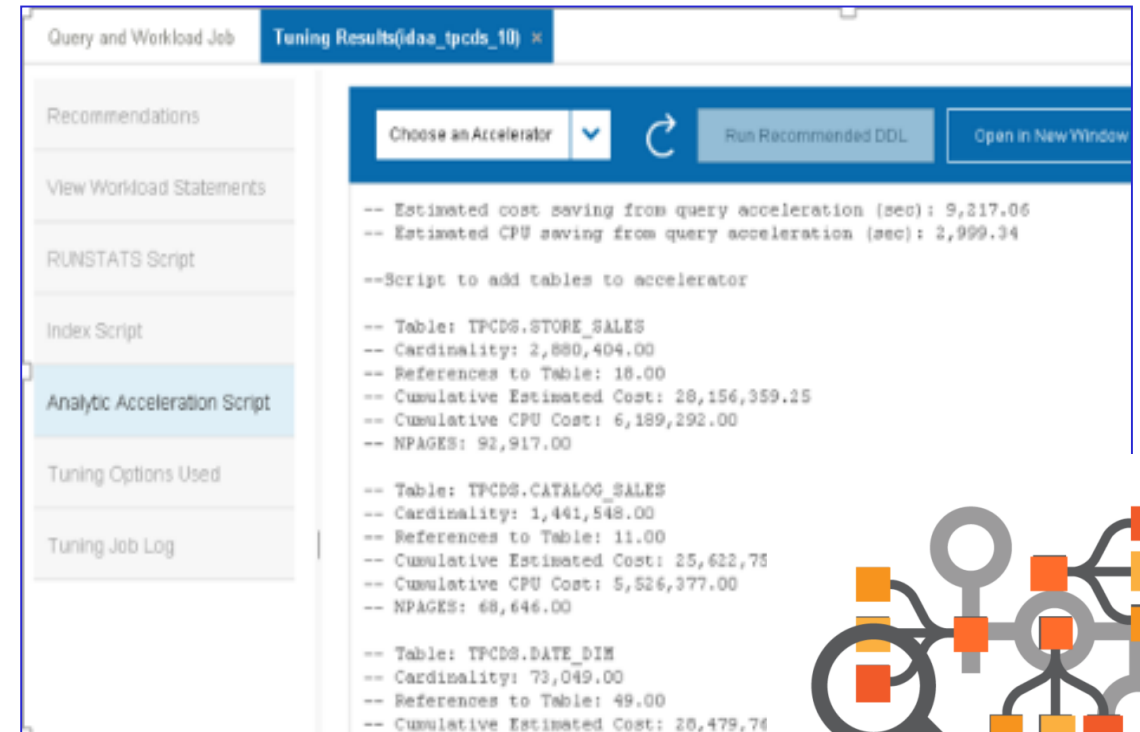
Optimizing the selection and tuning of accelerated workloads

■ Workload Analytics Accelerator Advisor

- Identify candidate queries and tables to be routed to the Accelerator
- Identify candidate tables to be routed to the accelerator
- Implement advisor-based tuning recommendations for mixed workloads of accelerated and un-accelerated queries
- Diagram accelerated queries in Access Plan Graphs
- Integrates with Query Monitor and OMPE for capturing query workloads for complete analysis
- Enable “what if” analysis

■ Benefits

- Shorten the process of selecting tables to be accelerated
- Visualize access paths of accelerated queries
- Increase productivity by working with accelerated queries through a unified interface
- Increase overall system capacity



analyze

Prevent problems before they impact the business

- ## Workload Comparison

- Enhancements
 - **Compare two different workloads**

DB2 Query Monitor for z/OS v3.3

Welcome SUDB101

About | Help | Logout | IBM Support

IBM

Home

Alerts

Activity Browser

Configuration

Action Console

Sources

CQM Subsystems

Targets

ZT01.DBQM

Intervals

1/4/2017, 6:55:31 AM

1/4/2017, 7:27:56 AM

DB2

All

RUN

Sources

CQM Subsystems

Targets

ZT01.DBQM

Intervals

1/4/2017, 6:15:12 AM

1/4/2017, 6:55:09 AM

DB2

All

Baseline

Ad Hoc

Compare

All

Filter

No Filter

TopN

No TopN

Views

Standard View

SQL

Show SQL

Time

All

Collect HostVars

Compare

Export

Chart

Details

Perspective

Summaries

Program

Command	Program	Row Status	Exec Count	Calls	Elapsed
EMPEND	EMPEND	RETAINED	-1,176	Old: 2,104 New: 928 Change: -1,176 Percent Change: -55.89%	Old: 7.426930 New: 4.535108 Change: -2.891822 Percent Change: -38.94%
DPTUPR	DPTUPR	RETAINED	-4,284	-2,112,745	-0.187570
PRJSEL	PRJSEL	RETAINED	-1,873	-1,037,176	-0.090531
STRSEL	STRSEL	RETAINED	-5,184	-1,690	-0.016235
PRJANO	PRJANO	RETAINED	-1,690	-1,242	-0.074590
CHCPADB	CHCPADB	RETAINED	-121	-164	-0.004759
EMPUPD	EMPUPD	RETAINED	-2,298	-904	-0.044093
DSNADMCD	DSNADMCD	RETAINED	-108	-2,489	-0.232317
EMPUPR	EMPUPR	RETAINED	-1,242	-104,280	-33.813864
DPTANO	DPTANO	RETAINED	-904	-30	-0.002591
DPTUPD	DPTUPD	RETAINED	-2,489		
PRJADD	PRJADD	RETAINED	-104,280		
CHCDLRDB	CHCDLRDB	RETAINED	-30		

Analysis

General

Delays

Locks

Buffer

Elapsed Time

Value -2.891822

Unaccounted Time

Value 0.000000

Overall 0.00%

Specific Elapsed

Value -0.101926

Overall 3.52%

CPU time

Value -1.263498

Overall 44.38%

Total Delay Time

Value -1.548210

Overall 53.54%

Other Delay

Value +0.000031

Overall -0.00%

Percent -0.00%

Total Lock Delay

Value +0.000001

Overall -0.00%

Percent -0.00%

Servtask Switch Delay Time

Range: 1-25 Total: 42 Selected: 1

Page 1

Rows 25





Host variable Collection & Selectivity Override



- Why did the DB2 Optimizer choose that path?
- Helps users improve query access plans for dynamic queries with parameter markers
- The selectivity override feature utilizes parameter marker information
- Users can deploy a selectivity profile generated by this function to create better access plans.

```
SELECT * FROM  
EMPLOYEES WHERE  
SALARY BETWEEN ?  
AND ?
```





Create a Baseline

- Run a test application now to get a baseline.
- Average execution time for this application is: 125ms
- Note this query is well tune before selectivity override analysis
- Remember this number

Execution Log

ID: 1484464703338
Name: Appl_Ad_Hoc_4_min_01
Database Name: WKLDDVR01
Start Time: 1484600890461
End Time: 1484601132118
Result: The job executed successfully.

Commands Executed:

```
cd /tmp;  
/usr/bin/java -jar /home/db2inst1/hol/queryTuning/App1_Ad_Hoc_01.jar 4 0 01;
```

Output from executing:

```
Connected to DEMOMVS  
Running program for: 4 minute(s)  
Starting time: 13:08:11:925  
.....  
Ending time: 13:12:12:005  
  
Execution count: 1908  
Average execution time: 125ms
```





Collect Activity

Summaries

Access and refine view of your system's query activity

Collection Period

Choose the Activity Browser data from a past time period

Collect Host Variables

Configure how the host variable information is captured

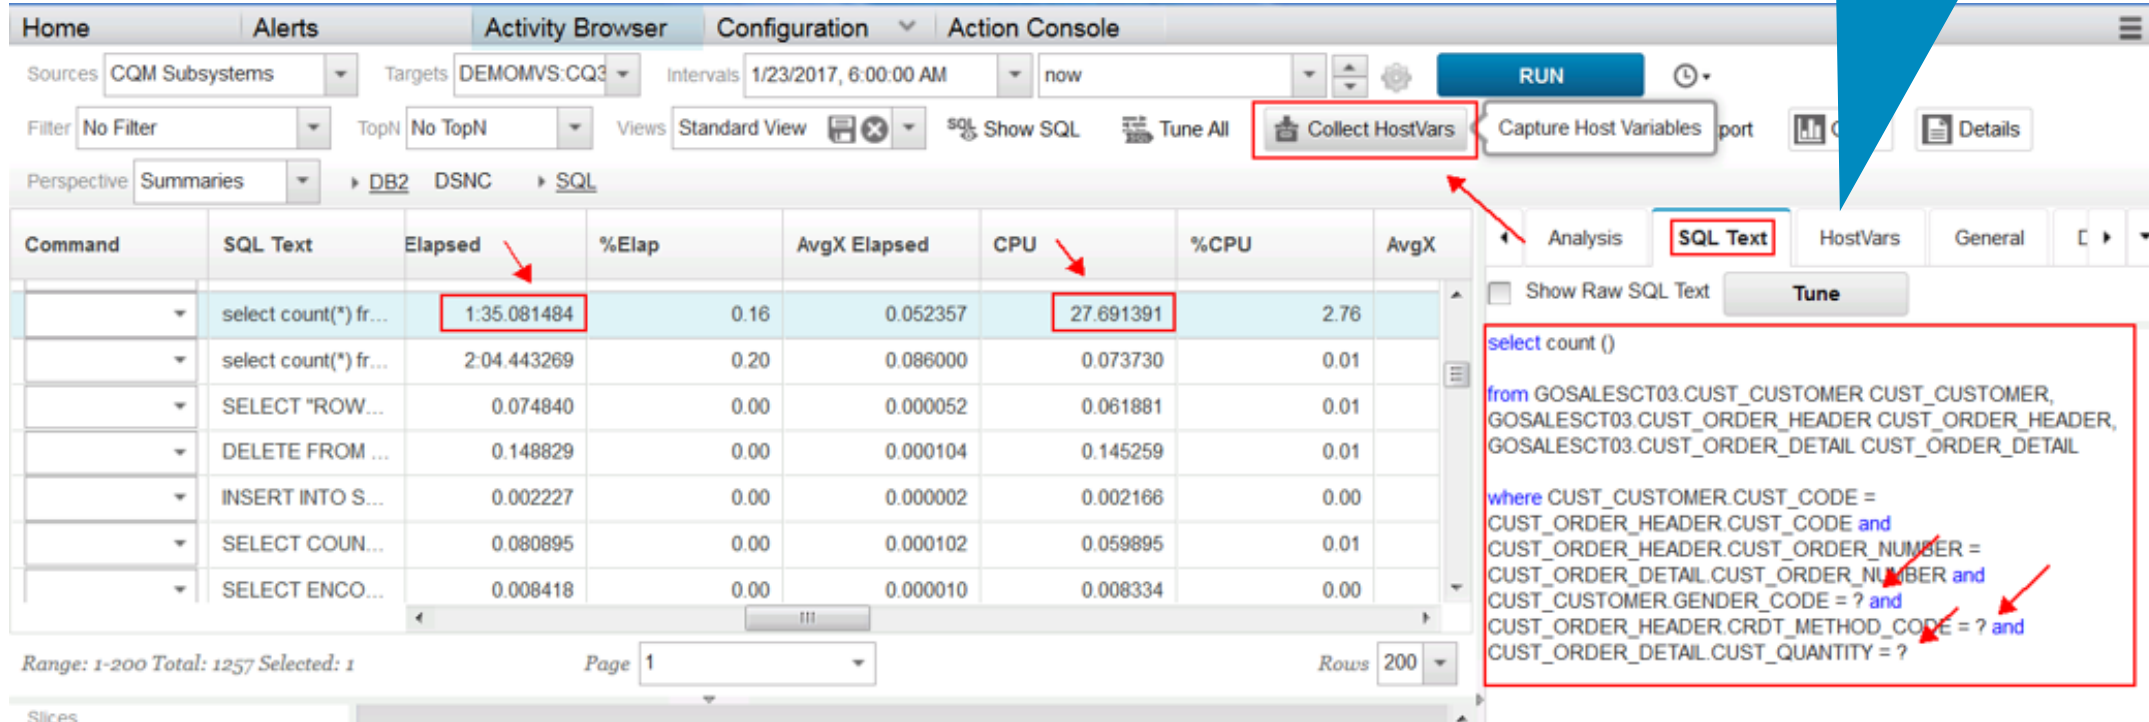
Request host variable collection



Analyze Hostvars details and identify candidate query

Tune or Tune all
For selectivity override
analysis

With parameter marker, high elapse time, CPU time, execution count etc.



The screenshot displays the IBM Analytics Activity Browser interface. The top navigation bar includes tabs for Home, Alerts, Activity Browser, Configuration, and Action Console. The Activity Browser tab is active, showing a table of SQL queries. The table has columns for Command, SQL Text, Elapsed, %Elap, AvgX Elapsed, CPU, %CPU, and AvgX. The first row is highlighted, showing a high elapsed time of 1:35.081484 and a CPU time of 27.691391. The 'Collect HostVars' button is highlighted in the top toolbar. The 'SQL Text' tab is selected, showing the query details with parameter markers.

Command	SQL Text	Elapsed	%Elap	AvgX Elapsed	CPU	%CPU	AvgX
select count(*) fr...	select count(*) fr...	1:35.081484	0.16	0.052357	27.691391	2.76	
select count(*) fr...	select count(*) fr...	2:04.443269	0.20	0.086000	0.073730	0.01	
SELECT "ROW...	SELECT "ROW...	0.074840	0.00	0.000052	0.061881	0.01	
DELETE FROM ...	DELETE FROM ...	0.148829	0.00	0.000104	0.145259	0.01	
INSERT INTO S...	INSERT INTO S...	0.002227	0.00	0.000002	0.002166	0.00	
SELECT COUN...	SELECT COUN...	0.080895	0.00	0.000102	0.059895	0.01	
SELECT ENCO...	SELECT ENCO...	0.008418	0.00	0.000010	0.008334	0.00	

Range: 1-200 Total: 1257 Selected: 1 Page 1 Rows 200

SQL Text

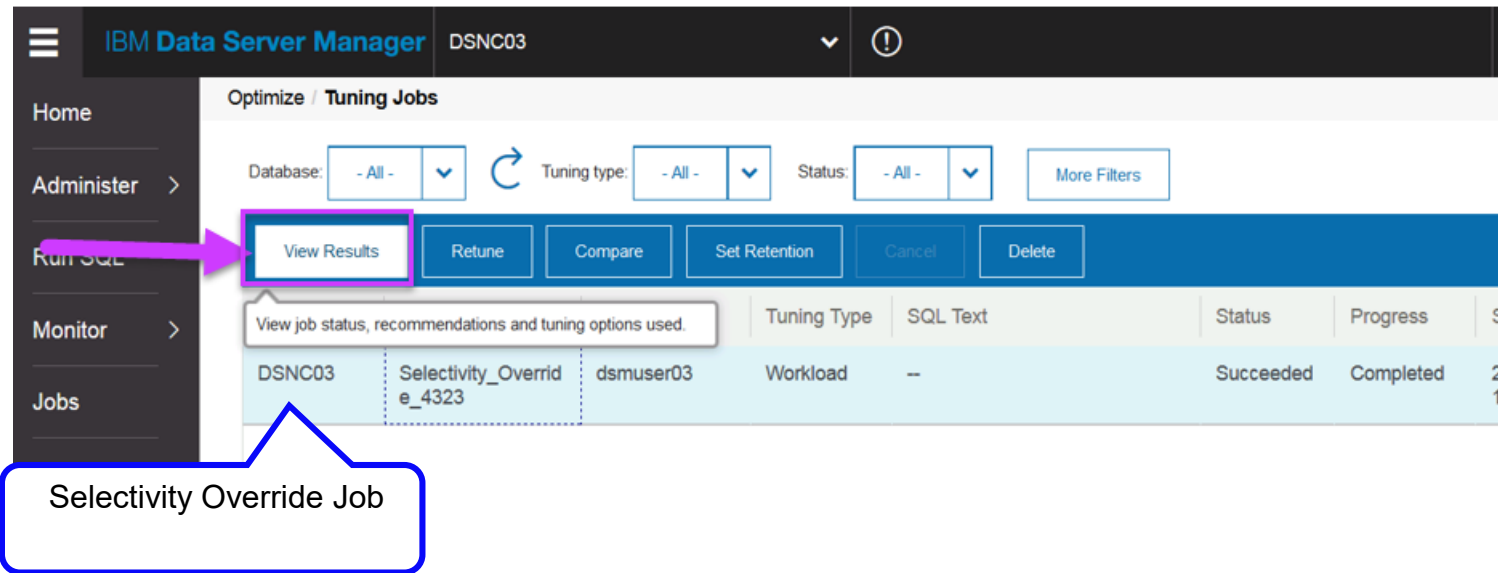
```
select count ()  
from GOSALESC03.CUST_CUSTOMER CUST_CUSTOMER,  
GOSALESC03.CUST_ORDER_HEADER CUST_ORDER_HEADER,  
GOSALESC03.CUST_ORDER_DETAIL CUST_ORDER_DETAIL  
where CUST_CUSTOMER.CUST_CODE =  
CUST_ORDER_HEADER.CUST_CODE and  
CUST_ORDER_HEADER.CUST_ORDER_NUMBER =  
CUST_ORDER_DETAIL.CUST_ORDER_NUMBER and  
CUST_CUSTOMER.GENDER_CODE = ? and  
CUST_ORDER_HEADER.CRDT_METHOD_CODE = ? and  
CUST_ORDER_DETAIL.CUST_QUANTITY = ?
```



42



Tune selectivity override



IBM Data Server Manager DSNCO3

Optimize / Tuning Jobs

Database: - All - Tuning type: - All - Status: - All - More Filters

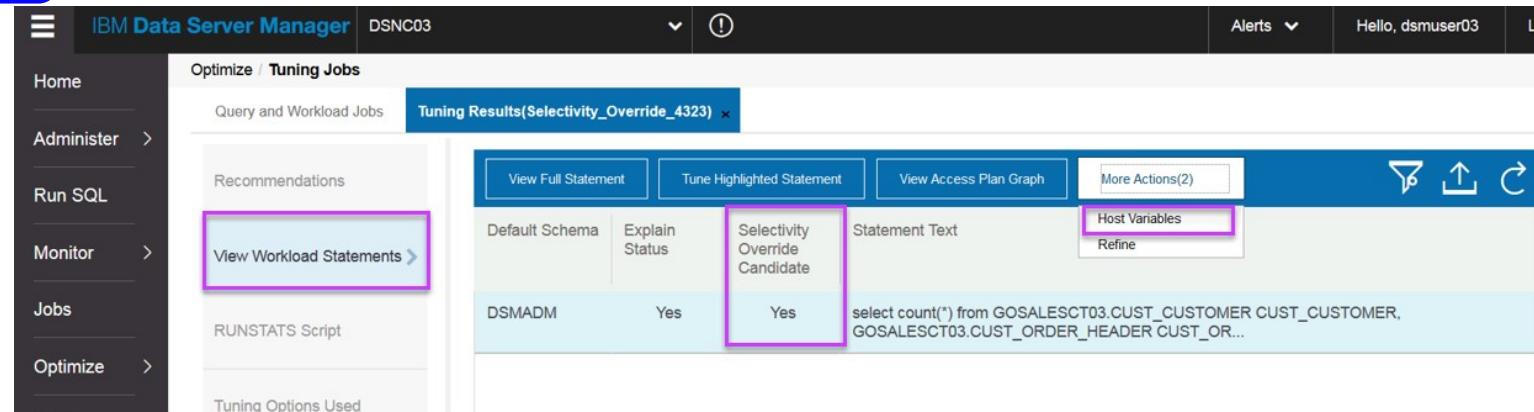
View Results Retune Compare Set Retention Cancel Delete

View job status, recommendations and tuning options used.

	Tuning Type	SQL Text	Status	Progress	St
DSNCO3	Selectivity_Override_4323	dsmuser03	Workload	--	Succeeded Completed 2012

Selectivity Override Job

- Go to View Workload Statements, you can see the query is Selectivity Override Candidate



IBM Data Server Manager DSNCO3

Optimize / Tuning Jobs

Query and Workload Jobs Tuning Results(Selectivity_Override_4323)

Recommendations

View Full Statement Tune Highlighted Statement View Access Plan Graph More Actions(2)

Default Schema Explain Status Selectivity Override Candidate Statement Text Host Variables Refine

DSMADM	Yes	Yes	select count(*) from GOSALEST03.CUST_CUSTOMER CUST_CUSTOMER, GOSALEST03.CUST_ORDER_HEADER CUST_OR...	
--------	-----	-----	--	--

RUNSTATS Script

Tuning Options Used

- Then, select Host Variables



Review analysis

- In this dialog, you can see:
 - parameter markers distribution
 - Weight of each parameter marker value set
- Select the sets (all) for Selectivity Override analysis
- Click Selectivity Override

Host variable/Parameter marker values

[Learn more](#)

Selectivity Override		View Full Statement					
<input checked="" type="checkbox"/>	Set	Count	Weight	Position	Predicate Text	Candidate	Value
<input checked="" type="checkbox"/>	▼	1	18	0.18			
				1	"CUST_CUSTOMER"."GENDER_CODE" = ?	Yes	1
				2	"CUST_ORDER_HEADER"."CRDT_METHOD_CODE" = ?	Yes	29
				3	"CUST_ORDER_DETAIL"."CUST_QUANTITY" = ?	Yes	1
<input checked="" type="checkbox"/>	▶	2	15	0.15			
<input checked="" type="checkbox"/>	▶	3	14	0.14			

Execution count: 100
Total literal value sets: 13 Selected: 13

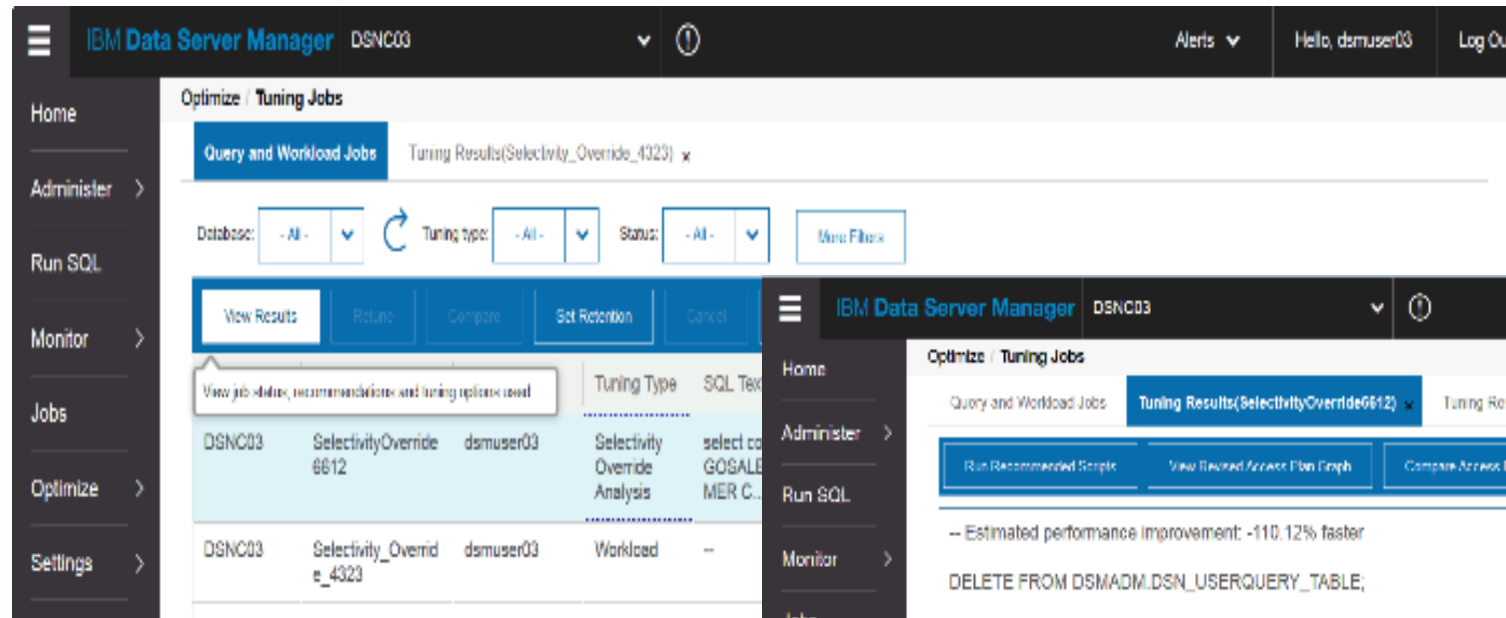
[Cancel](#)

optimize



View results and deploy the selectivity profile

- A selectivity override analysis job is created
 - Click View Results when it is completed



IBM Data Server Manager DSNCO3

Optimize / Tuning Jobs

Query and Workload Jobs Tuning Results(Selectivity_Override_4323) x

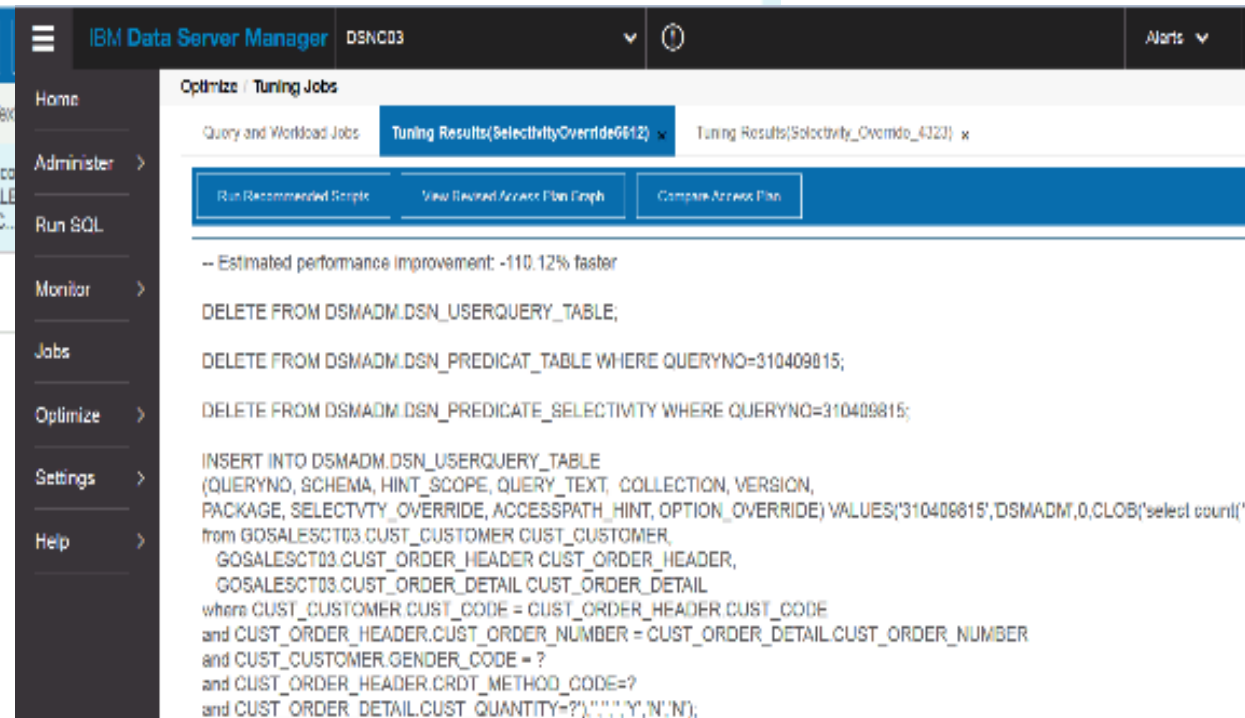
Database: - All - Tuning type: - All - Status: - All - View Filter

View Results Retune Compare Set Retention Cancel

View job status, recommendations and tuning options used

Database	Job Name	User	Tuning Type	SQL Text
DSNCO3	SelectivityOverride 6612	dsmuser03	Selectivity Override Analysis	select co...
DSNCO3	Selectivity_Override_4323	dsmuser03	Workload	--

- Run recommended scripts
- Flush the statement cache



IBM Data Server Manager DSNCO3

Optimize / Tuning Jobs

Query and Workload Jobs Tuning Results(SelectivityOverride6612) x Tuning Results(Selectivity_Override_4323) x

Run Recommended Scripts View Recommended Access Plan Compare Access Plan

-- Estimated performance improvement: -110.12% faster

```
DELETE FROM DSMADM.DSN_USERQUERY_TABLE;

DELETE FROM DSMADM.DSN_PREDICAT_TABLE WHERE QUERYNO=310409815;

DELETE FROM DSMADM.DSN_PREDICATE_SELECTIVITY WHERE QUERYNO=310409815;

INSERT INTO DSMADM.DSN_USERQUERY_TABLE
(QUERYNO, SCHEMA, HINT_SCOPE, QUERY_TEXT, COLLECTION, VERSION,
PACKAGE, SELECTVTY_OVERRIDE, ACCESSPATH_HINT, OPTION_OVERRIDE) VALUES('310409815', DSMADM, 0, CLOB('select count(*)
from GOSALESC03.CUST_CUSTOMER CUST_CUSTOMER,
GOSALESC03.CUST_ORDER_HEADER CUST_ORDER_HEADER,
GOSALESC03.CUST_ORDER_DETAIL CUST_ORDER_DETAIL
where CUST_CUSTOMER.CUST_CODE = CUST_ORDER_HEADER.CUST_CODE
and CUST_ORDER_HEADER.CUST_ORDER_NUMBER = CUST_ORDER_DETAIL.CUST_ORDER_NUMBER
and CUST_CUSTOMER.GENDER_CODE = ?
and CUST_ORDER_HEADER.CRD_METHOD_CODE=?
and CUST_ORDER_DETAIL.CUST_QUANTITY=?'), 'Y', 'N', 'N');
```





Compare against baseline

- Run the test application again run after Selectivity Override analysis
- Average execution time for this application is: 92ms
- Improvement of 26%
- On an already well-tuned query!!

Execution Log

ID: 148461703
Name: Appl
Database Name:
Start Time: 148461703
End Time: 148461703
Result: The job executed successfully.

➤ 26% faster!!

Commands Executed:

```
cd /tmp;  
/usr/bin/java -jar /home/db2inst1/hol/queryTuning/Appl_Ad_Hoc_01.jar 4 0 01;
```

Output from executing:

```
Connected to DEMOMVS  
Running program for: 4 minute(s)  
Starting time: 18:17:01:823  
.....  
Ending time: 18:21:01:913
```

Execution count: 2595


Average execution time: 92ms



apply

IBM delivers complete DB2 performance management

- Reduce costs of DB2 for z/OS and applications
 - Improve performance of all package applications
 - Tune performance of query warehouse
- Identify and solve faster closing the loop on problem resolution
- Replace ad-hoc methods with integrated solutions for scalable, robust approach to performance management
- Improve performance and time to resolution by up to 50%
- Speed DB2 and application migration with comprehensive comparison capabilities



Learn how to use Data Server Manager to tune your DB2 on z/OS Queries

Now available on IDUG.org



IDUG
Leading the DB2 User
Community since 1988

IDUG Data Tech Summit
May 1-2, 2017 | Anaheim, California

Meet us at IDUG NA

ABOUT IDUG **MEMBERSHIP** **EVENTS** **FORUMS** **RESOURCES** **ARTICLES & CONTENT** **GET INVOLVED**

IDUG Content Blog

Use Data Server Manager to tune your DB2 on z/OS Queries for Free



Saghi Amirsoleymani

IBM® Data Server Manager is an integrated database management tools platform for DB2® for z/OS databases. It was designed to be simple to set up, easy to use, and enterprise ready with the ability to manage hundreds of databases.

Data Server Manager Base Edition offers database development, administration and basic performance capabilities to all DB2 for Linux, UNIX and Windows and DB2® for z/OS® clients at no charge.

For DB2® for z/OS®, Data Server Manager Base also offers advanced capabilities with centralized configuration management available in the DB2 Administration Solution Pack for z/OS. DB2 Query Workload Tuner for z/OS takes advantage of Data Server Manager Architecture either as an individual product or component of the DB2 Performance Solution Pack for z/OS.

Query optimization with Data Server Manager

Data Server Manager 2.1 is now available at no additional cost to help you optimize queries to

Recent Stories

R You Ready to be a Data Scientist?

Updated 10:06AM CST, Mon Feb 13th, 2017



In this popular session from the 2016 IDUG North American Conference, IDUG President Paul Turpin answers questions about R, Data Science and DB2.

Learning Machine Learning

Updated 3:01PM CST, Thu Feb 2nd, 2017



Machine Learning is a rapidly growing field of the computer science, which has

Thank
You