sqlAdria

SQLADRIA SEMINAR

Ljubljana, 15 April 2011

Sponsored by

IBM

SQLAdria Seminar

10.00 Opening

10.15 Content Management (Digital archiving)

Grgo Miočić

IBM

Content Management (Content Manager on Demand)

Grgo Miočić

IBM

11.30 Coffee Break

11.45 SQL Performance Good and Bad Access Paths

Sašo Prek

IBM

Workload Tuning

Sašo Prek

IBM

13.15 Lunch

14.30 User Topic: DB2 Explain

Bernard Zver

Informatika

15.00 Open Forum

Sponsored by:

Ljubljana, 15th April 2011

Hotel Slon

Participants of the SQLAdria Seminar
Ljubljana, 15th April 2011

	Name	Surname	Company	E-mail
1	Bernard	Zver	Informatika d.d.	bernard.zver@informatika.si
2	Mateja	Jankovič	Informatika d.d.	mateja.jankovic@informatika.si
3	Matej	Jozelj	Informatika d.d.	matej.jozelj@informatika.si
4	Jana	Perko	Informatika d.d.	jana.perko@informatika.si
5	Vida	Labus	Informatika d.d.	vida.labus@informatika.si
6	Marjana	Glavan	Informatika d.d.	marjana.glavan@informatika.si
7	Lovro	Pavšič	Informatika d.d.	lovro.pavsic@informatika.si
8	Sanja	Vičić	Apis IT d.o.o.	sanja.vicic@apis-it.hr
9	Nina	Petanjek Fureš	Apis IT d.o.o.	nina.petanje-fures@apis-it.hr
10	Lidija	Penezić	Apis IT d.o.o.	lidija.penezic@apis-it.hr
11	Biljana	Dragišić	Apis IT d.o.o.	biljana.dragisic@apis-it.hr
12	Borislav	Njegovan	Apis IT d.o.o.	borislav.njegovan@apis.hr
13	Igor	Mlakar	Nova ljubljanska banka d.d.	igor.mlakar@nlb.si
14	Branko	Božič	Nova ljubljanska banka d.d.	branko.bozic@nlb.si
15	Marija	Ankon	Nova ljubljanska banka d.d.	marija.ankon@nlb.si
16	Franci	Drnovšek	Nova ljubljanska banka d.d.	franci.drnovsek@nlb.si
17	Roman	Bučar	Nova ljubljanska banka d.d.	roman.bucar@nlb.si
18	Stane	Berdajs	Nova ljubljanska banka d.d.	stane.berdajs@nlb.si
19	Igor	Miškulin	Nova ljubljanska banka d.d.	igor.miskulin@nlb.si
20	Matej	Kafadar	Nova ljubljanska banka d.d.	matej.kafadar@nlb.si
21	Samo	Herman	Nova ljubljanska banka d.d.	samo.herman@nlb.si
22	Tomaž	Strgar	Nova ljubljanska banka d.d.	tomaz.strgar@nlb.si
23	Tatjana	Lajovic Šink	Nova ljubljanska banka d.d.	tatjana.lajovic@nlb.si
24	Borut	Zavodnik	Nova ljubljanska banka d.d.	borut.zavodnik@nlb.si
25	Goran	Bavčar	Nova ljubljanska banka d.d.	goran.bavcar@nlb.si
26	Zdravko	Kelšin	Nova ljubljanska banka d.d.	zdravko.kelsin@nlb.si
27	Alojz	Mestek	ZZZS Slovenija	alojz.mestek@zzzs.si
28	Olga	Dragan - Frlež	ZZZS Slovenija	olga.dragan.frlez@zzzs.si
29	Dragan	Momčilović	ZZZS Slovenija	dragan.momcilovic@zzzs.si
30	Tadej	Poljšak	ZZZS Slovenija	tadej.poljsak@zzzs.si
31	Janez	Miko	ZZZS Slovenija	janez.miko@zzzs.si
32	Mateja	Ošina	ZZZS Slovenija	mateja.osina@zzzs.si
33	Tatjana	Šumec	ZZZS Slovenija	tatjana.sumec@zzzs.si
34	Željko	Blagaić	HZMO	zeljko.blagaic@mirovinsko.hr
35	Iva	Mesarić	HZMO	iva.mesaric@mirovinsko.hr
36	Goran	Petrinjac	HZMO	goran.petrinjac@mirovinsko.hr
37	Ana	Ramljak	Zagrebačka banka	ana.ramljak@unicreditgroup.zaba.hr
38	Tomislav	Žarković	Zagrebačka banka	tomislav.zarkovic@unicreditgroup.zaba.hr
39	Zoran	Heged	Zagrebačka banka	zoran.heged@unicreditgroup.zaba.hr
40	Goran	Mendaš	Zagrebačka banka	goran.mendas@unicreditgroup.zaba.hr
41	Damir	Jelić	Zagrebačka banka	damir.jelic@unicreditgroup.zaba.hr
42	Mirna	Kos	Zagrebačka banka	mirna.kos@unicreditgroup.zaba.hr
43	Željen	Stanić	CA	zeljen.stanic@ca.com
44	Marko	Koprivnikar	IBM Slovenija d.o.o.	marko.koprivnikar@gmail.com
45	Sašo	Prek	IBM Slovenija d.o.o.	saso.prek@si.ibm.com
46	Grgo	Miočić	IBM Hrvatska d.o.o.	grgo.miocic@hr.ibm.com
47	Olivera	Stanić	CTK SQLAdria	ostanic@ctk-rijeka.hr

SQLADRIA SEMINAR

Ljubljana, 15 April 2011

Content Management (Digital archiving)

Grgo Miočić
IBM

Sponsored by

IBM FileNet Advance Case Management

Grgo Miocic

IBM Advance Case Management

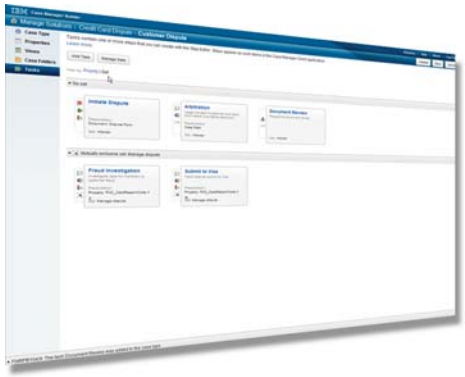

Case Design for the Business

- Easy to use, wizard-driven
- Comprehensive across case assets
- Leverage templates for a fast start
- Significantly shortens time-to-value for case-style applications

Case Application User Experience

- Role-based and personalized
- Flexible and extensible
- Provides deep context for case work
- Brings people, process and information together to drive case progression and better outcomes

When do I need Case Management?

A solution pattern where…

- There are **collaborative**, ad-hoc processes
- Activities are **event-driven**
- Work is **knowledge** intensive
- **Content** is essential for decision making
- **Outcomes** are goal-oriented
- The judgment of people impact how the **goal** is achieved
- **Process** is often not predetermined

For challenges such as…

- Complex decision making
- Complaint or dispute management
- Contract management
- Lending applications
- Benefits enrolment
- Rate case management
- …that are both **horizontal** and **vertical** in nature

That requires unique capabilities from…

- Content management
- Business process management
- Collaboration tools
- Social software
- Business rules
- Analytics

Different types of 'Cases' across industries
…. Similar patterns and challenges

Invoice, contract, employee, vendor, customer, project, change request, complaint, exception, incident, audit, etc.

INSURANCE

- Policy
- Underwriting
- Claim
- Annuity

BANKING

- Loan
- Dispute
- Mortgage
- Account
- Credit card
- Personal line
- Investment

HEALTHCARE

- Patient case
- Member
- Provider

GOVERNMENT

- Benefits enrolment
- Grant
- Court case
- Citizen
- Taxpayer

ENERGY

- Rate case
- Claim
- Permit
- Land
- Property

What makes IBM Case Manager Solution?

Full Case Context

Documents
- Forms
- Correspondence
- Receipts
- Statements

History
- Case Created
- Document Added
- Comment
- Review started

Properties
- Customer Name
- Priority
- Transaction IDs
- Account Number

Tasks
- Initial Case Review
- Email the Customer
- Investigate Claim
- Issue Refund

What makes IBM Case Manager Solution?

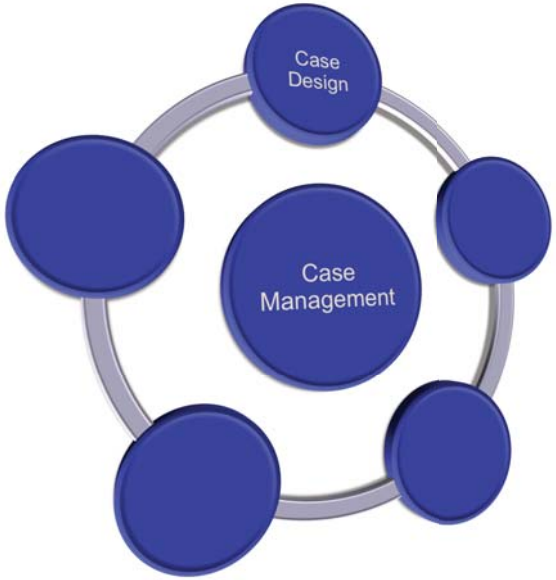

Rapid application case design for the business

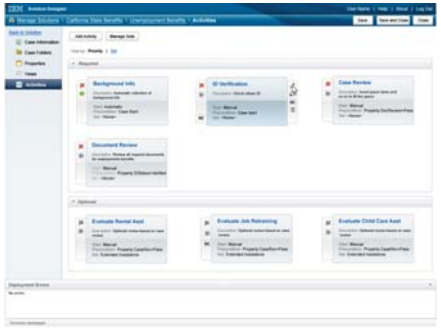

Easy to use, "Interview Mode" (wizard-driven)
Case designer allows a business user to very quickly build a solution

Comprehensive across case assets
Case designer can provide 360⁰ view of case

Leverage templates for a fast start
Represent industry best practices

Significantly shortens time-to-value for case-style applications

Bottom line:
Deliver end users the solution in a fraction of time of other approaches

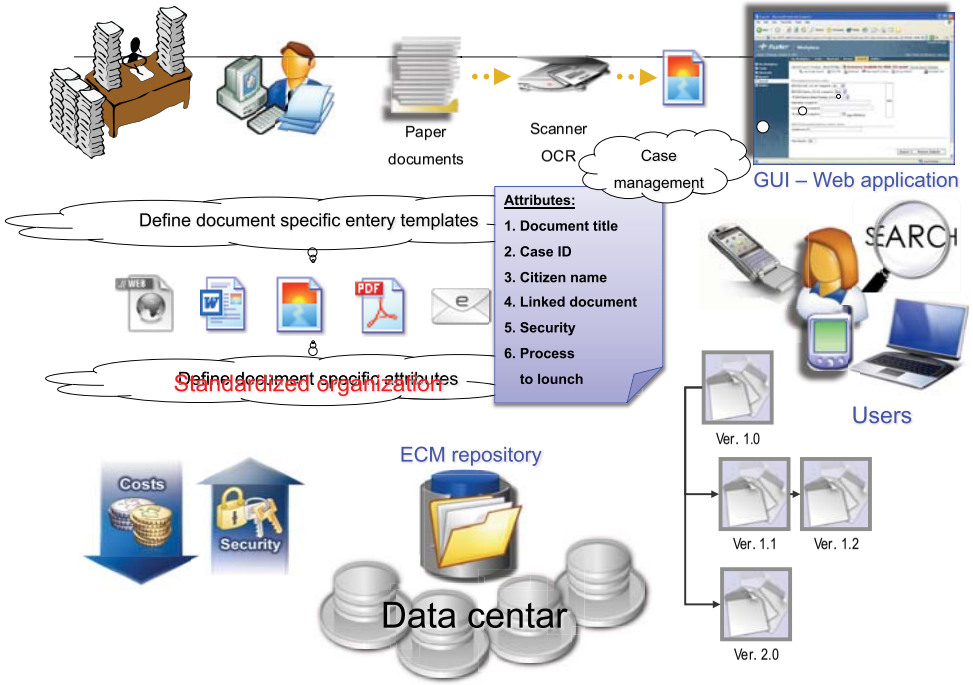

"Active Content"="Activating" Work

IBM

- Event driven architecture
  - Immediate reaction on business events
- Reaction can be:
  - Start process
  - Send e-mail
  - Change security
  - ….

Essential for Case Management applications

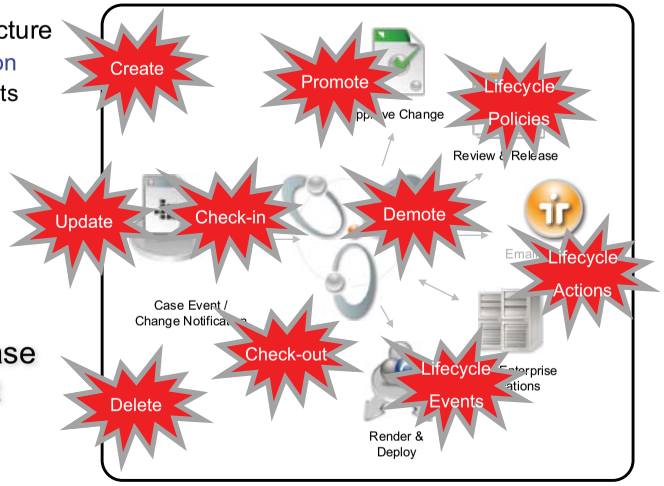

Case Design & Business Process Management

IBM

Business analysts can define many aspects of a case management solution using a single, targeted tool…

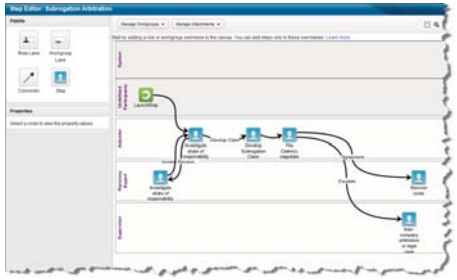

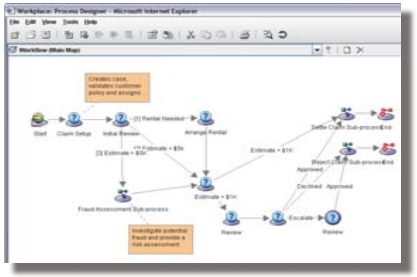

… technical experts can work with the exact same solution using specialized tools to build data integrations and complex decision logic

Claims Processing Case Design

IBM

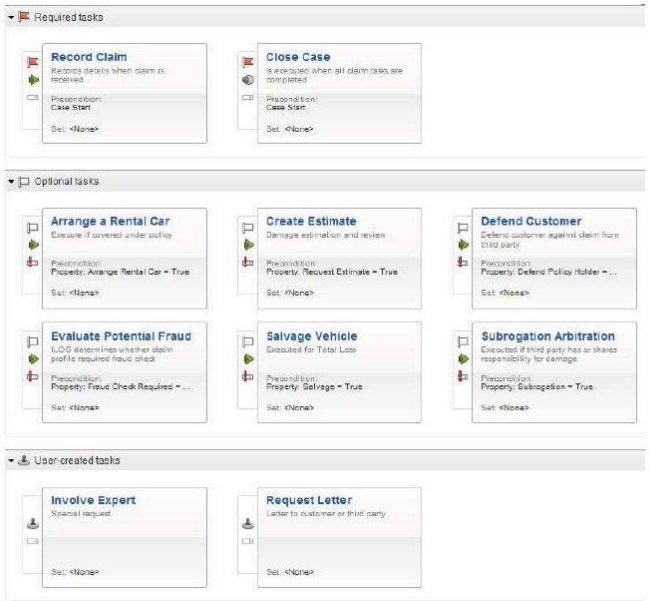

Ad Hoc Tasks - Cases are Dynamic

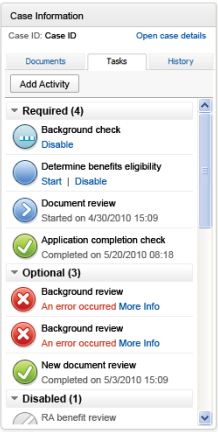

Case workers can view and start new Tasks for a case, changing how that case is handled on-the-fly

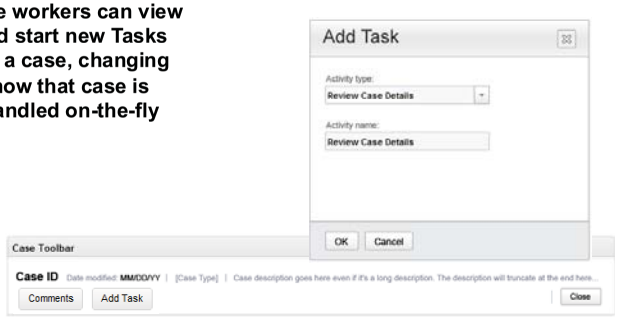

Case workers can even add new Tasks to a case

Advanced Rule Management

Decision Rule

```
if the yearly repayment of 'the loan' is more than
    the yearly income of 'the borrower' * 0.3
then
  add "Too big Debt-To-Income ratio" to the
     messages of 'the loan' ;
  reject 'the loan' ;
```

Decision Table

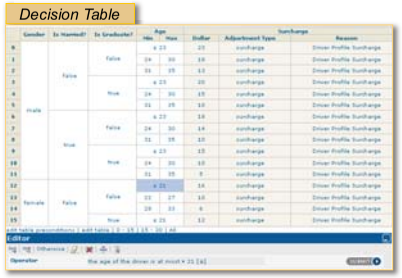

A statement of business logic that:

- ✓ Is not cryptic code!
- ✓ Can be authored by business users
- ✓ Is self-documented, executable, reusable

Decision Tree

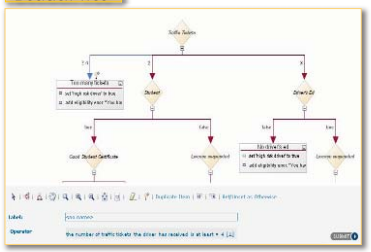

What makes IBM Case Manager Solution?

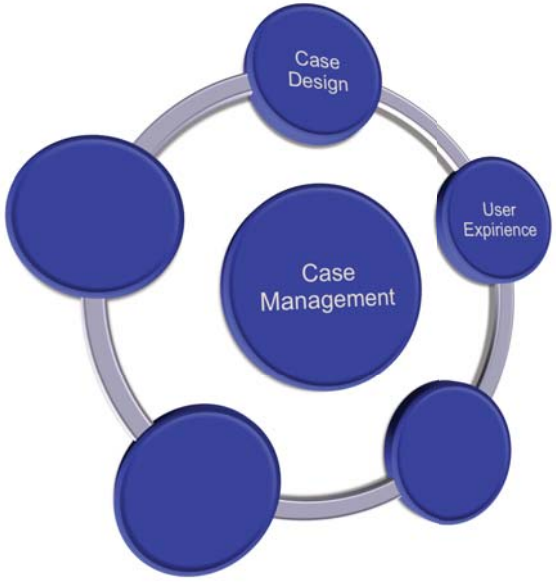

Case worker user experience

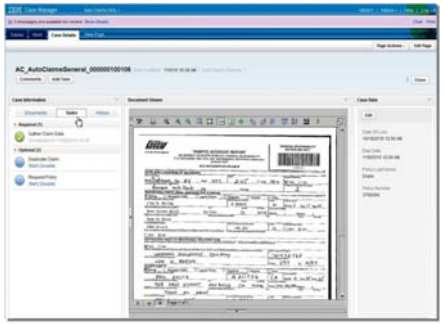

Role-based and personalized
End user gets exactly the information they need to progress the case

Flexible and extensible
Can be configured to meet unique business requirements

Provides deep context for case work
No more disjointed jumping between application

Brings people, process and information together to drive case progression and better outcomes

Bottom line:
A case worker has all the information they need to improve case outcomes

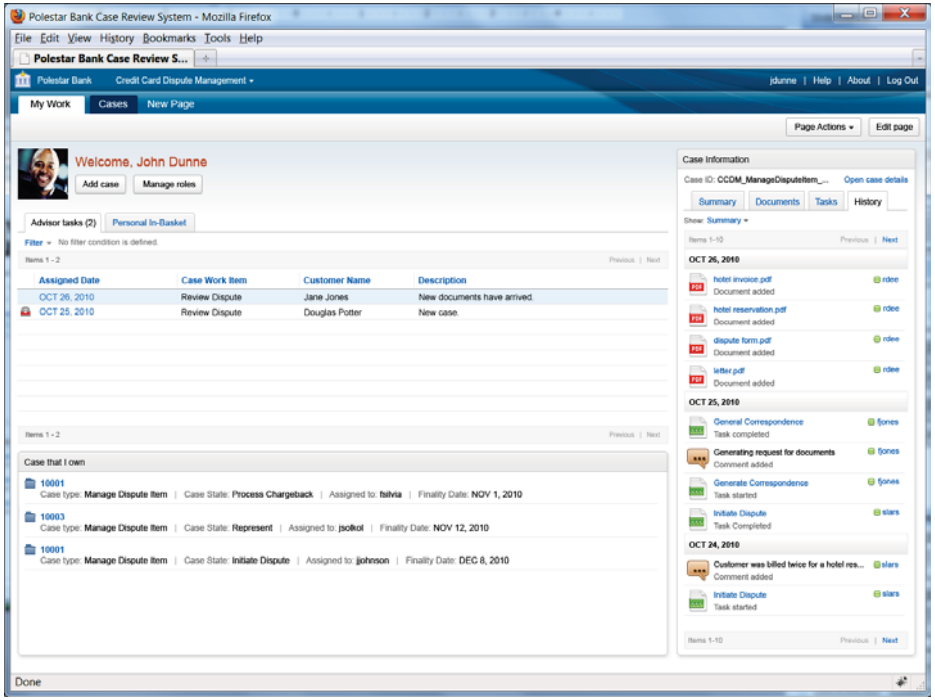

Add Fraud Case

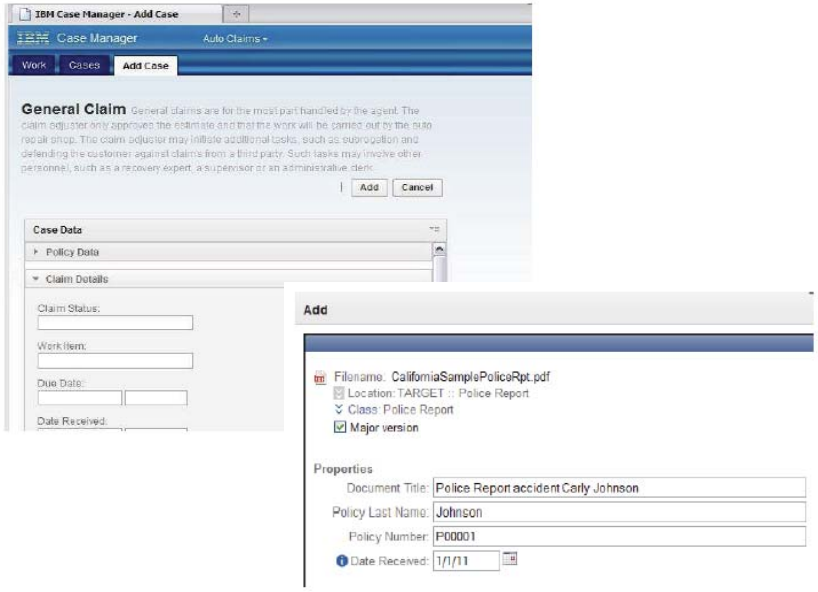

Case Tasks Inbox

1. Tasks can be sorted
2. Task can be filtered
3. Imortant data can be exposed for quick review
4. Decisions are faster

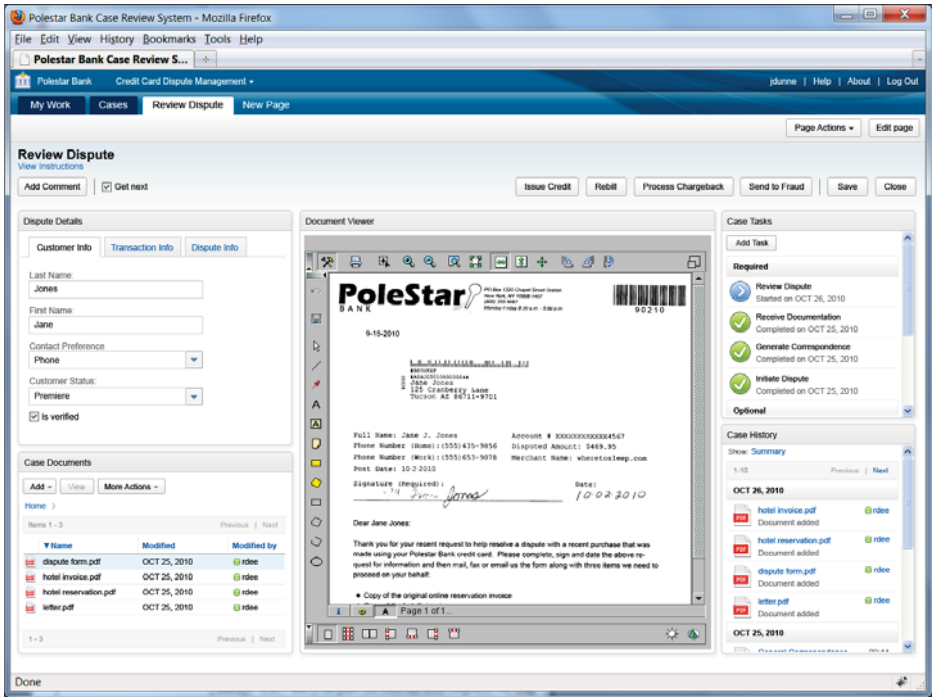

Searching Documents – Investigate

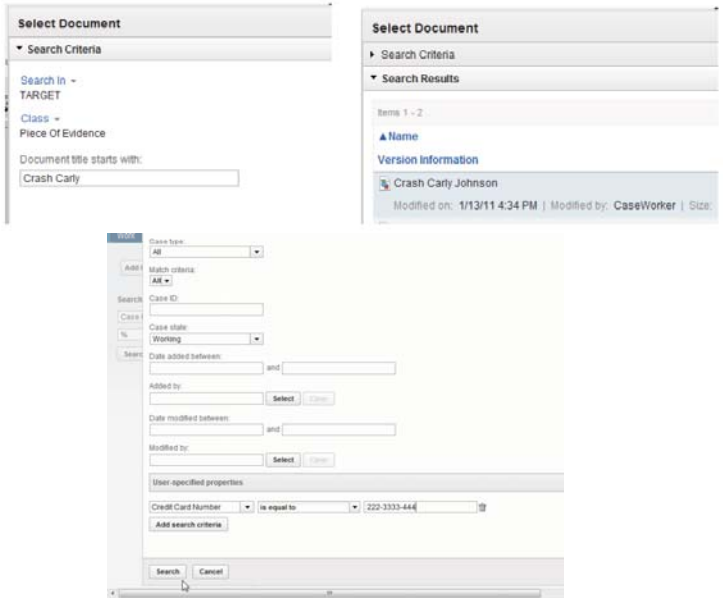

P8 eForms Integration

- eForm can be embedded in a step page
- Business value
  - Rapid solution development
  - Catch errors at the source
  - Reduce system complexity

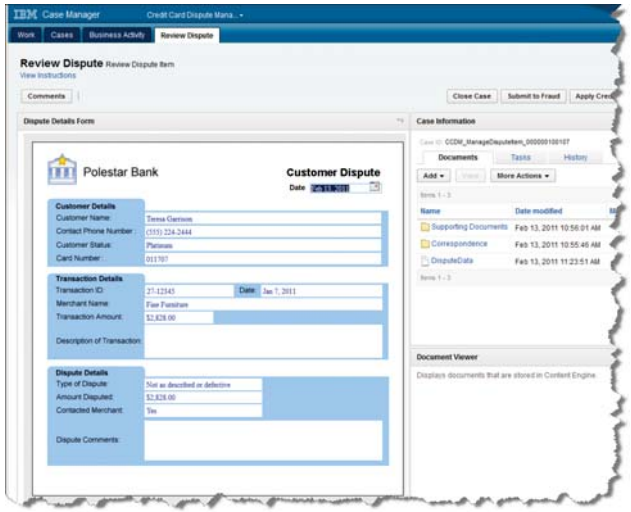

What makes IBM Case Manager Solution?

Cases are Collaborative

Case workers can locate and chat with each other directly through the Case Manager runtime

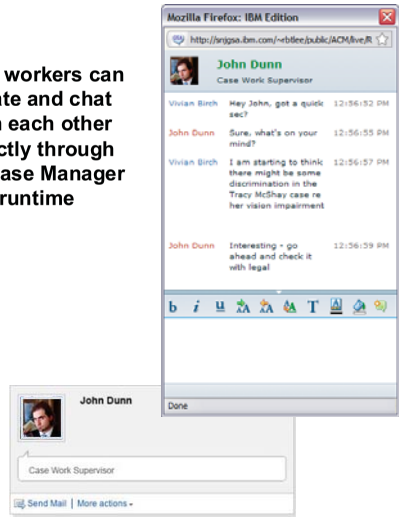

Shared comments and case history provide context

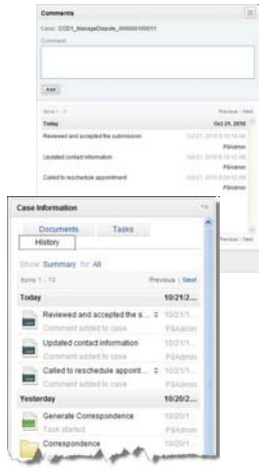

Simplified Business and IT Collaboration

Comments and sharing let business analysts and technical experts have a single, common view of a solution

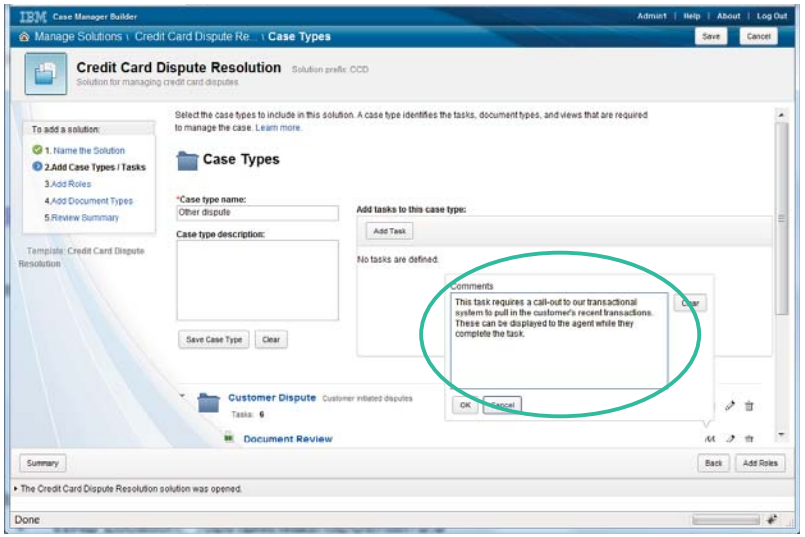

Content Collection for e-Mail

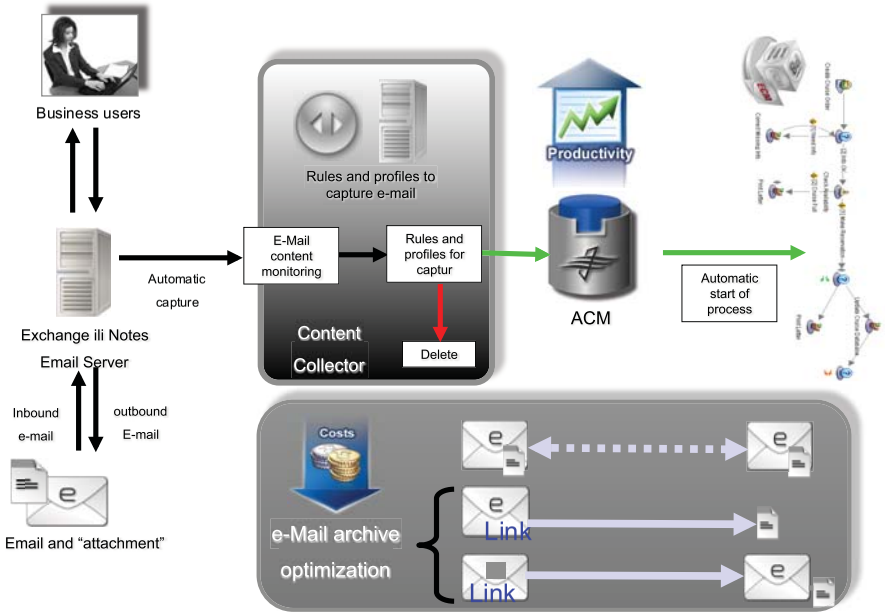

What makes IBM Case Manager Solution?

Powerful analytics for better case outcomes

Comprehensive reporting and analysis
Gives case managers visibility across all information types to assess and act quickly

Real-time dashboards
Understand issues before they become a problem

Unique content analytics for discovering deeper case insight

Bottom line:
Case managers need insight in order to impact results.

Case real-time acitivity monitoring

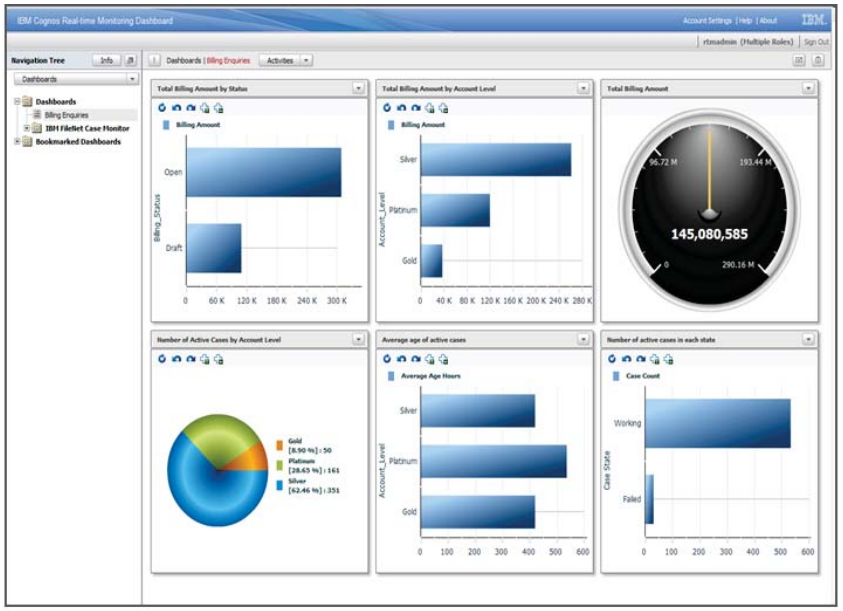

Case Analytics

Out of the box reports and content analytics enable business users to

- **review operational performance and**
- **determine the root causes of inefficiencies**

What makes IBM Case Manager Solution?

Integration & process orchestration

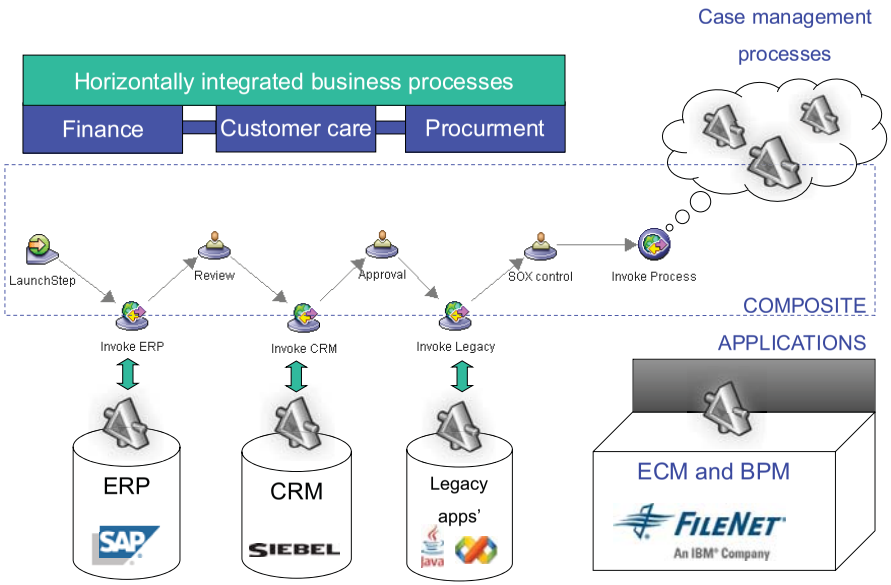

WebSphere Process Server Integration

- Associate Case Manager Tasks to WPS components
- Support for web services based authoring (WID) integration to WPS components

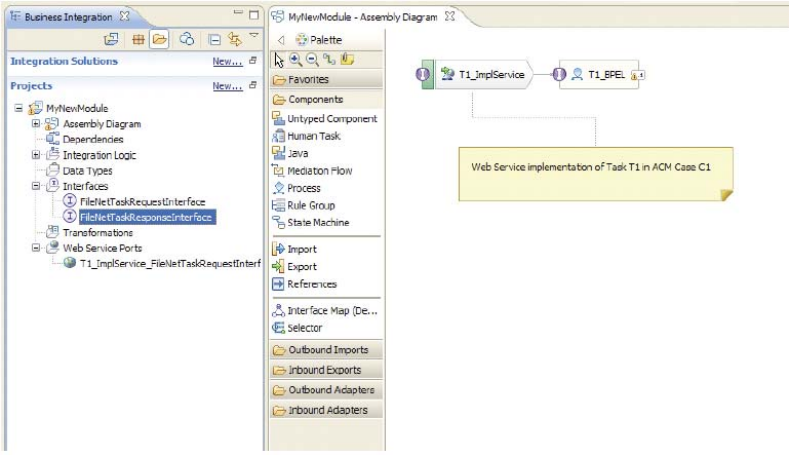

How IBM case management can help customers

"We need to treat business processes as assets. If you have ownership of your processes and manage them, you get a better return."

Rusty Williams
CIO
Good Samaritan

- ✓ *Improve customer service*
- ✓ *Improve case outcomes*
- ✓ *Speed delivery*
- ✓ *Create competitive differentiation*
- ✓ *Improve information accuracy*
- ✓ *Manage risk*
- ✓ *Assure best practices*
- ✓ *Foster collaboration*
- ✓ *Reduce errors*
- ✓ *Reduce costs*

© 2011 IBM Corporation

SQLADRIA SEMINAR
Ljubljana, 15 April 2011

Content Management (Content Manager on Demand)

Grgo Miočić
IBM

Sponsored by

Enterprise Content Management

IBM Content Manager on Demand

Grgo Miočić

Introducing Content Manager OnDemand

Archiving, Web Delivery, Distribution, Reprint, and much more…

Created initially as a Services Offering at the request of our AFP Print customers in the early 1990s, "OnDemand" has grown to be an integral part of IBM's Enterprise Content Management and Discovery portfolio.

✓OnDemand is the clear Enterprise Report Management leader

✓Proven technology that works on a wide variety of platforms

✓Flexible installation options

✓Large customer base, large skill base

✓Meets your Report Management needs today and tomorrow

Types of Content

Faxes

Complex Print Documents

PC Documents

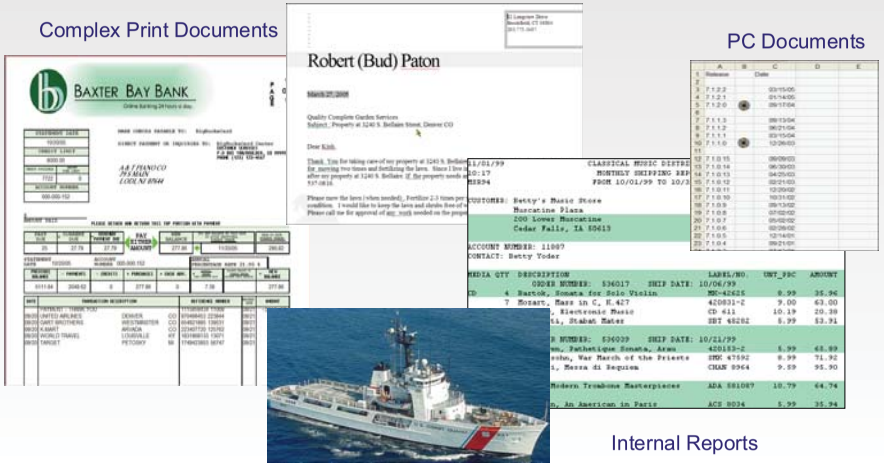

Internal Reports

Scanned Images

Content/Image versus Archive

P8 & CM8

- **Work in progress**
- **Content/Image System**
- **Paper intensive**
- **Individual documents**
  - Word
  - Scanned
  - Fax
- **Document management**
- **Creation**
- 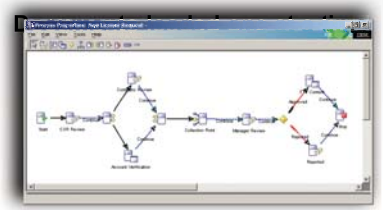

CMOD

- **After Process**
- **Long term archive**
- **Large datastreams with potentially 1,000,000+ documents daily**
- **Many documents loaded at the same time**
- **Little or no workflow**
- **NO report creation capability**
- **Documents loaded in batch**

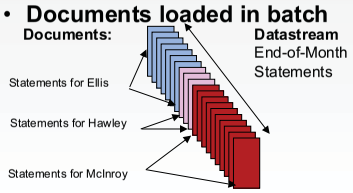

OnDemand Overview

- Client/Server document archival and retrieval system
  - **Servers**
    - Windows ™ - (WIN3K)
    - Linux – (RedHat, SuSE)
    - UNIX - (AIX,HP and Sun)
    - iSeries (i5/OS)
    - zSeries mainframe (z/OS, Linux for System z)
  - **Clients**
    - WEBi, Web Administration, Portlets (8.3 only)
    - Windows ™ XP SP2 or later, Vista
    - CICS/3270 client available as services offering
    - Internet clients CGI, Servlets, Java, IBM Integrator for Content (IICE)
    - Structured API's (z/OS)
    - DDE and OLE API's

iSeries
pSeries
xSeries
zSeries

OnDemand Server Structure

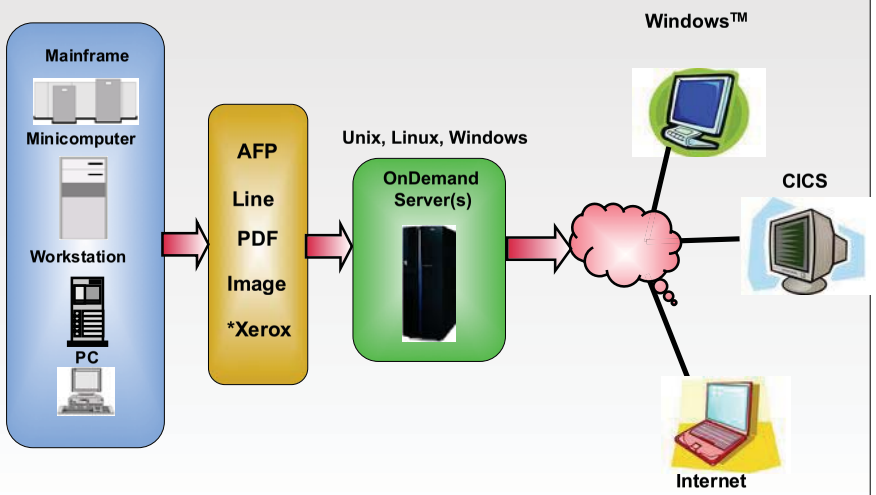

*Requires: Printstream Transform

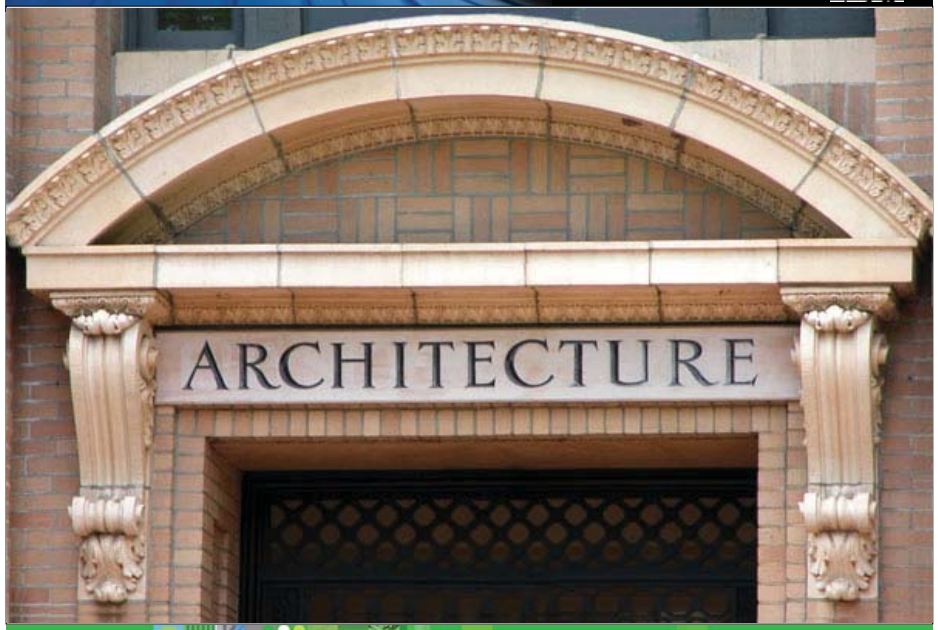

OnDemand Document Capture

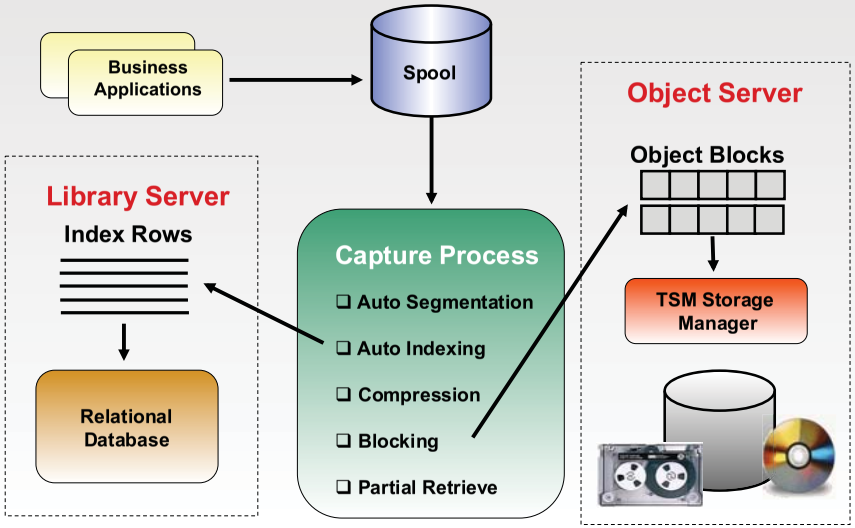

SEGMENTATION COMPRESSION INDEXING

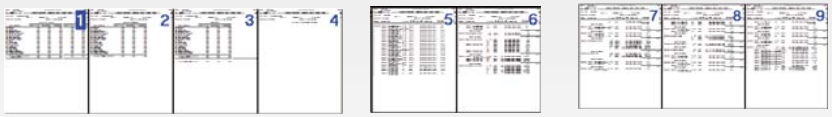

John Smith
123-11-2001

Ron Smith
523-22-6678

William Smith
709-03-4332

DATABASE

John Smith	123-11-2001
Ron Smith	523-22-6678
William Smith	709-03-4332

Storage Manager

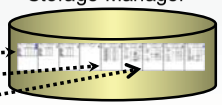

Searching f

Accou
12

Amou

Da
01/01

Hit list

ATLAS INC.

Invoice Number: 234568001

Date: 06/24/2006

Customer:
Bud Smith
1289 Smith Road
Boulder, CO 80301

Customer Number:
123-44-5555

Item	Description	Qty	Price
1	Appraisal	1	1700.00
2	Drawing	2	900.00
3	Wheel Rudder	3	400.00

Tax 180.00

Total Due: 500.00

e – meta data

	Date	Amount Due
	01/03/08	180.51
	08/19/07	0.0
	06/24/06	500.00
	12/25/05	207.82

ching…

ue Date Due

06/24/07

On-the-fly Conversion

CMOD

DB2 – Oracle – SQL Server

Name	SSN	Object	Offset
John Smith	709-03-4332	1FAA	0
Ron Smith	123-11-2001	1FAA	22560
William Smith	444-00-1289	1FAA	35120

Library Server

Object Server

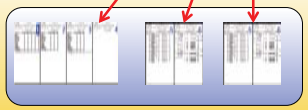

AFP Documents

Web Server
WAS
ODWEK
WEBi

AFP2PDF
Line2PDF

Browser Client
Viewing with
Adobe Reader

PDF Documents

INDEXING

- ACIF Indexer
  - Supports 370 and ASCII line data and AFP data
- OS/390 Indexer
  - Line and AFP data on mainframe
- PDF Indexer
  - Parses PDF files extracts index data at known locations
- Generic Indexer
  - The generic index file contains index field-id's and values and offset and file information about the associated object.
  - Used to load user misc. file types into OnDemand

Document Indexing and Segmentation

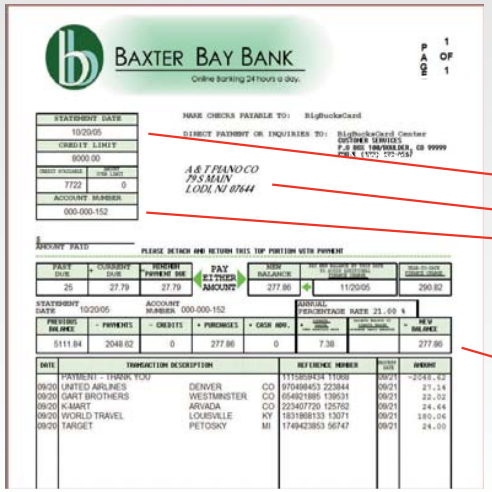

Segment print data by invoice number

Index Segments By
- Date
- Customer Name
- Customer Account

Add statement balance as additional display information

Graphical Indexing

D:\arsdata\loan.lin - Warning! This data should match the data being loaded.

```
REPORT    D94100100                PENNANT NATIONAL BANK                    DATE   10/01/94
BANK      001                                                               TIME   16:03:46
FROM      10/01/94                                                            MODE   9
TO        10/01/94                LOAN DELINQUENCY REPORT                     PAGE 00001

   LOAN        CUSTOMER              LOAN        DELINQUENT     DELINQUENT     DELINQUENT
  NUMBER        NAME                AMOUNT         30 DAYS        60 DAYS        90 DAYS

 0000010000  MCMULLIGAN, PATRICK  $10000000.00  $      50.00  $      50.00  $        .00
 0000010001  ABBOTT, DAVID        $   11000.00  $     100.00  $     200.00  $        .00
 0000010002  ABBOTT, DAVID        $   12000.00  $     140.00  $        .00  $        .00
 0000010003  ABBOTT, DAVID        $   13000.00  $     150.00  $        .00  $        .00
 0000010005  ROBINS, STEVEN       $     500.00  $      50.00  $        .00  $        .00
 0000010006  PALMER, ARNOLD       $    1000.00  $      75.00  $     150.00  $     225.00
 0000010007  PETERS, PAUL         $     650.00  $      50.00  $        .00  $        .00
 0000010008  ROBERTS, ABRAHAM     $    9000.00  $     120.00  $        .00  $        .00
 0000010009  SMITH, RANDOLPH      $    8000.00  $     115.00  $        .00  $        .00
 0000010010  KLINE, PETER         $    8500.00  $     110.00  $        .00  $        .00
 0000010017  WILLIAMS, ALFRED     $   10000.00  $      50.00  $      50.00  $        .00
 0000010019  JAMES, TIMOTHY       $   11000.00  $     100.00  $     200.00  $        .00
 0000010022  THOMAS, JAMES        $   12000.00  $     140.00  $        .00  $        .00
 0000010026  ROBBINS, KARL        $   13000.00  $     150.00  $        .00  $        .00
 0000010029  MILLER, FREDERICK    $     500.00  $      50.00  $        .00  $        .00
 0000010033  DAVIDSON, ALBERT     $    1000.00  $      75.00  $     150.00  $     225.00
 0000010049  STEVENS, MARY        $     650.00  $      50.00  $        .00  $        .00
 0000010050  MICHAELS, LOUISE     $    9000.00  $     120.00  $        .00  $        .00
 0000010051  ABEL, CHARLIE        $    8000.00  $     115.00  $        .00  $        .00
 0000010056  BAKER, THOMAS        $    8500.00  $     110.00  $        .00  $        .00
 0000010101  TAYLOR, ADRIANNE     $   13000.00  $     150.00  $        .00  $        .00
 0000010111  MILLER, ROBERT       $     500.00  $      50.00  $        .00  $        .00
 0000010123  DAVID, NEIL          $    1000.00  $      75.00  $     150.00  $     225.00
 0000010132  STEVENS, SUSAN       $     650.00  $      50.00  $        .00  $        .00
 0000010133  MITCHELL, PAMELA     $    9000.00  $     120.00  $        .00  $        .00
```

Ready | Display | Page 1 of 4 | 100%

OnDemand Administration

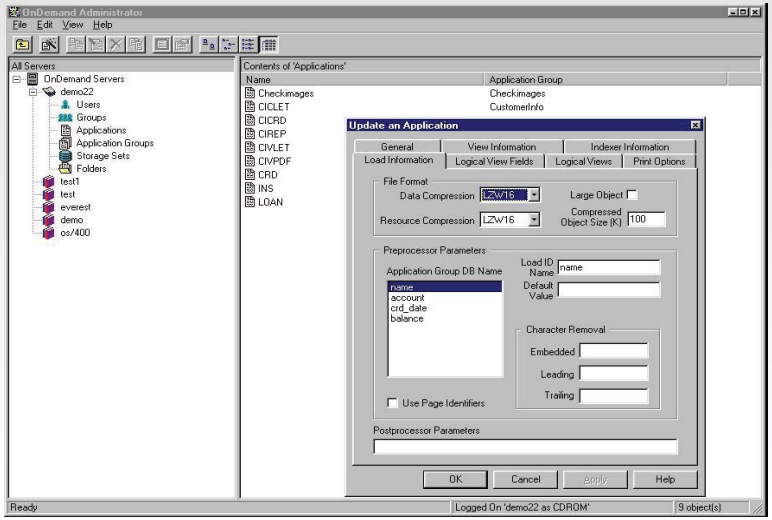

Search/Retrieval User Interface

- **Integrated retrieval/viewing interface**
- **Provides a fill-in-the-blanks query interface**
- **Supports viewing multiple documents seamlessly**
  - AFP, PDF, Line Data, and Image documents
  - Common user interface for multiple document formats
- **Supports AFP resource objects including Overlays and Page Segments**

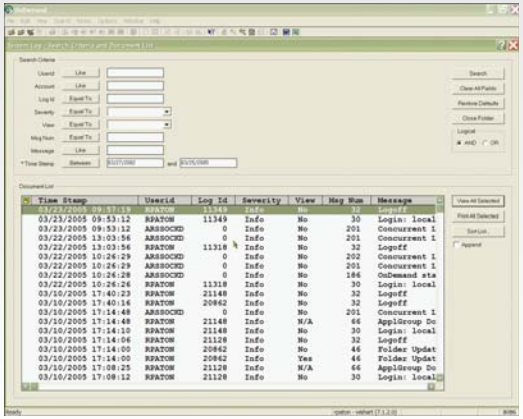

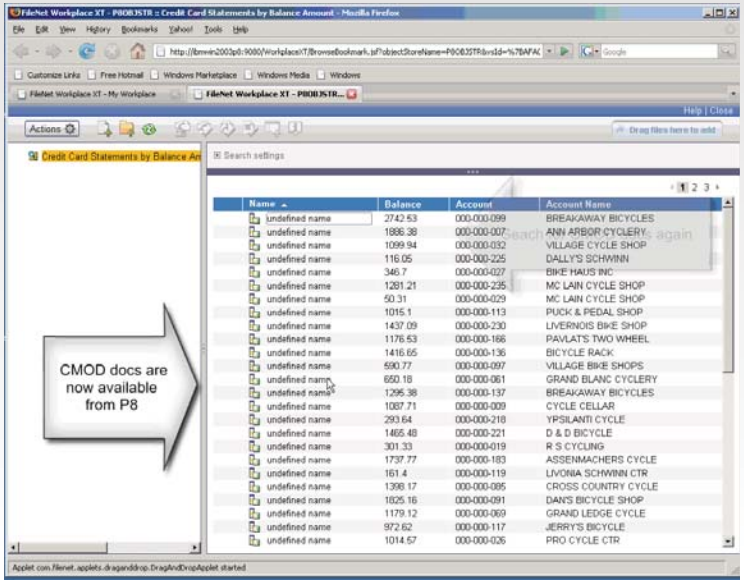

What is Driving Organizations to CM OnDemand?

- **Significant ROI**
  - quickly by transferring costly high volume print output to electronic information capture and presentation
- **Eliminate Print**
  - reduce or totally eliminate printing cost
- **Storage Savings**
  - compression, compression, compression
- **Electronic Statements**
  - customer self service, reduced mail costs, printing
- **Information access**
  - Workers spend 35% of their time looking for information

SQLADRIA SEMINAR
Ljubljana, 15 April 2011

SQL Performance
Good and Bad Access Paths

Sašo Prek
IBM

Sponsored by

Saso Prek
IBM SWG

SQL performance (Access Paths) ?

Smarter software for a smarter planet

Information Management

© 2010 IBM Corporation

Why did the DB2 for z/OS Optimizer choose that access path?

Smarter software for a smarter planet

Information Management

© 2010 IBM Corporation

Agenda

- Why did the optimizer choose that access path?
- Assignment of individual FFs
- Combining FFs
- How can Optim Query Tuner help?

© 2010 IBM Corporation

Why did the optimizer choose that access path?

- This is a cost based optimizer
  - –It chose what it believes to be the lowest access path
- If the access path is not optimal, then the real question is
  - –Why did the optimizer think that was the lowest cost access path?

OR

–What doesn't the optimizer know???

© 2010 IBM Corporation

What types of queries are at risk?

- Do statistics allow distinction between the choices?
- Are there indexes that support an efficient path?
- Is the optimizer able to apply the filtering early?

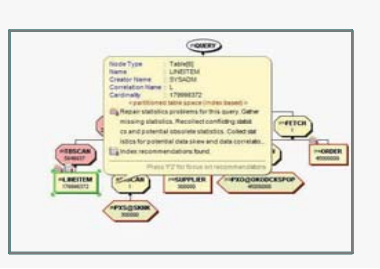

- If no……
  - You're more likely to have performance regressions
    - From minor changes in statistics
    - At release or maintenance upgrade times

© 2010 IBM Corporation

Two main (optimizer) topics for discussion

- Assignment of individual FFs
  - Challenging with host vars or parameter markers
    - Data Skew unknown
    - Range predicate estimates are optimistic
- How optimizer combines multiple FFs
  - Dependent on correlation statistics
    - Or less common – multi-column histograms/frequencies

© 2010 IBM Corporation

Agenda

- Why did the optimizer choose that access path?
- Assignment of individual FFs
- Combining FFs
- How can Optim Query Tuner help?

© 2010 IBM Corporation

Problem Query

- Identified by
  - Pro-active analysis, or
  - Reactive analysis (user complaint?)

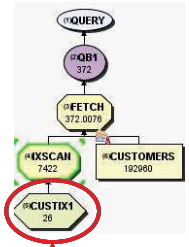


```
SELECT *
FROM CUSTOMERS
WHERE CITY      = ?
  AND STATUS    = ?
  AND BIRTHDATE < ?

Index CUSTIX1*(CITY)
Index CUSTIX2 (BIRTHDATE, STATUS)
```


© 2010 IBM Corporation

Why does this perform poorly?

- Run Counts
  - –At whole query level
  - –At "per object level"
  - –Requires literal values
    - Don't have representative literals?
      - Consider worst case scenario

```
SELECT COUNT(*)
FROM CUSTOMERS
WHERE CITY      = ?
  AND STATUS    = ?
  AND BIRTHDATE < ?

Index CUSTIX1*(CITY)
Index CUSTIX2 (BIRTHDATE, STATUS)
```


© 2010 IBM Corporation

What's the worst case scenario for each index choice?

- Run a count per index
  - –CUSTIX1 (CITY)

```
SELECT CITY, COUNT(*)
FROM CUSTOMERS
GROUP BY CITY
ORDER BY 2 DESC;
```

	CITY		
1_	NEW YORK	147456	76.4%
2_	CHICAGO	21504	11.1%
3_	ATLANTA	7968	4.1%
4_	JACKSONVILLE	7872	4.1%
5_	LAS VEGAS	6144	3.2%

.....only top 5 of 26 shown

© 2010 IBM Corporation

What's the worst case scenario for each index choice?

- Run a count per index
  - –CUSTIX2 (BIRTHDATE, STATUS)
    - Start with individual predicates

```
SELECT COUNT(*) FROM CUSTOMERS
WHERE BIRTHDATE < '9999-12-31'

+------------+
|   192960   | ← 100% of rows qualify for BIRTHDATE
+------------+
SELECT STATUS, COUNT(*)
FROM CUSTOMERS
GROUP BY STATUS
ORDER BY 2 DESC
```

	STATUS		
1_	Y	183313	95%
2_	N	9647	5%

© 2010 IBM Corporation

Counts vs Estimates

- Counts show
  - –CUSTIX1 (CITY) could qualify 76%
    - •Estimate is 3.85%
  - –CUSTIX2 (BIRTHDATE, STATUS) could qualify 95%
    - •Estimate is 5%

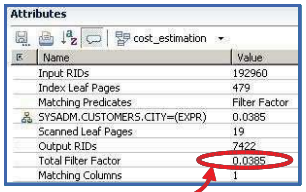

CUSTIX1 = 3.85%

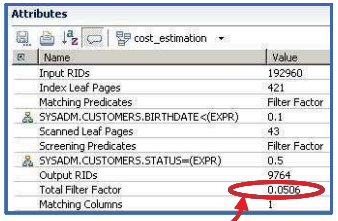

CUSTIX2 = 5%

© 2010 IBM Corporation

What does the optimizer know?

- CUSTIX1 (CITY) estimate
  - CITY has COLCARDF=26
    - Equals predicate (CITY = ?) assumes 1/26 = 0.0385

      That FF estimate used for all calculations

      » Output RIDs = 192960 * 0.0385 ~= 7422

      » Scanned leaf pages = 479 * 0.0385 ~= 19

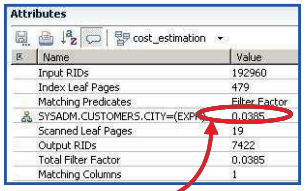

CUSTIX1 = 3.85%

Could be up to 76%

© 2010 IBM Corporation

What does the optimizer know?

- CUSTIX2 (BIRTHDATE, STATUS)
  - BIRTHDATE < ? estimate is 0.1 (10%)
    - That FF estimate used for estimating scanned pages (& rows)

      Scanned leaf pages = 421 * 0.1 ~= 43
  - STATUS = ?
    - COLCARDF = 2, equals predicate assume ½ = 0.5

    43 = 10% of 421 = 42,1 = 43

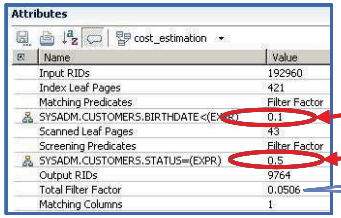

BIRTHDATE < ?

STATUS = ?

Estimate is 5%, could be up to 95%

© 2010 IBM Corporation

How did optimizer estimate filtering on BIRTHDATE < ?

Default filter factors for interpolation (from DB2 V9 Perf.Mon. and Tun.Guide,SC18-9851-09, str. 279)

COLCARDF	Factor for Op	Factor for LIKE/BETWEEN
>=100,000,000	1/10,000	3/100,000
>=10,000,000	1/3,000	1/10,000
>=1,000,000	1/1,000	3/10,000
>=100,000	1/300	1/1,000
>=10,000	1/100	3/1,000
>=1,000	1/30	1/100
>=100	1/10	3/100
>=2	1/3	1/10
=1	1	1
<=0	1/3	1/10

Note: Op is one of these operators: <, <=, >, >=.

1/10 = 0.1
BIRTHDATE FF

COMMENT: This is DB2's documented guess for an impossible to estimate Filter factor. Used for host variables/parameter markers.

© 2010 IBM Corporation

Is that estimate accurate?

- Range predicate with parameter marker/host variable
  - Use default interpolation filter factor chart
    - COLCARDF 456 --> FF = 1/10
  - In reality, could qualify anywhere from all to no rows

BIRTHDATE < ?

- How many people here born before host variable?

  What if value is '1920-01-01'?

  What if value is '1990-01-01'?

  Cannot accurately estimate without literal value

© 2010 IBM Corporation

How does Optim Query Tuner help?

- Statistics advisor concluded that……
  - REOPT bind parameter required

    Query contains host variables or parameter markers

    And thus, literals unknown

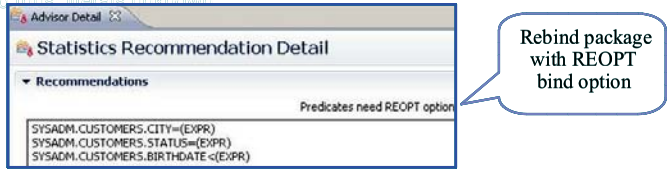

  - And after providing literals…..

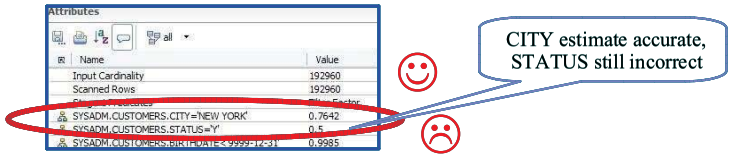

© 2010 IBM Corporation

Did Query Tuner recommend anything else?

- Statistics advisor also recommended FREQVAL

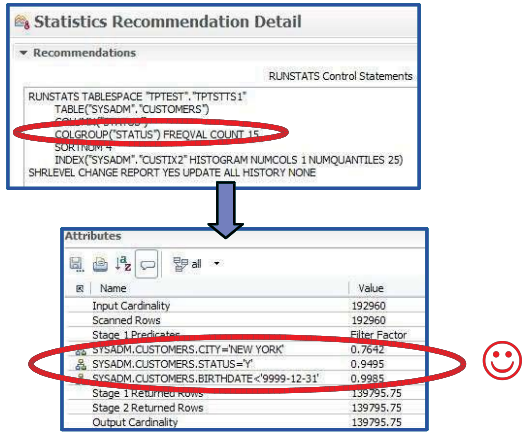

© 2010 IBM Corporation

Regardless of literals – can I simplify the competition?

- Current access path depends on literals
  - Instead, provide 1 clear plan!!!
    - Create index on CITY, STATUS & BIRTHDATE
  - If not possible…..
    - REOPT could be employed

      Ensure FREQVAL stats are collected
    - Use literal values if always searching for same value (STATUS='N')

```
WHERE CITY      = ?
  AND STATUS    = ?
  AND BIRTHDATE < ?

Index CUSTIX1 (CITY)
Index CUSTIX2 (BIRTHDATE, STATUS)
```

*NOTE: Indexes that lead with range predicate columns will always pose risk for optimizer

© 2010 IBM Corporation

Agenda

- Why did the optimizer choose that access path?
- Assignment of individual FFs
- Combining FFs
- How can Query Workload Tuner help?

© 2010 IBM Corporation

How does optimizer combine FFs?

- We know each zipcode isn't in each city
  - But does optimizer know this???

```
SELECT * FROM CUSTOMERS
WHERE CITY = ?      ← Matching
  AND ZIPCODE = ? ← Screening
```

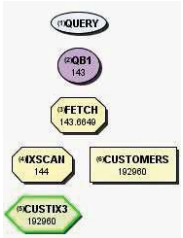


```
Index CUSTIX3 (CITY, CUSTNO, ZIPCODE) ← Winner
Index CUSTIX4 (ZIPCODE)
```

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	1854
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.CITY=(EXPR)	0.0385
Scanned Leaf Pages	72
Screening Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Output RIDs	144
Total Filter Factor	0.0007
Matching Columns	1

1 / 26 = 0.0385

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	461
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Scanned Leaf Pages	9
Output RIDs	3711
Total Filter Factor	0.0192
Matching Columns	1

1 / 52 = 0.0192

© 2010 IBM Corporation

Combining FFs

- Optimizer chooses larger index
  - Due to much better filtering
    - 144 rows qualified vs 3711 for ZIPCODE alone
      Is this filtering accurate?
    - Despite 72 vs 9 leaf pages to be read

```
Index CUSTIX3 (CITY, CUSTNO, ZIPCODE) ← Winner
Index CUSTIX4 (ZIPCODE)
```

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	1854
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.CITY=(EXPR)	0.0385
Scanned Leaf Pages	72
Screening Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Output RIDs	144
Total Filter Factor	0.0007
Matching Columns	1

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	461
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Scanned Leaf Pages	9
Output RIDs	3711
Total Filter Factor	0.0192
Matching Columns	1

© 2010 IBM Corporation

How did optimizer estimate qualified rows?

- # qualified rows from index
  - Input rows * matching filtering * screening filtering

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	1854
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.CITY=(EXPR)	0.0385
Scanned Leaf Pages	72
Screening Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Output RIDs	144
Total Filter Factor	0.0007
Matching Columns	1

192960 * 0.0385 * 0.0192 = 142.64 (~ 144 ☺)

© 2010 IBM Corporation

How did optimizer estimate qualified rows?

- # qualified rows from index
  - Input rows * matching filtering * screening filtering

Attributes

Name	Value
Input RIDs	192960
Index Leaf Pages	461
Matching Predicates	Filter Factor
SYSADM.CUSTOMERS.ZIPCODE=(EXPR)	0.0192
Scanned Leaf Pages	9
Output RIDs	3711
Total Filter Factor	0.0192
Matching Columns	1

192960 * 0.0192 = 3705. (~ 3711 ☹

© 2010 IBM Corporation

Are CITY & ZIPCODE correlated or independent?

- Run individual column counts

```
SELECT COUNT(DISTINCT CITY)     = 26 CITIES
      ,COUNT(DISTINCT ZIPCODE) = 52 ZIPCODES
```

- Run count of number of distinct groups

```
SELECT COUNT(*) FROM
(SELECT DISTINCT CITY, ZIPCODE FROM T1) AS X
                        = 52 ← CITY/ZIP Combinations
```

- Calculate product of columns = 26 * 52 = 1,352
  - If combined count < product, cols are correlated
    - 52 < 1352, thus columns are correlated

© 2010 IBM Corporation

And how correlated or how independent?

- Given CITY & ZIPCODE
  - Highest individual cardinality is 52 (ZIPCODE)
  - Product of 2 columns = 26 * 52 = 1352
  - Thus…..
    - Column group (KEYCARD) for CITY/ZIPCODE must fall between

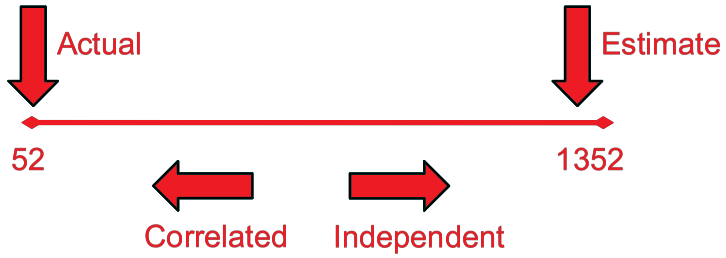

© 2010 IBM Corporation

What about KEYCARD?

RUNSTATS sets the following columns to LARGE:
- CARD in SYSTABLES
- CARD in SYSINDEXPART
- FAROFFPOS in SYSINDEXPART
- NEAROFFPOS in SYSINDEXPART
- FIRSTKEYCARD in SYSINDEXES
- FULLKEYCARD in SYSINDEXES

- By default, RUNSTATS only collects
  - FIRSTKEYCARDF → (CITY,CUSTNO,ZIPCODE)
  - FULLKEYCARDF → (CITY,CUSTNO,ZIPCODE)

- KEYCARD option collects
  - 2NDKEYCARDF (for 3+ col index)
  - 3RDKEYCARDF (for 4+ col index) → (CITY,CUSTNO,ZIPCODE)
  - etc.

```
RUNSTATS TABLESPACE TESTDB.TESTTS
         INDEX(CUSTIX3) KEYCARD
```

- But we need CITY, ZIPCODE without CUSTNO

© 2010 IBM Corporation

So it's the fault of the index?

- Because there is no index leading with CITY, ZIPCODE
  - Correlation information (KEYCARD) is not collected by RUNSTATS

- Index CUSTIX3 (CITY, CUSTNO, ZIPCODE)
  - Firstkeycardf = 26
  - Fullkeycardf = 192960
  - Colcardf of CUSTNO = 192960

  - High cardinality CUSTNO, destroys the correlation of CITY, ZIPCODE
    - Thus - nothing to stop CITY & ZIPCODE being assumed independent

© 2010 IBM Corporation

How do I collect correlation information?

- RUNSTATS COLGROUP can collect correlation for
  - non-consecutive indexed columns,
  - or non-indexed columns

```
RUNSTATS TABLESPACE TESTDB.TESTTS
         TABLE(SYSADM.CUSTOMERS)
         COLGROUP(CITY, ZIPCODE)
```

  - Or, create an index with CITY, ZIPCODE (either order)
    - And specify KEYCARD on RUNSTATS

      If > 2 columns in index

```
DSNU616I  ) 098 13:36:15.04 DSNUSUCD - SYSCOLDIST CATALOG STATISTICS FOR
                                       CITY,ZIPCODE
                                 CARDINALITY =  5.2E+01
```


© 2010 IBM Corporation

What happens once correlation is known?

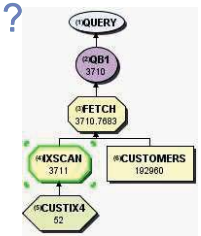

- With correlation statistics
  - Both indexes now estimated to qualify 3711 rows

```
Index CUSTIX3 (CITY, CUSTNO, ZIPCODE)
Index CUSTIX4 (ZIPCODE)                ← Winner
```

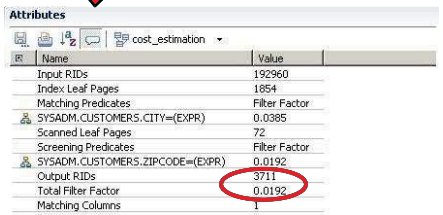

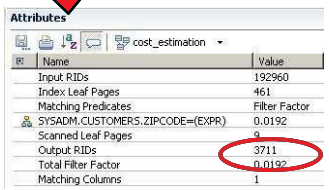

© 2010 IBM Corporation

How does Optim Query Tuner help?

- Statistics Advisor Recommendations
  - COLGROUP suggested

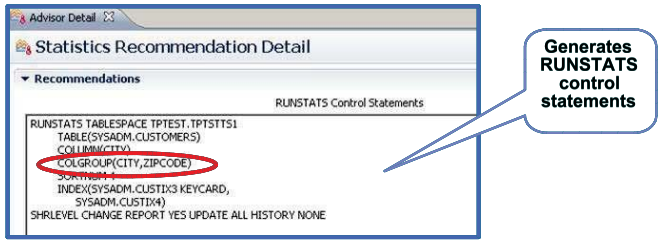

© 2010 IBM Corporation

Why did the optimizer??? - Conclusions

- Common reasons for optimizer making the wrong choice
  - Mis-estimating individual FFs
    - Unknown data skew
    - Overly optimistic estimates for range predicates
    - More likely to occur….

      Literals unknown (host variables/parameter markers)

      Statistics missing
  - Incorrectly combining individual FFs
    - Unknown correlation
    - More likely to occur….

      when filtering predicates are not supported by (matching) indexes

- *Same problems exist for single or multi-table queries

© 2010 IBM Corporation

Why did the optimizer??? - Conclusions

- If there exist access path choices that can fall into traps
  - Risk of regression after RUNSTATS/migration etc
- Provide clear choices
  - Statistically
    - Allow statistics to accurately differentiate choices
    - Use Statistics Advisor from OQT/OQWT
  - OR, Indexing to provide clear choice
    - Use Index Advisor from OQT/OQWT

© 2010 IBM Corporation

Agenda

- Why did the optimizer choose that access path?
- Assignment of individual FFs
- Combining FFs
- How can Optim Query Tuner help?

© 2010 IBM Corporation

© 2010 IBM Corporation

IBM Optim Query Tuner Offerings

Optimize performance and reduce costs

- Optimize performance
  - Improve customer satisfaction
  - Improve organizational productivity
  - Satisfy service level objectives
  - Fit workloads into maintenance windows
- Reduce costs
  - Increase overall system capacity
  - Reduce chargebacks
  - Simplify analysis for DBAs, developers, designers, and others
  - Improve statistics collection efficiency
  - Reduce time to respond

© 2010 IBM Corporation

Single Query and Workload Offerings

Development Environment

Optim Query Tuner

Tune SQL pre-production while costs and impact are low

Production Environment

Optim Query Workload Tuner

Optimize workload for peak performance

© 2010 IBM Corporation

Workload Tuning vs. Query Tuning

- Optimization decisions are based on tradeoffs
  - Statistics – CPU costs vs. query savings
  - Indexing – query speed vs resource and transaction
- Workload tuning speeds up analysis
  - Analyzes multiple queries at once
- Workload tuning consolidates and optimizes recommendation for overall workload
  - Statistics recommendations
  - Index recommendations

Speed up analysis, optimize design,

and balance resource usage

© 2010 IBM Corporation

Expert-Enabled Query Tuning Out of the Box

- Identify query candidates
  - DB2 catalog
  - Dynamic statement cache
  - Development Studio hot spots
  - Query or performance monitors
- Facilitate analysis
  - Query formatting
  - Query annotation
  - Access path visualization and annotation
- Get expert tuning advice
  - Improve query design
  - Improve statistics quality
  - Improve database design

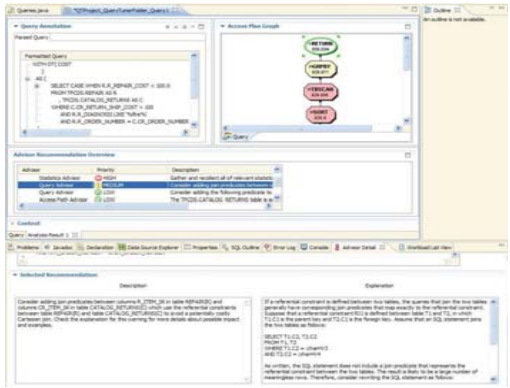

© 2010 IBM Corporation

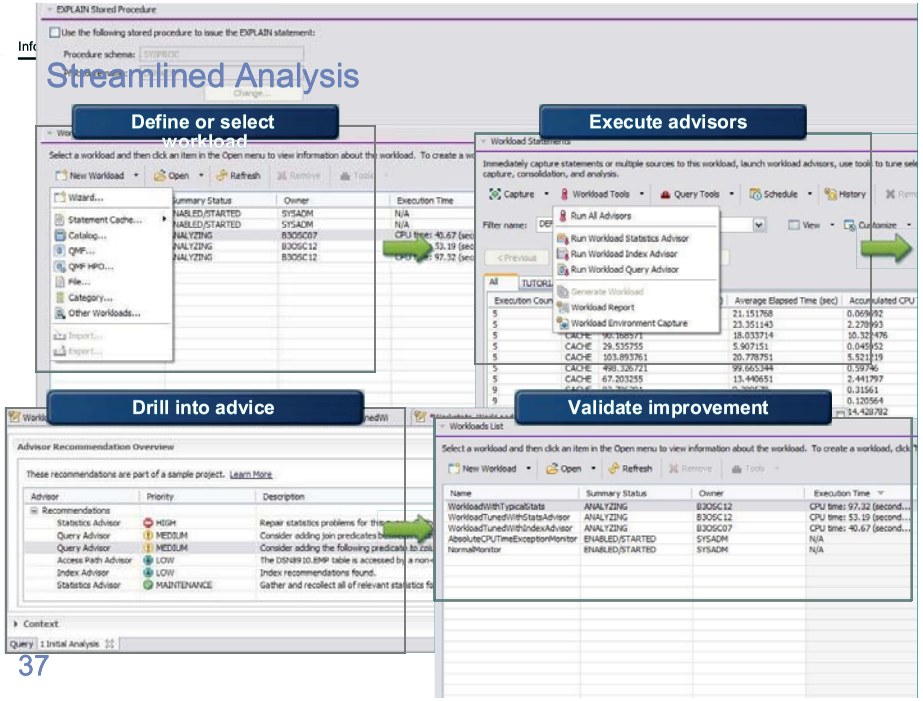

Indexing Advice to Improve Query Efficiency

- Improve query efficiency
  - Indexing foreign keys in queries that do not have indexes defined
  - Identifying index filtering and screening
  - Support for index only access
  - Indexing to avoid sorts
- 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
  - Utilize virtual index capabilities built into the DB2 engine

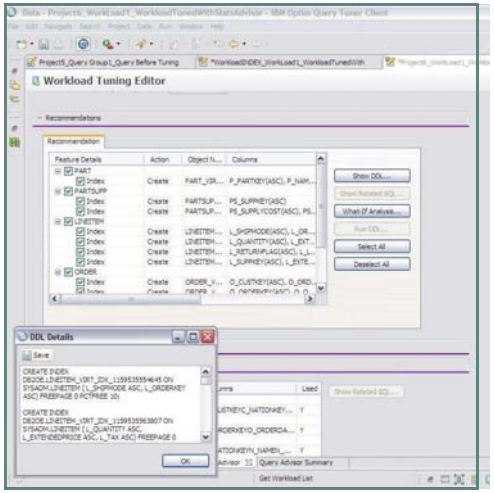

© 2010 IBM Corporation

Improve Query Design

Guard against errors and oversights:
Further constrain query, increase index utilization, and reduce data reads

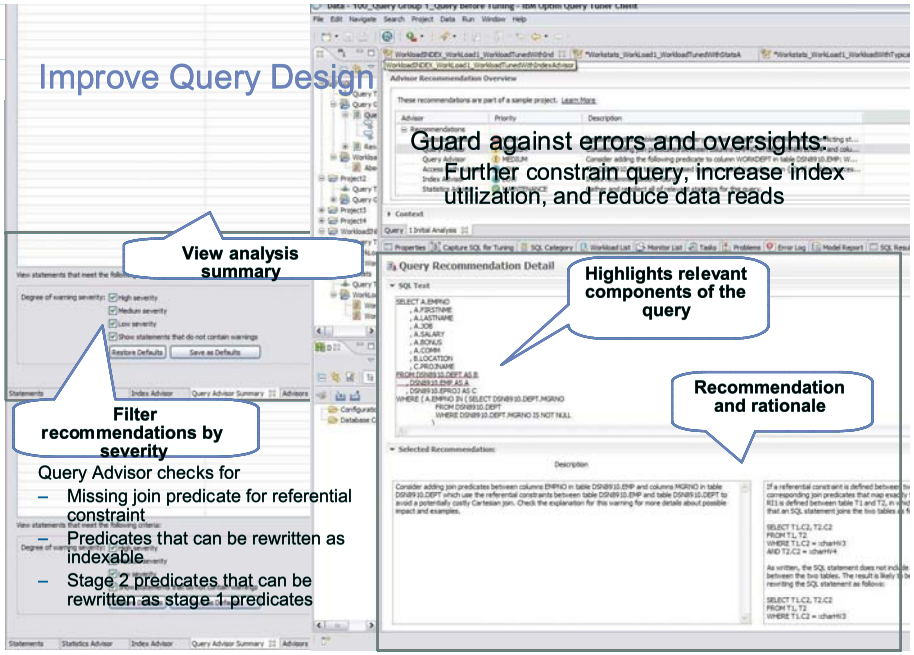

Query Advisor checks for
- Missing join predicate for referential constraint
- Predicates that can be rewritten as indexable
- Stage 2 predicates that can be rewritten as stage 1 predicates

© 2010 IBM Corporation

What types of queries are at risk?

- Do statistics allow distinction between the choices?
  Solution — Statistics Advisor
- Are there indexes that support an efficient path?
  Solution — Index Advisor
- Is the optimizer able to apply the filtering early?
  Solution — Query Advisor

© 2010 IBM Corporation

Tuning SQL with Optim Query Tuner, Part 1: Understanding access paths

New! Article

– *Learn how to monitor and tune queries and workloads to improve application performance*

- If you are a developer, DBA, or query tuning specialist, it is critical that you understand the basics of access paths so that you can precisely tune queries and query workloads before they cause problems in your production environment. This basic understanding, coupled with the visualization and tuning advice provided by IBM® Optim® query tuning solutions, can help make you more efficient at this task. This article provides conceptual background on access paths, shows you how to read an access path graph, and walks through the access path graph to demonstrate critical information regarding access path selection. The article concludes with a sample scenario that demonstrates how to use query annotation, a feature in Optim Query Tuner that helps you with query analysis by providing vital statistical information directly in the SQL statement.
- http://www.ibm.com/developerworks/data/library/techarticle/dm1006optimquerytuner1/index.html

© 2010 IBM Corporation

Resources

- Integrated Data Management Community
  - http://www.ibm.com/developerworks/spaces/optim
- IBM Optim Query Workload Tuner web page
  - https://www-01.ibm.com/software/data/db2imstools/db2tools/opti-expert-zos/
- SQL Tuning, Not just for hardcore DBAs anymore
  - http://www.ibmdmmagazinedigital.com/dmmagazine/
- Redbooks at www.redbooks.ibm.com
  - DB2 9 for z/OS Performance Topics SG24-7473
  - IBM DB2 9 for z/OS: New Tools for Query Optimization SG24-7421

© 2010 IBM Corporation

Contents of Eclipse-based Query Tuning offerings

	Data Studio	Optim Query Tuner for z/OS	Optim Query Workload Tuner for z/OS
Queries from all sources	✓	✓	✓
Reports	✓	✓	✓
Query Formatter	✓	✓	✓
Access Plan Graph	✓	✓	✓
Query Statistics Advisor	✓	✓	✓
Query Annotation		✓	✓
Visual Plan Hint		✓	✓
Query Index Advisor		✓	✓
Query Advisor		✓	✓
Access Path Advisor		✓	✓
Workload Statistics Advisor			✓
Workload Index Advisor			✓
Workload Query Advisor			✓

© 2010 IBM Corporation

Why did the optimizer choose that access path?

Saghi Amirsoleymani
amirsole@us.ibm.com
Terry Purcell
tpurcel@us.ibm.com

© 2010 IBM Corporation

SQLADRIA SEMINAR
Ljubljana, 15 April 2011

Workload Tuning

Sašo Prek
IBM

Sponsored by

Saso Prek
IBM SWG

Workload Management and DB2

Content

- Governor – RLF
- Workload Overview
- Important WLM Concepts And Terminology
- How DB2 Exploits Workload Manager
- Understanding WLM and how it impacts DB2
- WLM/DB2 Usage Recommendations
- Summary

© 2011 IBM Corporation

Workload Management

Workload Management (DB2)

=

Priotirizing Resources

=

z/OS Workload Managemnt

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility

The Resource Limit Facility (RLF) is a DB2 function, defined at DB2 install time, which provides the ability to prevent SQL from starting,

which we call ***preemptive governing***,

or

it can stop activity after a certain amount of resources are consumed, which we call

reactive governing.

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility

Predictive governing

Set warning and error thresholds by which the governor can inform users (via your application programs) that a certain processing limit might be exceeded for a particulary dynamic SELECT, INSERT, UPDATE, or DELETE statement.

Reactive governing

Stop a currently executing dynamic SQL statement (SELECT, INSERT, UPDATE, or DELETE) that exceeds the processor limit.

- Restrict bind and rebind activities to avoid performance impacts on production data.
- Restrict particular parallelism modes for dynamic queries.

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility

START RLIMIT ID=...

```
DSNT704I  -DSN1 SYSIBM.DSNRLST01 HAS BEEN STARTED FOR THE RESOURCE
LIMIT FACILITY
DSNT704I  -DSN1 SYSIBM.DSNRLMT01 HAS BEEN STARTED FOR THE RESOURCE
LIMIT FACILITY
DSN9022I  -DSN1 DSNTCSTR 'START RLIMIT' NORMAL COMPLETION
***
```

Starts the governor and identifies a resource limit specification table. You can also use START RLIMIT to switch resource limit specification tables.

DISPLAY RLIMIT

```
DSNT700I  -DSN1 SYSIBM.DSNRLST01 IS THE ACTIVE RESOURCE LIMIT
SPECIFICATION TABLE
DSNT700I  -DSN1 SYSIBM.DSNRLMT01 IS THE ACTIVE RESOURCE LIMIT
SPECIFICATION TABLE
DSN9022I  -DSN1 DSNTCDIS 'DISPLAY RLIMIT' NORMAL COMPLETION
***
```

Displays the current status of the governor. If the governor has been started, the output from the command also identifies the resource limit specification table.

STOP RLIMIT

Stops the governor and removes any set limits.

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility

Table 29.1. The Columns of the RLST

Name	Definition
AUTHID	Identifies the primary authorization ID of the user to whom the limit set by this row applies. If blank, this row applies to all primary authorization IDs at the location specified by the LUNAME column.
PLANNAME	Specifies the plan name for which the limit set by this row applies. If blank, this row applies to all plan names at the location specified by the LUNAME column. PLANNAME is valid only when RLFFUNC is blank. If RLFFUNC contains a value, the column must be blank or the entire row is ignored.
ASUTIME	Specifies the maximum number of CPU service units permitted for any single dynamic SQL statement. If NULL, this row does not apply a limit. If less than or equal to 0, this row indicates that dynamic SQL is not permitted.
LUNAME	The logical unit name of the site where the request originated. If blank, this row applies to the local site. If PUBLIC, this row applies to all sites.
RLFFUNC	Indicates the type of resource this row is limiting: blank = row governs dynamic SQL reactively by plan name 1 = row governs BIND for plans or packages in collections 2 = row governs dynamic SQL reactively by collection and package names 3 = row disables query I/O parallelism 4 = row disables query CP parallelism 5 = row disables Sysplex query parallelism 6 = row governs dynamic SQL predictively by plan name 7 = row governs dynamic SQL predictively by collection and package names If any other values are in this column, the row is ignored.
RLFBIND	Indicates whether the BIND command is permitted. The value N indicates that BIND is not allowed; any other value means that the BIND command is allowed. Valid only when RLFFUNC equals 1.

GOVERNOR - RLF – Tuning Strategy

You can create multiple RLSTs, with each controlling resources in a different manner. Some reasons for doing this are as follows:

- To control the same resources in different RLSTs with different limits.
- To control different resources in different RLSTs.
- To eliminate resource control for a plan or package from a certain RLST, thereby removing the limit.
- To control one type of limiting separately from another type; for example, to control binds in one RLST, plans and packages in another, and users in another. However, this is impractical because only one RLST can be active at any given time.

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility - DSNRLMTxx

```
CREATE TABLE authid.DSNRLMTxx      (RLFFUNC    CHAR(1)
       RLFEUAN    CHAR(32)    NOT NULL WITH DEFAULT,
       RLFEUID    CHAR(16)    NOT NULL WITH DEFAULT,
       FLFEUWN    CHAR(18)    NOT NULL WITH DEFAULT,
       RLFIP      CHAR(254)   NOT NULL WITH DEFAULT)
       ASUTIME    INTEGER,
       RLFASUERR  INTEGER,
       RLFASUWARN INTEGER,
       RLF_CATEGORY_B CHAR(1) NOT NULL WITH DEFAULT)
         IN DSNRLST.DSNRLSxx;
```

Column name	Description
RLFFUNC	Specifies how the row is used. The values that have an effect are: '8' The row reactively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by client information (RLEUID, RLFEUAN, RLFEUWN, and RLFIP). '9' The row predictively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by client information (RLEUID, RLFEUAN, RLFEUWN, and RLFIP). All other values are ignored.
RLFEUAN	Specifies an application name. A blank value in this column means that the row applies to all application names from the location specified in RLFIP.
RLFEUID	Specifies an end user's user ID. A blank value means that the limit specifications in this row apply to every user ID for the location that is specified in RLFIP.

DSNRLMTxx Resource limit table can be used to limit the amount of resources used be dynamic queries that run on middleware servers

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility - DSNRLSTxx

```
CREATE TABLE authid.DSNRLSTxx
      (AUTHID   VARCHAR(128)     NOT NULL WITH DEFAULT,
       PLANNAME CHAR(8)       NOT NULL WITH DEFAULT,
       ASUTIME  INTEGER,
       -------3-column format --------
       LUNAME   CHAR(8)       NOT NULL WITH DEFAULT,
       -------4-column format ---
       RLFFUNC  CHAR(1)
       RLFBIND  CHAR(1)
       RLFCOLLN VARCHAR(128)     NOT NULL WITH DEFAULT,
       RLFPKG   VARCHAR(128)  NOT NULL WITH DEFAULT),
       -------8-column format --------
       RLFASUERR  INTEGER,
       RLFASUWARN INTEGER,
       RLF_CATEGORY_B CHAR(1) NOT NULL WITH DEFAULT)
       -------11-column format --------
        IN DSNRLST.DSNRLSxx;
```

Column name	Description
RLFFUNC	blank The row reactively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by plan name. '1' The row reactively governs bind operations. '2' The row reactively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by package or collection name. '3' The row disables query I/O parallelism. '4' The row disables query CP parallelism. '5' The row disables Sysplex query parallelism. '6' The row predictively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by plan name. '7' The row predictively governs dynamic SELECT, INSERT, UPDATE, MERGE, TRUNCATE, or DELETE statements by package or collection name. All other values are ignored.

DSNRLSTxx Resource limit tables can be used to limit the amount of resource used by queries

© 2011 IBM Corporation

GOVERNOR - RLF – Resource Limit Facility

```
Table: SYSIBM.DSNRLST01

AUTHID :
PLANNAME : DISTSERV
ASUTIME : 1547388
LUNAME : PUBLIC
RLFFUNC : 2
RLFBIND :
RLFCOLLN :
RLFPKG : SQLLF000
RLFASUERR : ?
RLFASUWARN : ?
RLF_CATEGORY_B :
```


© 2011 IBM Corporation

GOVERNOR - RLF – Calculating ASUTIME

For example, if processor A is rated at 900 service units per second and you do not want any single dynamic SQL statement to use more than 10 seconds of processor time, you could set ASUTIME as follows:

ASUTIME time = 10 seconds 900 service units/second = 9000 SU

CPU Model Number	No. of Engines	Service Units per SRM Sec	Estimated MIPS rate
z/900 2Cx models:			
2064-2C1	1	14692.38	303.0
2064-2C2	2	13961.61	575.0
2064-2C3	3	13377.93	830.0
2064-2C4	4	13082.58	1073.0
2064-2C5	5	12638.23	1305.0
2064-2C6	6	12345.68	1576.0
2064-2C7	7	12048.19	1736.0
2064-2C8	8	11756.06	1940.0
2064-2C9	9	11461.32	2129.0
2064-210	10	11165.39	2308.0
2064-211	11	10869.56	2478.0
2064-212	12	10723.86	2638.0
2064-213	13	10430.25	2790.0
2064-214	14	10139.42	2933.0
2064-215	15	9939.75	3067.0
2064-216	16	9696.97	3192.0

© 2011 IBM Corporation

Why Workload Manager (WLM)?

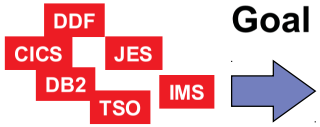

Goal Prioritization

What does z/OS need to accomplish objectives?
How important is the workload to the business?

- **Complexity of systems has increased dramatically**
  - In the past priority managed by PARMLIB options
  - Too much work to 'micro-manage' each z/OS system
  - Multiple LPARs, Data sharing, large sophisticated workloads
- **Workload Manager improves the performance management process**
  - Prioritize workload based upon goals and business objectives
  - Let the system optimize and prioritize resource management
  - Make sure that the most important workload gets the appropriate resource

© 2011 IBM Corporation

DB2 Integration in z/OS

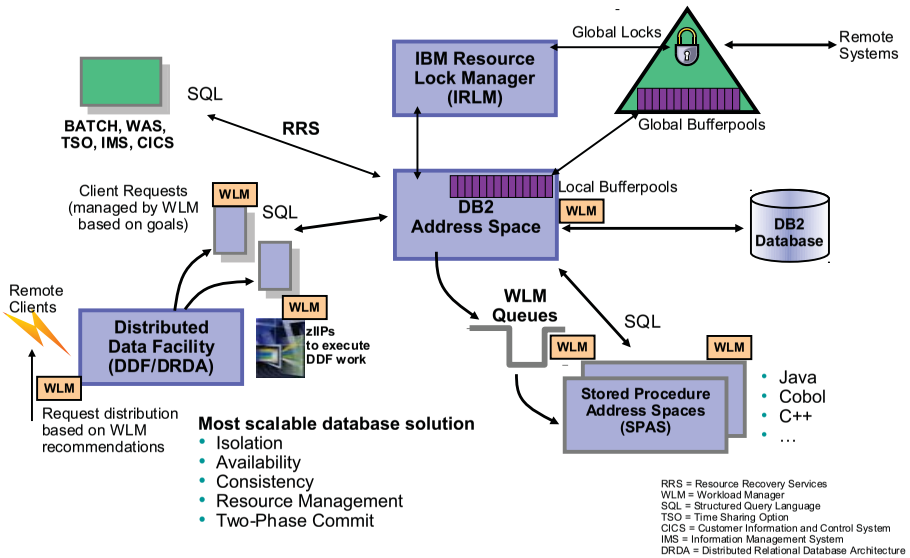

© 2011 IBM Corporation

DB2 And Workload Manager

- **Workload Manager (WLM) is the priority and resource manager for z/OS and implicitly for DB2 as well**
- **Workload Manager manages CPU, I/O, and memory resources as needed by the workload**
- **WLM manages critical DB2 address spaces**
  - DB2 subsystem address space(s)
  - DB2 Stored Procedure address spaces
- **These resources have an impact on how DB2 applications perform**
- **Workload Manager is used to manage DB2 workflow**
  - DB2 Distributed and DB2 Stored Procedure workload
  - Priority and performance of allied address spaces (CICS, IMS, batch, TSO) that call DB2

© 2011 IBM Corporation

Workload Manager
Terminology

- **A SERVICE DEFINITION consists of one or more SERVICE POLICIES**
- **A service policy contains several WORKLOADS**
- **Each workload consists of one or more SERVICE CLASSES**
- **Each service class has at least one PERIOD and each period has one GOAL**
- **There are five types of goals**
  - System, Average Response Time, % Response Time, Execution Velocity, Discretionary
- **Address spaces and transactions are assigned to service classes by CLASSIFICATION RULES**

SERVICE DEFINITION
SERVICE POLICIES
WORKLOADS
SERVICE CLASSES
PERIOD
GOAL
CLASSIFICATION RULES

© 2011 IBM Corporation

Key WLM Concepts

- **WLM is built upon two key concepts**
  - ***DEFINITION*** - WLM provides mechanisms to categorize, prioritize, and manage workload
    - These are the service definitions managed by the WLM dialogs
  - ***FEEDBACK*** - The components managed by WLM provide information (samples) to help WLM determine how well it is doing
    - These are feedback mechanisms provided by each component, subsystem, and operating system

© 2011 IBM Corporation

WLM
Service Classes Categorize Workload

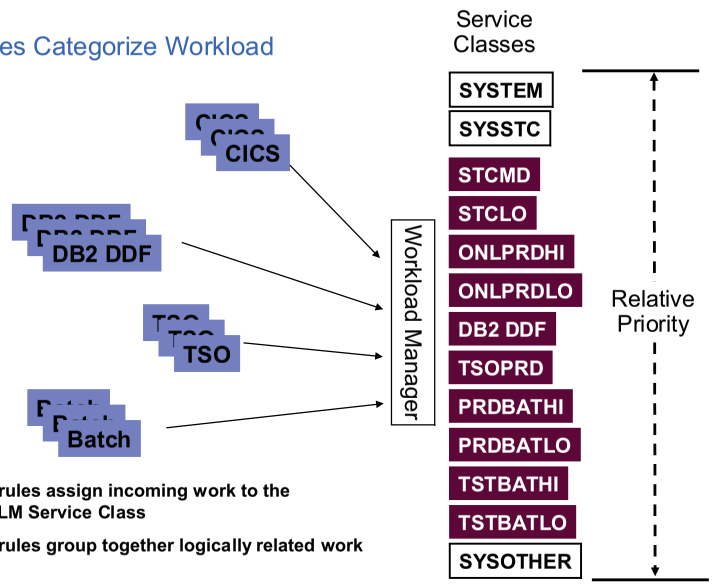

- **Classification rules assign incoming work to the appropriate WLM Service Class**
- **Classification rules group together logically related work**

© 2011 IBM Corporation

Workload Manager
Service Classes And Goals

- **z/OS resources assigned based upon goals defined in WLM**

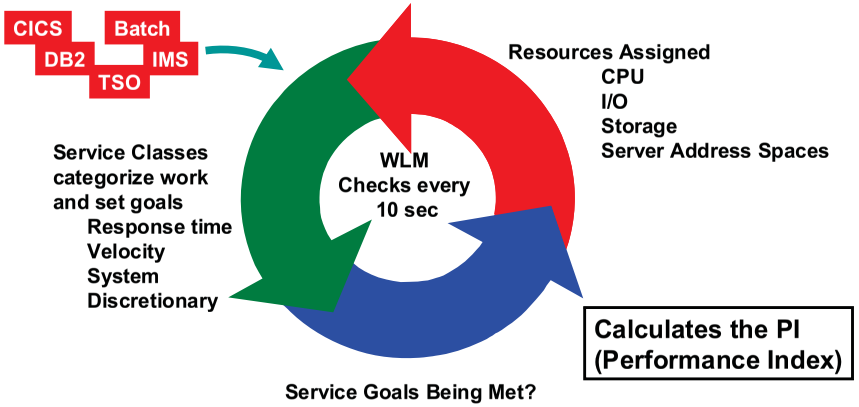

© 2011 IBM Corporation

Understanding WLM Goals
The Performance Index

- **Service Class periods are compared by calculating a Performance Index (PI) for each**
- **PI gives WLM a common way to track how well the work is doing regardless of goal type**
- **Importance parameter**
  - Defined as part of the Service Class - 1 (high) to 5 (low)
  - Assigned to a Service Class Period
  - A way to prioritize critical goals
  - For work at the same importance level, WLM attempts to equalize the PIs

• **The PI equals 1**	=>	**The work in the period is meeting its goal exactly**
• **The PI is less than 1**	=>	**The work is doing better than its goal**
• **The PI is more than 1**	=>	**The work is missing its goal**

© 2011 IBM Corporation

The Importance Of Importance

- **Importance parameter**
  - A way to prioritize critical goals
- **Not analogous to Dispatching Priority**
  - What is the importance of achieving the goal?
- **WLM attempts to meet importance 1 goals first, and so on…….**
- **Helps WLM determine donors and receivers of resources**
  - Donors – workload that can give up resource
  - Receivers – workload that needs resource

© 2011 IBM Corporation

Types Of WLM Goals
Velocity Goals

- **The percentage of time workload is ready and able to run, and is not delayed for lack of resources**
  - Example - Velocity of 50 means that 50% of the time resources should be available for work to run
- **Velocity goals measure of acceptable delay based on samples**
  - High velocity goals (example 90) in general are unreasonable
  - Use relatively higher velocity goals for DBM1 and SSAS
  - At first glance easy to set
- **Velocity goals actually requires more analysis**
  - Systems, environments, and workloads will change over time
  - Velocity goals require regular review
  - As systems change, velocity goals may need 'fine-tuning'

© 2011 IBM Corporation

Types Of WLM Goals
Response Time Goals

- **Average response time**
  - Average response time for a given set of transactions
  - Include queue time and execution time
- **Percentile response time**
  - Percentile of transactions that need to complete within a desired response time
  - Reduces the impact of 'outliers'
- **Rule of thumb**
  - Work should have at least 10 completions in a 20 minute time frame to have adequate samples
- **Consider Response time goals where possible for DB2 workloads**
  - DDF requests, even batch jobs

© 2011 IBM Corporation

Special Service Classes
SYSTEM, SYSSTC, & SYSOTHER

- **SYSTEM**
  - For selected high priority system address spaces
  - Get highest CPU and I/O dispatching priority in system
- **SYSSTC**
  - For selected high priority started tasks and workload
  - Second highest priority behind SYSTEM
  - Place very high importance workload items here
    - DB2 Example - place IRLM here
- **SYSOTHER**
  - Unclassified work falls here
  - Bottom of the resource food chain
- **IRLM** – very high – **SYSSTC**
- **VTAM** – very high – **SYSSTC**
- **CICS - ..... – ONLxxxx**
- **DDF - ……. –** in the middle

© 2011 IBM Corporation

Specifying The WLM Objectives

```
File  Utilities  Notes  Options  Help
-------------------------------------------------------------------------------
Functionality LEVEL011          Definition Menu          WLM Appl LEVEL013
Command ===> _____________________________________________________________

Definition data set  . . : none

Definition name  . . . . . DEMO      (Required)
Description  . . . . . . . No overrides

Select one of the
following options. . . . . ___  1.  Policies
                                2.  Workloads
                                3.  Resource Groups
                                4.  Service Classes
                                5.  Classification Groups
                                6.  Classification Rules
                                7.  Report Classes
                                8.  Service Coefficients/Options
                                9.  Application Environments
                                10. Scheduling Environments
```

All the various WLM constructs on z/OS are defined using the WLM ISPF dialogs

- **WLM provides an ISPF interface to define and manage the WLM Service Definition**
- **Note – z/OS Management Facility V1.12 provides a new management interface**

© 2011 IBM Corporation

Workload Manager As A DB2 Priority Mechanism

```
 Subsystem-Type  Xref  Notes  Options  Help
-------------------------------------------------------------------------------
              Create Rules for the Subsystem Type        Row 1 to 5 of 5

 Subsystem Type . . . . . . . . DDF    (Required)
 Description  . . . . . . . . . Distributed DB2
 Fold qualifier names?  . . . . Y  (Y or N)

 Enter one or more action codes: A=After  B=Before  C=Copy  D=Delete
 M=Move I=Insert rule IS=Insert Sub-rule  R=Repeat

          -------Qualifier-------------          -------Class--------
 Action    Type       Name     Start               Service     Report
                                          DEFAULTS: PRDBATCH    _______
  ___  1  SI         DB2P      __                  PRDBATCH    _______
  ___  2   CN          ONLINE  __                  PRDONLIN    _______
  ___  2   PRC         PAYPROC __                  PRDONLIN    _______
  ___  2   UI          SYSADM  __                  PRDONLIN    _______
  ___  2   PK          QMFOS2  __                  PRDQUERY    _______
  ___  1  SI         DB2T      __                  TESTUSER    _______
  ___  2
PR         PAYPROCT __                  TESTPAYR    _______
****************************** BOTTOM OF DATA ******************************
```

Examples of Thread Attributes

- AI (Accounting Information)
- CI (Correlation Information)
- CN (Collection Name)
- CT (Collection Type)
- LU (LU Name)
- NET (Net ID)
- PK (Package Name)
- PN (Plan Name)
- SI (Subsystem Instance)
- UI (Userid)

- All DB2P applications accessing their first SQL package in the collection ONLINE are in service class PRDONLIN.
- All DB2P applications that call stored procedure PAYPROC first are in service class PRDONLIN.
- All work performed by DB2P user SYSADM is in service class PRDONLIN.
- Users other than SYSADM that run the DB2P PACKAGE QMFOS2 are in the PRDQUERY class. (The QMFOS2 package is not in collection ONLINE.
- All other work on the production system is in service class PRBBATCH.
- All users of the test DB2 system are assigned to the TESTUSER class except for work that first calls stored procedure PAYPROCT, which is in service class TESTPAYR.

© 2011 IBM Corporation

DB2 Workload Priority
How Does WLM Assign Priority To DB2 Workload?

- **The priority of the DB2 workload will vary depending upon the origin of the workload**
- **DB2 workload originating from a local application (IMS, CICS, TSO, Batch)**
  - Priority is inherited from the invoking application
  - This applies to Stored Procedures invoked locally
- **DB2 Distributed requests (DDF)**
  - Priority controlled by DDF Service Class definitions
- **DB2 Stored Procedure request via DDF**
  - Priority controlled by Service Class definitions

© 2011 IBM Corporation

About Enclaves

- **Enclaves represent a "business unit of work"**
- **Enclaves are managed separately from the address space**
- **Enclaves can include multiple SRBs/TCBs**
  - Can span multiple address spaces
  - Can have many enclaves in a single address space
  - Assigned by WLM to a service class
- **Enclaves are managed separately from the address space**
- **DB2 exploits the WLM Enclave interface**
  - DB2 workload may create/delete the enclave, join an enclave, etc.

© 2011 IBM Corporation

DDF Goal Considerations

- **DDF workload may call for a combination of Velocity and Response time goals**
  - DDF address space versus DDF workload

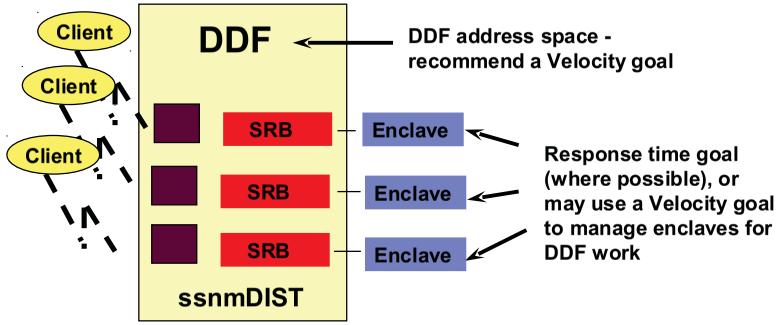

© 2011 IBM Corporation

DDF Stored Procedure Priority

- **When a Stored Procedure is called from DDF thread**
  - DB2 references the enclave created for the DDF request for Stored Procedure
  - Stored Procedure priority is the priority of the DDF request

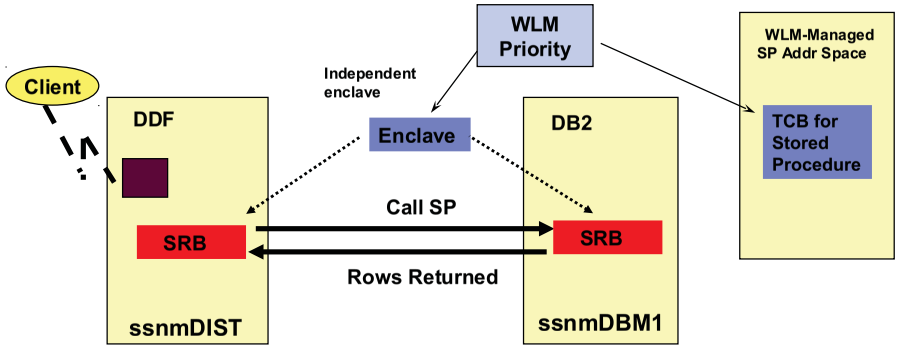

© 2011 IBM Corporation

Stored Procedure Priority
Called From A Local Application

- **When a Stored Procedure is called from an application on z/OS**
  - DB2 creates an enclave for use by the Stored procedure
  - Stored Procedure priority is the priority of the calling application address space

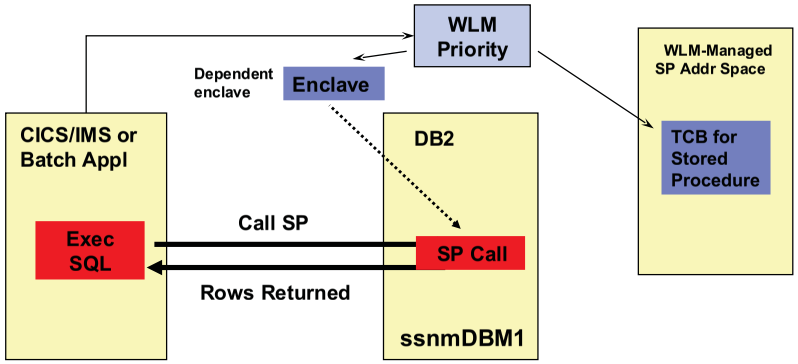

© 2011 IBM Corporation

Service Classes And Thread Priorities

```
                ZENCLD   VTM     O2       V310./C D81G 02/22/06  7:13:42   2
> Help PF1                                                              Back PF3

>       THREAD INFORMATION:  Enter a selection letter on the top line.

> A-THREAD DETAIL B-LOCK COUNTS C-LOCK WAITS     D-LOCKS OWNED  E-GLOBAL LOCKS
> F-CURRENT SQL   G-SQL COUNTS  H-DISTRIBUTED    I-BUFFER POOL  J-GROUP BP
> K-PACKAGES      L-RES LIMIT   M-PARALLEL TASKS N-UTILITY      O-OBJECTS
> P-CANCEL THREAD Q-DB2 CONSOLE R-DSN ACTIVITY   S-APPL TRACE   *-ENCLAVE
=================================================================================
                        ENCLAVE DETAIL INFORMATION
 PLAN
+ Thread:  Plan=DEMO1     Connid=DB2CALL  Corrid=DEMOENCL     Authid=CXE12AUR
+ Attach:  BATCH          JOB Name=DEMOENCL                   JOB Asid=  49
+ Package: DEMO1          Collection=
 enc
+ ENCLAVE TOKEN:     2000000001          Enclave Type:       Dependent
+ Owning System:     SP12                Owning Job:         DEMOENCL
+ WLM Mode:          Goal                Enclave CPU Time    00:00:00.000
+
+ SERVICE PERIOD INFORMATION
+ Period(s) for Service Class BATCH:     2
+ Current Period for This Thread:        1
+ Performance Index This Period:            .20
```

Service Class

Enclave type

Period

PI

- **Depending upon how an enclave is created (local allied address space or via DDF) controls what service class, etc that is assigned to a given thread**

© 2011 IBM Corporation

How DB2 Interacts With WLM To Assign I/O Priorities

- **DB2 informs z/OS about which address space's priority is to be associated with a particular I/O request**
  - WLM handles the management of the request

Table 8-1 How read I/O priority is determined

Request type	Synchronous reads	Prefetch reads
Local	Application's address space	Application's address space
DDF or sysplex query parallelism (assistant only)	Enclave priority	Enclave priority

Table 8-2 describes to which enclave or address space DB2 is associated with the I/O write requests.

Table 8-2 How write I/O priority is determined

Request type	Synchronous writes	Deferred writes
Local	Application's address space	ssnmDBM1 address space
DDF or sysplex query parallelism (assistant only)	DDF address space	ssnmDBM1 address space

Table from SG24-6472-03

© 2011 IBM Corporation

WLM Contention Management

- **WLM Contention Management helps addressing chronic or long lasting contention situations**
  - WLM provides interfaces to allow resource managers (for example – DB2) to signal contention situations
  - WLM has had the ability to promote (increase the DP) for a short duration to resolve the issue
- **DB2 example scenario**
  - Lock/latch contention in DB2 may impact performance
  - Often contention may be resolved with a short boost of resource
  - DB2 may notify WLM if a contention occurs
  - WLM may optionally raise the priority for the holder to complete the work
- **WLM can promote units of work for longer periods of time, and promote them to the priority of the highest-priority units of work waiting for a resource they are holding.**

© 2011 IBM Corporation

WLM Considerations – Example Stored Procedures

- **The original assumption**
  - All work requests inserted by DB2 (example – Stored Procedures) were independent requests
- **The reality**
  - Procedures may recursively call other procedures
  - The processing may be inter-dependent
- **The newer logic**
  - DB2 tells WLM about dependent stored procedure requests
  - WLM gives dependent requests priority
    - WLM may make adjustments, if needed

© 2011 IBM Corporation

WLM Considerations For Nested Stored Procedure Requests

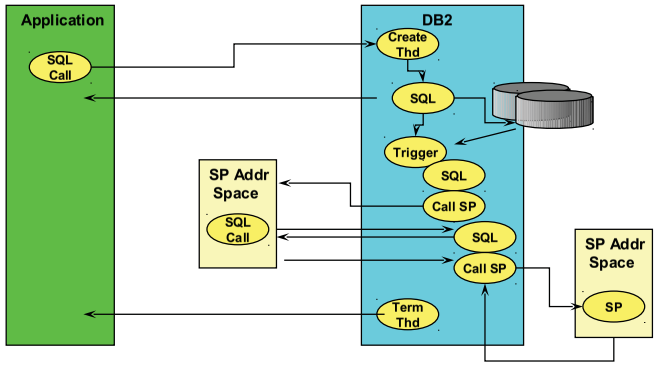

- **Triggers, Stored Procedures, and UDFs actions may be nested, sometimes multiple layers of nesting**
- **DB2 tells WLM about dependent stored procedure requests**
  - WLM may give dependent requests priority, if needed
  - WLM may start server regions more aggressively, if needed

© 2011 IBM Corporation

DB2 And zIIP Processors

- **Work on z/OS may have all or a portion of its resource usage on an enclave Service Request Block**
  - Enclave SRB work may be directed to the zIIP
- **Certain types of DB2 work may take advantage of zIIP**
  - DRDA - Queries that access DB2 for z/OS via DRDA over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP
  - Complex parallel queries
  - DB2 utilities for index maintenance
    - LOAD, REORG, and REBUILD
- **WLM and new enclave structures to manage zIIP related workload – work dependent enclave**
- **DB2 V10 – Sequential prefetch eligible for zIIP processor**

© 2011 IBM Corporation

About Work-dependent Enclaves

- **Extension to an independent, dependent, or other work-dependent enclave**
  - Extends the transaction the creating enclave.
- **Allows control of zIIP offload by entitled products.**

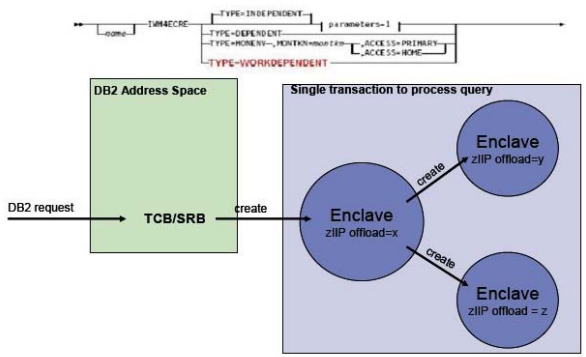

© 2011 IBM Corporation

DB2 DDF Considerations
Things To Note

- **For DDF inactive threads**
  - Consider a two-period service class with a response time goal where 80-90% of the transactions complete in first period
- **For DDF active threads**
  - Consider velocity goals and use a single-period service class
- **Look for overly simplistic Service Class definitions**
  - Example – type DDF and nothing more than DB2 subsystem name
    - Does little to exploit the ability of WLM to prioritize DB2 workloads
    - Some workloads will inherently be more important than others
- **Look for Distributed workloads that run longer than expected but use little resource**
  - Indicative of workload that drops into less than optimal Service Classes – example SYSOTHER

© 2011 IBM Corporation

DB2 Velocity Goals
Setting Optimal Goals

- **Use Velocity goals for always running and long running work**
- **Use a Velocity goal for the DB2 DDF address space**
  - DDF address space has internal tasks that govern thread creation that should have high performance goal
- **Use a Velocity goal for**
  - DB2 address spaces (SSAS and DBM1)
  - CICS and IMS regions (if not using response time goals)
- **Velocity goals at first glance seem easy to set**
  - Require more ongoing review
  - Should be validated as the operating environment changes – changes to operating system, hardware, and workload

© 2011 IBM Corporation

DB2 Response Time Goals
Setting Optimal Goals

- **Use Response Time goals when possible**
  - Less need for ongoing maintenance and review
  - WLM will manage resources dynamically to achieve goals
- **Response Time goals work well for certain types of DB2 workloads**
  - DB2 Distributed workloads in e-business and WebSphere transactional type workloads
  - Transactional type workloads in general including distributed workloads that invoke Stored Procedures
  - Repetitive workloads that have multiple events for WLM to measure and manage

© 2011 IBM Corporation

WLM
General Recommendations

- **KISS method (Keep It Simple Stupid)**
  - Avoid overly complex implementations
    - Too many Application Environments and Service Classes can result in an over abundance of server address spaces
    - Increases number of queues that WLM must manage
- **Avoid too many service classes**
  - WLM analyzes service classes in a round-robin manner
  - Too many and WLM is unable to manage them all effectively
  - Note – With z/OS V1.11 WLM is implementing features like the skip clock function

© 2011 IBM Corporation

What's Improved In z/OS V1.12

- **z/OS Management Facility (z/OSMF) V1.12 (5655-S28) provides an improved GUI management interface for z/OS**
- **A new system management task, Workload Management (policy editor), can simplify the creation, modification, and review of z/OS WLM service definitions**
- **A new application, Sysplex Status and Monitoring Desktops tasks, can provide real-time status of resources of all your servers, sysplexes, as well as Linux images from one location**
- **The Configuration Assistant for the z/OS Communications Server (available since z/OSMF V1.11) is updated with support for IP security**
- **The Incident Log capability (available since z/OSMF V1.11)**

© 2011 IBM Corporation

WLM Tools

http://www.ibm.com/servers/eserver/zseries/zos/wlm/tools/

- WLM Service Definition Formatter - GUI
- WLM Work Queue Viewer - ISPF
- WLM Service Definition Editor - GUI

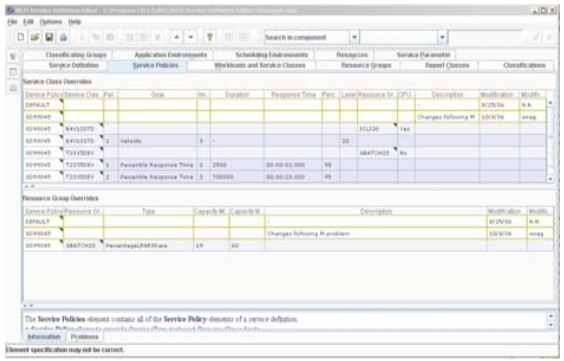

WLM Tools

- **WLM Policy Editor is available on the z/OS Management Facility**
  - All the same function as in the Web-download tool and many new features
  - Direct access to the WLM Couple Data Set to install/extract service definitions. No need to FTP WLM policy files!
  - Activation of service policies and monitoring of the WLM status in the sysplex
- **Requires z/OSMF V1.12 and z/OS V1.12**

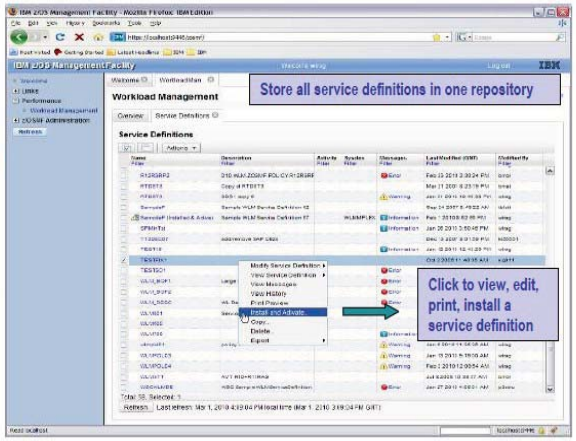

© 2011 IBM Corporation

WLM Tools

Editing service definitions

- **Simplified creation, modification and review of service definitions**
  - Policy elements are presented in tables
  - Tables can be edited, filtered, and sorted
  - Best-practice hints are displayed automatically
  - Several service definitions can be opened simultaneously
  - Cut, Copy, Paste of policy elements between service definitions

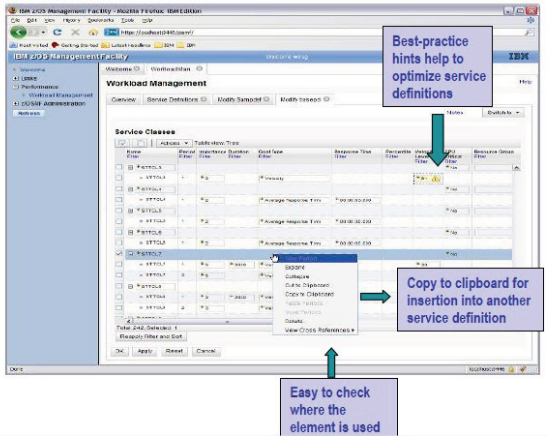

© 2011 IBM Corporation

Summary

Workload Manager As The Priority Manager Of DB2

- **Workload Manager (WLM) is the priority and resource manager for z/OS and DB2**
- **WLM manages critical DB2 address spaces**
- **These resources have an impact on how DB2 applications perform**
- **Workload Manager is used to manage DB2 workflow**
- **Important to have a basic understanding of WLM**
- **WLM is constantly being enhanced to provide new features and functions**

© 2011 IBM Corporation

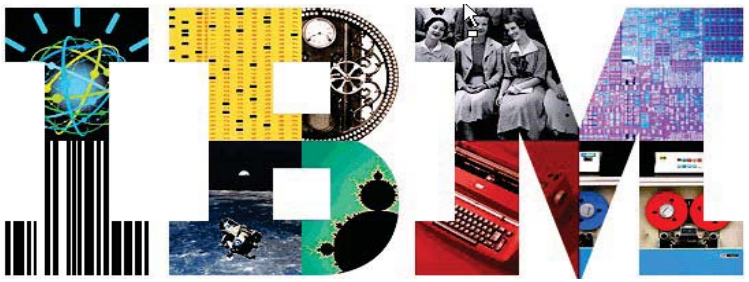

© 2011 IBM Corporation

Back-up slides

© 2011 IBM Corporation

Classifying Work

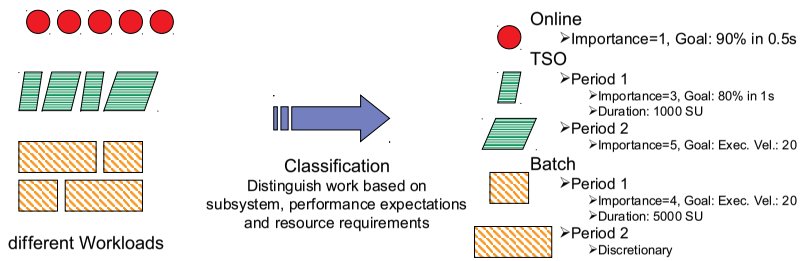

- **Work needs to be classified**
  - In order to distinguish it and in order to control it
  - Result
    - Work is classified into service classes
    - A service class describes the performance characteristics for work with same behavior and same performance expectations

© 2011 IBM Corporation

Determine Goal Achievement

Performance Index (PI)

- **Indicator how well work (service class) is doing**
  - Two exist per service class
    - Sysplex and Local PI
- **Basically**
  - If PI > 1: service class is eligible for help
  - If PI < 1: service class is a potential donor
  - Independent from business importance
- **How is a service class selected**
  - By Business Importance
  - By worst PI

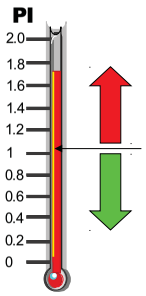

$$\text{Response Time Goal} \quad : \quad PI = \frac{\text{Actual Achieved Response Time}}{\text{Response Time Goal}}$$

$$\text{Execution Velocity Goal} \quad : \quad PI = \frac{\text{Execution Velocity Goal}}{\text{Actual Achieved Execution Velocity}}$$

© 2011 IBM Corporation

Example: How does WLM work

- **Test scenario**
  - CICS, IMS, TSO and Batch competing for resources
  - Start
    - System is utilized at about 90%
    - CICS, IMS and TSO show a constant load
    - Batch shows nearly no load
  - Test
    - Start a lot of batch work
    - System gets up to 100% utilized
    - Observe how the system behaves

Service Class	Per	Goal	Value	Imp
CICS	1	Avg RT	0.1 s	2
IMS	1	Avg RT	10 s	3
TSO	1	Avg RT	0.1 s	3
	2	Avg RT	1.0 s	3
	3	Avg RT	3.0 s	4
Batch	1	ExVel	10	5

© 2011 IBM Corporation

Management of Business Units of Work

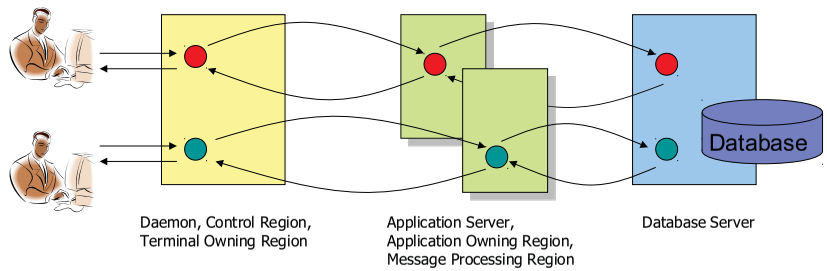

- **Subsystems and Middleware communicate with WLM**
  - WLM understands the flow of work (Business Units of Work)
  - WLM is able to manage the Business Units of Work
- **Work requests span multiple address spaces**
  - Typically for Transaction Monitors like CICS, IMS, Websphere, DB2 Stored Procedures, etc
- **WLM provides interfaces to give these subsystems the ability**
  - To tell WLM when transactions arrive and when they end
  - To tell WLM which address spaces, TCBs and SRBs execute programs on behalf of the transactions
- **Two facets**
  - Management for CICS and IMS
  - Management for exploiter of Enclave Services (Websphere, DB2 Stored Procedures, DB2 Distributed Data Facility)

© 2011 IBM Corporation

Subsystem Participation

Application, Subsystem	Classification	Monitoring Services	Enclaves	Queue/Server Management	Routing/DNS Services	Special Functions
CICS	X	X				
IMS	X	X				
DB2	X	X (as extension)	X (*R/D)	X	X	Bufferpool Management Contention Management
Websphere (MQSeries Workflow) Scalable Webserver	X	X (reporting only)	X (*S, *L)	X	X	
SAP R/3	X		X			
JES2/3	X			X (Batch)		Scheduling Environments
VTAM, TCP/IP					X	
Netview	X		X			
Oracle 9i	X		X			
TSO, APPC, OMVS, STC, LSFM	X					

© 2011 IBM Corporation

Websphere z/OS Integration

Merges the best 30 years of critical transaction Monitors and the J2EE runtime

- Isolation
- Availability
- Consistency
- Resource Management
- Two-Phase Commit

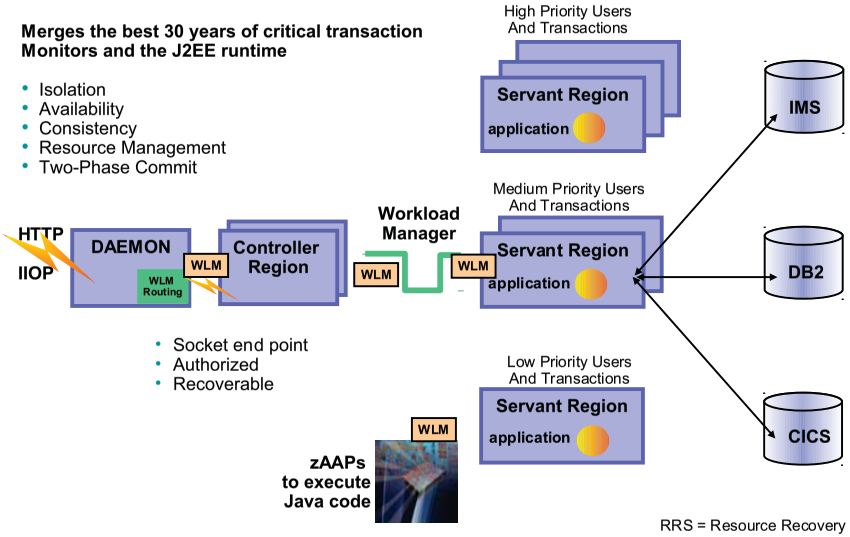

- Socket end point
- Authorized
- Recoverable

RRS = Resource Recovery Services

© 2011 IBM Corporation

SQLADRIA SEMINAR

Ljubljana, 15 April 2011

User Topic: DB2 Explain

Bernard Zver
Informatika

Sponsored by

DB2. Universal Database for z/OS

DB2 Explain

Bernard Zver

SQLAdria

Ljubljana, 15. April 2011

Agenda

- **DB2 terms**
  - DB2 Optimizer
  - Filter Factor
  - Clustering
  - 3R
  - Join method, Star Join
  - Stage predicates
  - Access Path
- **Explain**
  - Explain tables
  - Explain query
  - Recommendation
  - New DB2 Features

DB2

- IBM announced DB2 V1.1 in year 1983
- Informatika
  - DB2 V2.2 for OS/390 - z/OS in year 1990

Don Haderle

Roger Miller

DB2 Facts

- That the majority of performance problems among applications are caused by poorly coded programs or improperly coded SQL.
- That poor performing SQL is responsible for as much as 80 percent of response-time issues.
- That most developers need educated in how to read and analyze the DB2 Explain output.

DB2 Optimizer

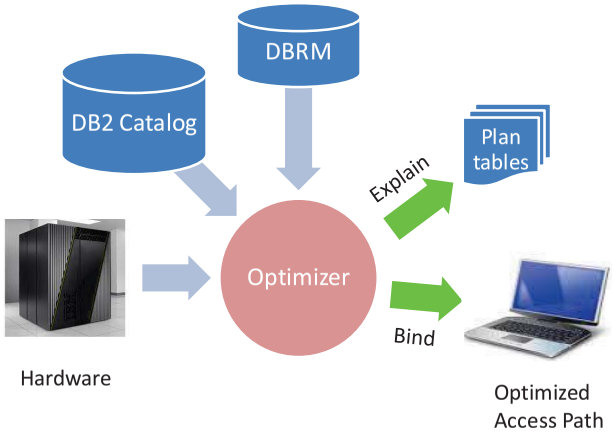

DB2 Optimization Process

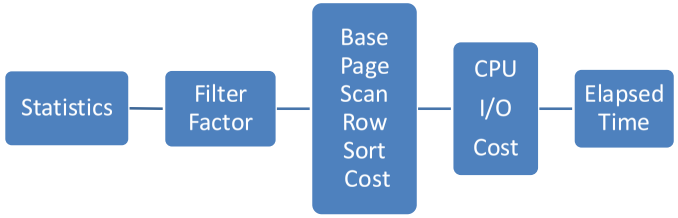

Filter Factor

The filter factor of a predicate is a number between 0 and 1 that estimates the proportion of rows in a table for which the predicate is true.

FF = 1/COLCARDF

sysibm.syscolumns

Column C1 in table T contains only five distinct values:
A, B, C, D, E.

C1 = 'E' FF = 0.2

Col = constant	1/colcardf	1/25	default
Col <> constant	1 - (1/colcardf)	1 - (1/25)	default

Clustering

CLUSTERRATIOF (sysibm.sysindexes) gives an indication of how closely the order of the index entries on the index leaf pages matches the actual ordering of the rows on data pages.

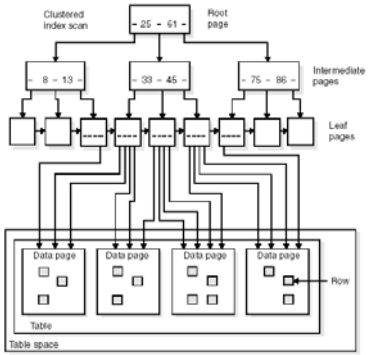

Clustering has a huge impact on a number of getpages that must be accessed.

50 rows using a **non-clustering** index could mean 50 getpages while 100 rows using a **clustering** index could mean just 1 or 2 getpages.

3R

- RUNSTATS
  - Run runstats after reorganizing the data or indexes
  - Mass INSERT/DELETE operations
- REORG
  - Using useful catalog queries DSN910.SDSNSAMP (DSNTESP)
  - Tablespace: Query 15
  - Index: Query 16
- REBIND
  - Rebind Plan/Package when performance degrades to an unacceptable level
  - Clusterratiof changes to less than 80%
  - NLEAF, NPAGES changes more than 20%
  - Cardinality changes more then 20%

Join more tables

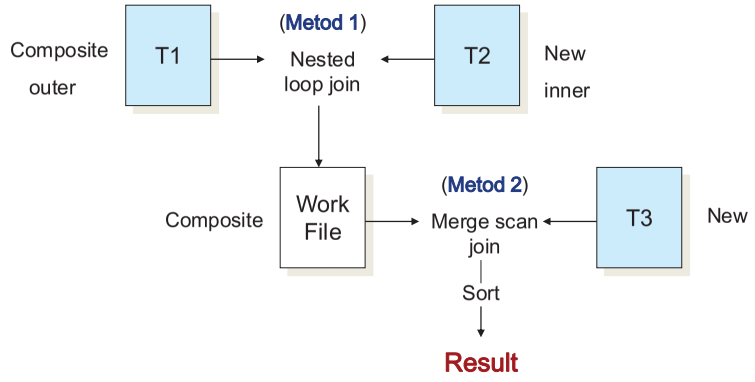

Join method

- Method=0

  DB2 accesses the first table (this method is allways)

- Method=1

  Nested loop join. DB2 scans outer table and for each row searches for matching rows in inner table.

- Method=2

  Merge scan join. DB2 scans both tables in the order of the join columns.

- Method=4

  Hybrid join.

Star Join

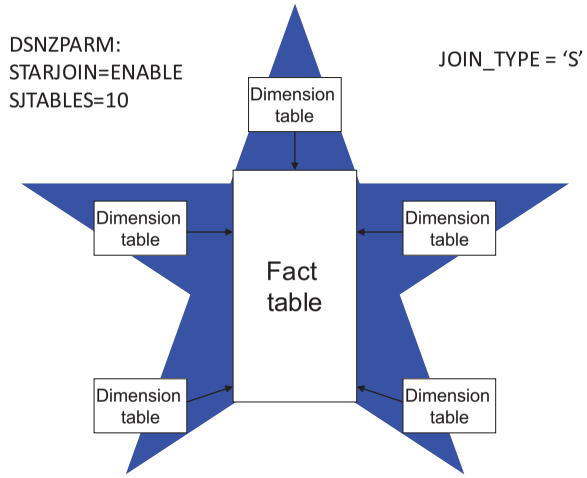

Stage Predicates

Predicate Type	Indexable?	Stage 1?
COL = value	Y	Y
COL = noncol expr	Y	Y
COL IS NULL	Y	Y
COL op value	Y	Y
COL op noncol expr	Y	Y
COL BETWEEN value1 AND value2	Y	Y
COL BETWEEN noncol expr1 AND noncol expr	Y	Y
value BETWEEN COL1 AND COL2	N	N
COL BETWEEN expression1 AND expression2	Y	Y
COL BETWEEN COL1 AND COL2	N	N
COL LIKE 'pattern'	Y	Y
COL IN (list)	Y	Y
COL <> value	N	Y
COL <> noncol expr	N	Y
COL IS NOT NULL	Y	Y
COL NOT BETWEEN value1 AND value2	N	Y

Access Path

- ACCESSTYPE is I, I1, N, MX(MI, MU), DX, …

  DB2 uses an index to access the table.
  - MATCHCOLS > 0 means matching index scan
  - Index screening

```
select * from T                Index on T(c1,c2,c3,c4)
where c1 = 7                   matchcols = 1
  and c3 > 0                   screening predicate
  and c4 = 5                   screening predicate
```

  - MATCHCOLS = 0 means nonmatching index scan (reads all RIDs).
  - INDEX_ONLY = Y

- ACCESSTYPE is **R** Table space scan.

  PAGE_RANGE = Y, DB2 reads only selected partition(s).

Explain tables

16 explain tables:

- plan_table
- dsn_statement_table
- dsn_predicat_table
- dsn_filter_table
- …

DSN910.SDSNSAMP(DSNTESC)

Runstat

```
EXPLAIN PLAN SET QUERYNO=1 FOR
SELECT NAME, CARDF, RECLENGTH
FROM SYSIBM.SYSTABLES
WHERE CREATOR='SYSADM'
  AND DBNAME='DBPAVTO'
  AND NAME LIKE 'TR9%'
  AND TYPE='T';
```


Explain query

```
select a.progname, method, creator, tname, accesstype, matchcols,
       accessname, indexonly, prefetch, procms
from plan_table a,
     dsn_statemnt_table b
where a.queryno = b.queryno
```

Progname	Method	Creator	Tname	Access type	Match cols	Access name	Index only	Prefetch	Procms
DSN@EP2L	0	SYSIBM	SYSTABLES	I	2	DSNDTX01	N		1

Explain query

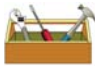

```
Select stage, keyfield, text,
       decimal(filter_factor,13,12) filterfactor
from nadi.dsn_predicat_table p,
     nadi.dsn_filter_table f
where p.predno = f.predno
```

Stage	Keyfield	Text	FilterFactor
MATCHING	Y	SYSIBM.SYSTABLES.CREATOR='SYSADM'	0,101008474826
STAGE1	Y	SYSIBM.SYSTABLES.DBNAME='DBPAVTO'	0,001056328183
MATCHING	Y	SYSIBM.SYSTABLES.NAME LIKE 'TR9%'	0,000284779351
STAGE1	N	SYSIBM.SYSTABLES.TYPE='T'	0,142857134342

Sql Tips

- year(birthdate) = 1987 — stage2
  index on expression year(birthdate) — indexable

  birthdate between '1987-01-01' and '1987-12-31' — stage1, indexable
- substr(lastname,1,1) = 'Z' — stage2
  lastname like 'Z%' — stage1, indexable
- dec_col152 = :hv_int_col — stage2
  dec_col152 = DECIMAL(:hv_int_col,15,2) — stage1, indexable
- char4_col > 'abcd_' — stage2
  char4_col > 'abcd' — stage1, i

Top Ten Development Best Practices

1. Minimize network calls
2. Minimize passes through the data
3. Put the work into the SQL, not the program
4. Unlearn the “flat file” mentality
5. Be sure data type and length match in predicates
6. Know your Stage 1, Stage 2, and Indexable predicates
7. Document your code
8. Always check the SQLCODE or GET DIAGNOSTICS
9. Analyze your access paths (and tune your SQL in test)
10. Avoid Bachelor Programming Syndrome

by Craig S. Mullins

Top Ten Common SQL Mistakes

1. Assuming an ORDER BY is not necessary for ordered results
2. Forgetting the NULL indicator
3. Asking for more data than you need (columns and/or rows, select *)
4. Coding predicates appropriately in Outer Joins
5. Not coding a cursor for a multi-row result
6. Recompiling but not rebinding
7. Forgetting to use single quotes around strings (instead of double quotes)
8. Trying to modify a Primary Key column
9. Forcing dynamic SQL into static SQL (sometimes hundreds of static SQL statements)
10. Incorrect expectations when using the NOT IN predicate with NULLs *

by Craig S. Mullins

NOT IN Predicate

*

T1

id	value
1	30
2	90
3	140

T2

id	name
1	V7
null	V8
4	V9

```
SELECT * FROM T1
WHERE id NOT IN ( SELECT id FROM T2 )
```

null could be any number

New DB2 V8/V9 SQL Features

- Set operations 'Intersect' and 'Except'
- Merge statement for 'Upsert' processing. Insert or Update
- OLAP features for Ranking and Numbering of data
- Native SQL Stored Procedures
- 'Instead of' Triggers
- New support and SQL for XML data
- Indexing on Expressions
- Skipped locked data
- Truncate statement

New DB2 V8/V9 SQL Features

- More Stage 1 predicates
- Multi Row Fetch, Update, and Insert
- Multiple Distincts
- Expressions in the 'Group By'
- Common Table Expression
- Sequences versus Identity Columns
- Materialized Query Tables (MQTs)
- Recursive SQL
- More efficient use of indexes. Forward and Backward scans
- New 'Get Diagnostics' for warning and error information
- Select from an Insert statement
- Order by collation_key (*col_name*,'UCA410_LSL_RSI')

New DB2 V8/V9 SQL Features

- Array host variables now supported for stored procedures
- Timestamp auto update for Inserts and Updates
- New DECFLOAT datatype
- Select from Update or Delete getting old or new values
- Fetch First, Order BY within subqueries
- REOPT AUTO
- Data Studio for Native Stored Procedures
- New XML functions and datatypes

SQL to XML

```
select
 xml2clob(
   xmldocument(
     xmlelement(name "Emp",
       xmlelement(name "Id", xmlattributes(empno),
         xmlelement(name "PersData",
           xmlforest(firstnme    as "Firstname",
                     lastname    as "Lastname",
                     birthdate   as "Birthdate"
                    )
                   ),
         xmlelement(name "JobData",
           xmlforest(workdept    as "Department",
                     job         as "Job",
                     hiredate    as "Hiredate"
                    )
                   )
                 )
               )
              ,xmlcomment('xml document')
              )
         ) AS "XML Rezultat"
from dsn8910.emp

XML Rezultat
------------
<LOB value>
<LOB value>
```

```
<Emp>
 <Id EMPNO="000180">
   <PersData>
       <Firstname>MARILYN</Firstname>
       <Lastname>SCOUTTEN</Lastname>
       <Birthdate>1949-02-21</Birthdate>
    </PersData>
    <JobData>
       <Department>D11</Department>
       <Job>DESIGNER</Job>
       <Hiredate>1973-07-07</Hiredate>
    </JobData>
  </Id>
</Emp>
<!-- xml document  -->
```


Literatura

DB2 for z/OS - Technical Resources

- http://www-01.ibm.com/support/docview.wss?rs=64&uid=swg27011656

DB2 9 for z/OS Technical Overview

- http://www.redbooks.ibm.com/abstracts/SG247330.html

- IDUG 2010