


## Intelligent Business Using DB2 for i

**Doug Mack**

**Mackd@us.ibm.com**

Follow us [@IBMpowersystems](#)

Learn more at [www.ibm.com/power](http://www.ibm.com/power)



# The Agenda

- **What the heck is Big Data?**
  
- **Now let's talk about Analytics for the i client**
  - DB2 for i Enablers for BI/Analytics
  
  - Architectural Considerations/Solutions
  
  - The BI Tooling

# Analytics is expanding from enterprise data to big data

## Volume

**12** terabytes  
of Tweets create daily

## Velocity

**5** million  
trade events per second

## Variety

**100's** video feeds  
from surveillance cameras

**350** billion  
meter readings per annum

**500** million  
call detail records per day

**80%** data growth  
are images, video, documents...

Suppose you had to scan 100Terabytes of data  
With 1 node of a system with 50MB per sec, it would take 23 Days.  
Imagine if you had 1000 nodes – It would be 30 minutes!

## Big Data For Crime Fighting

- **Memphis Police Department - Blue CRUSH (Criminal Reduction Utilizing Statistical History) crime-fighting methodology**
  
- **Enables rich picture of the various circumstances surrounding a particular crime:**
  - *Where did it occur?*
  - *What was the weather?*
  - *What was the weather the day or week before?*
  - *Did the crime take place over a three-day weekend?*
  - *Was there a payday for government workers?*
  - *Was there a gun show at the local convention center?*
  - *Was there a recruitment day at a local university?*

### Memphis PD: Using analytics to fight crime before it happens.

By following crime trends, Memphis Police can address, and even prevent, new ones.

### How Memphis PD Got Smarter

Larry Godwin, Director of Police Services for Memphis, pioneered a way to focus existing patrol resources more intelligently—which has reduced overall crime volume by 30%.

## New Health Care Applications

Healthcare is “dying of thirst in an ocean of data”

*“Medicine has become too complex. Only about 20% of the knowledge clinicians use today is evidence-based.”*

**Steven Shapiro**

Chief Medical & Scientific Officer  
University Pittsburgh Medical Center

Medical info is doubling every 5 years

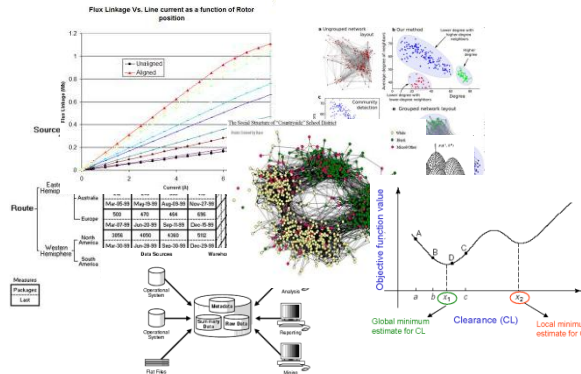
81% of physicians spend < 5 hrs / month reading medical journals

1.5M errors in the way medications are prescribed, delivered and taken

\$750B, or 30 cents of every dollar, is wasted in US alone

# Scalable: Different types of analytics require a **scalable** IT infrastructure

**Different types of analytics (OLTP, Data warehouse, Streaming Data, OLAP, Operational Analytics, Ad-hoc reporting, time series, deep analytics)...**



...need to access data differently and require compute *and* storage resources that are distinctly different and often highly scalable.



**Analytics challenge:** Vestas is addressing the challenge of wind turbine placement by analyzing petabytes of varied data such as weather reports, tidal phases, geospatial and sensor data, satellite images, deforestation maps, and weather modeling to pinpoint installation. The output drives the business model for their client’s wind turbine farms.

**IT Architectural approach:** Vestas realized early on that a new IT architecture was crucial to achieving this objective. To that end, they deployed an IBM system capable of 150 trillion calculations per second, combined with a distributed parallel file system with near storage compute capabilities. All within a highly energy efficient footprint.

**Benefit:** The analysis, which used to take weeks, can now be done in less than one hour.

*“Using IBM software and systems, we can now answer these questions quickly to identify new markets for wind energy and help our clients meet aggressive renewable energy goals.” – Lars Christian Christensen, Vice President of Plant Siting and Forecasting, Vestas Technology R&D*

[Vestas IBM Video](#)


## So What is IBM Doing about Big Data?

- **Creating offerings and services around key technologies**
  - Large Scale Data warehousing
  - Statistical Analysis
  - Stream Analysis
  - Hadoop
  - Massively parallel processing clustered configurations
  - Query Acceleration
- **It addresses the EXPONENTIAL growth of data !**
  - 90% of data has been created in the last 2 years. WOW
    - Sensors gather climate information
    - Posts to social web sites
    - Cell phone GPS signals
    - Web page interaction



# Example: IBM Systems for Smarter Analytics

**IBM Systems is the market leader in support of analytical workloads**



**Workloads:**

- Data Analysis/Data Mining
- Data Warehousing/Data Mart
- Scientific/Engineering/Industrial R&D

IDC Server Workload study 2011



## What the Analysts are saying

*“When we commenced this exercise we expected to find that there were some areas in which IBM excelled and others in which Oracle did so. We have been surprised to find that that is not the case and that the **IBM Smart Analytics System out-competes Oracle Exadata in almost every area we have examined**” Philip Howard, Bloor Research*

*Three-year costs for Smart Analytics System 7700 are **43 and 40 percent less than those for Oracle and Teradata systems** respectively. Source: Cost/Benefit Case for Enterprise Warehouse Solutions, International Technology Group*

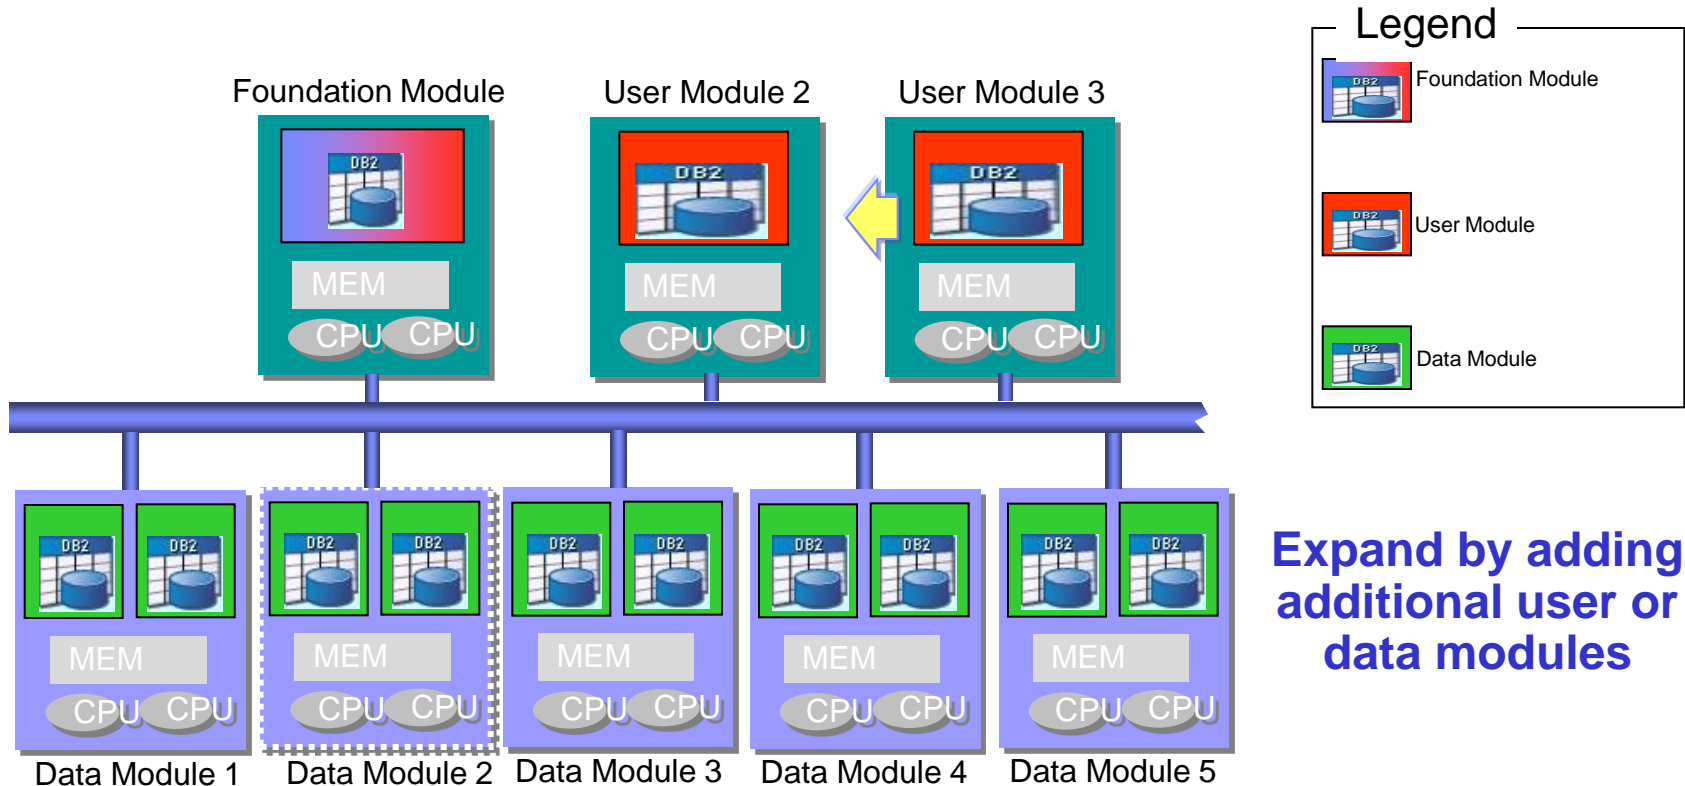
## IBM and Analytics at a glance:

- More than **\$14B** in Acquisitions Since 2005
- More than **10,000** Technical Professionals
- More than **7,700** Dedicated Consultants
- **Largest** Math Department in Private Industry
- More than **27,000** Business Partner Certifications



# Scaling out to support more data or more users

Shared Nothing, Massively Parallel Design



**Expand by adding additional user or data modules**

- **Balanced system design**
  - System modules with optimal processor, memory, and I/O specifications
- **Scale-out by adding additional system modules**
  - Which always include balanced I/O
- **Proven “best practice“ for large scale data warehousing**

**Power is performance redefined**

## Big Data Technologies

### ▪ Info Streams

- IBM Research developed language to analyze web data in real time
- Info Streams is a specific proprietary language using commodity hardware
- Examples of applying Infostreams
  - USC project to monitor the sentiment of voters in the recent presidential election campaigns – tweet by tweet by tweet !!
  - Challenge...how do you deal with 60,000 tweets per minute ? WOW !!!

### ▪ Big Insights puts an enterprise solution together around Hadoop

- Hadoop is a framework from Apache for running applications (aka jobs) on large clusters built on commodity hardware capable of processing petabytes of data
  - Hadoop implements a computational paradigm named **Map/Reduce**, where the application is divided into self contained units of work, each of which may be executed or re-executed on any node in the cluster
  - The goal is to move the processing closer to the data – not the reverse !!
- Big Insights simplifies the environment with installation, visualizations, enterprise level security and availability, etc.
  - Big Sheets Visualizations; Adapters for DB2, Netezza, etc.
  - PIG, or PIGLatin is the programming language; HIVE is the SQL like query language

### ▪ Social Media Analytics

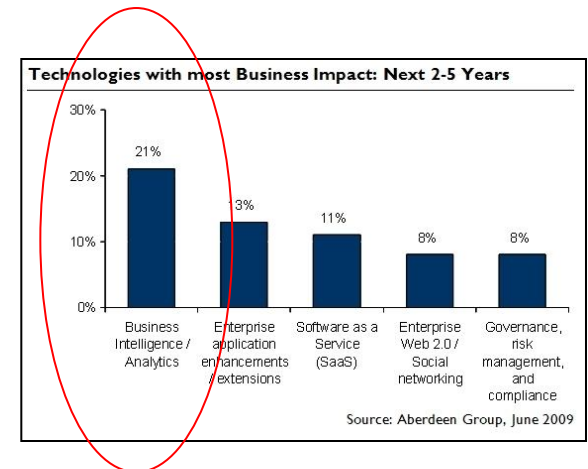
- Also known as Cognos Insights
  - Analyzes web data (what it can....e.g., facebook pages are blocked) to understand “pulse” of the brand or company

# The Agenda

- What the heck is Big Data?
  
- **Now let's talk about Analytics for the i client**
  - DB2 for i Enablers for BI/Analytics
  
  - Architectural Considerations/Solutions
  
  - The BI Tooling

## What about the i Client?

- **SMB companies have similar requirements for analytics, but more constraints:**
  - Budget is smaller
  - Staff is limited– Admin, DBAs
  - *“The full capabilities of most BI tools are overkill for the typical SMB (Small and Medium Business).” – Aberdeen Group*
- **The i market is 85% SMB**
  - Heavily relying on old technology (Query/400) for reporting
- **Requirements for the IBM i SMB BI Client**
  - Low cost
  - Simplified Infrastructure
  - Leverages existing environment (DB2 for i)
  - Reduced dependency on I/T
  - Without sacrificing the benefits of enterprise BI



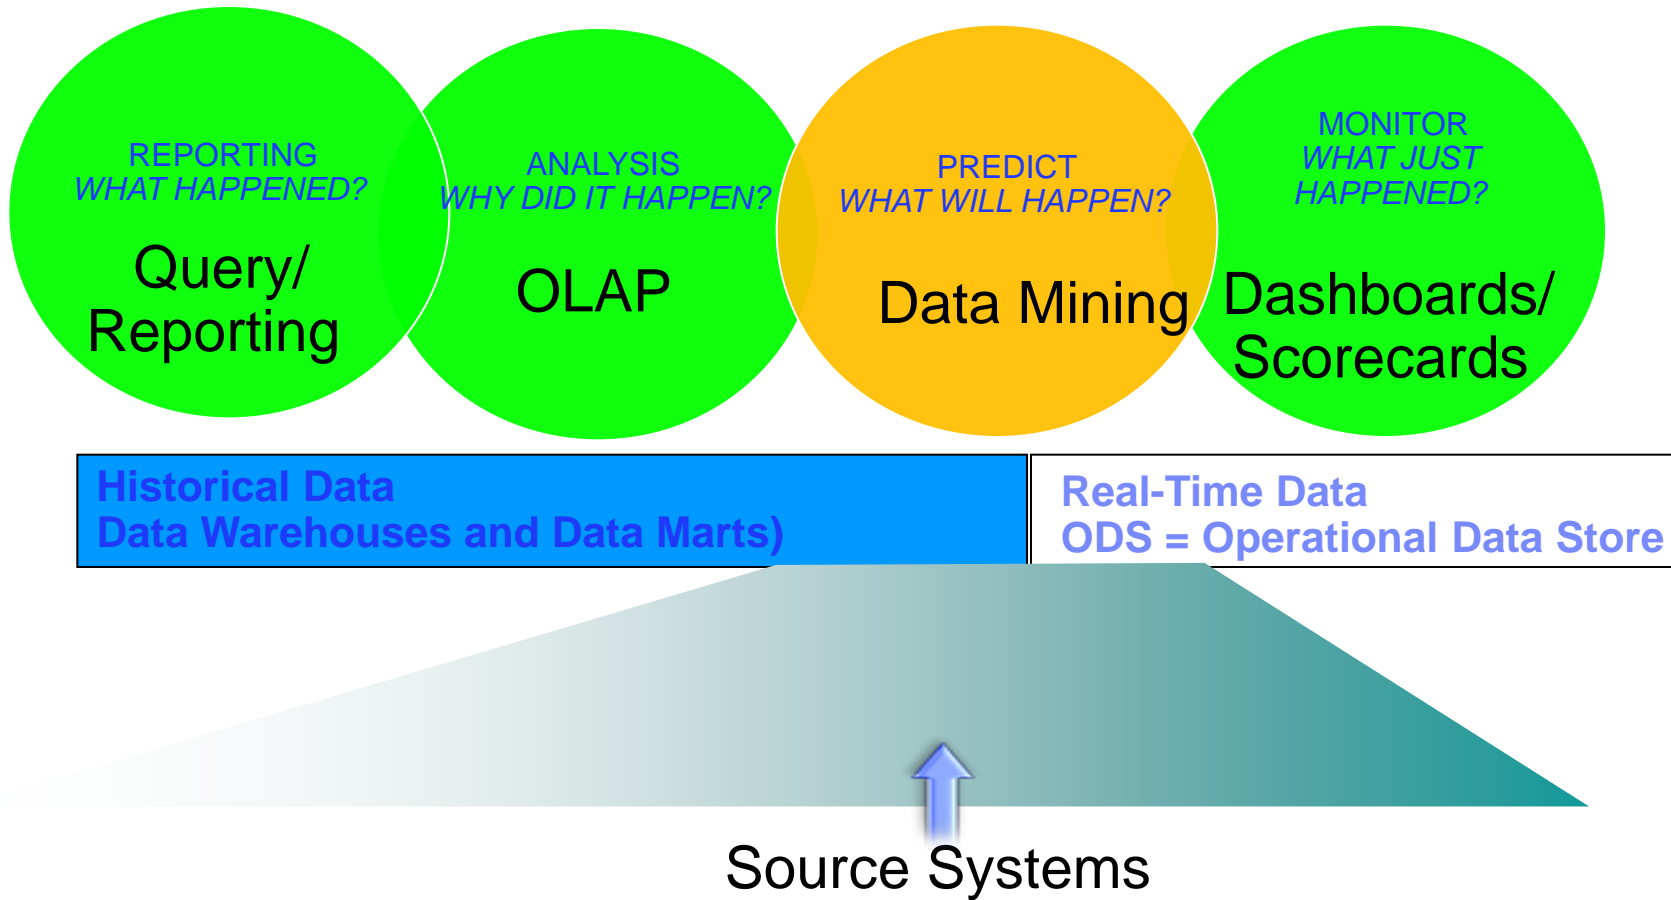
“The IT departments at midsize companies are usually fairly small. They’re often very clever at getting things done more efficiently and with less expense than their larger brethren, but they have their limits. For example, if an SMB’s IT department has 10 people and three or more of them would be needed on a BI team, that’s a substantial percentage of the total IT head count”- Claudia Imhoff

[http://searchbusinessanalytics.techtarget.com/news/2240113703/Copin-g-with-midmarket-BI-project-management-challenges-Tips-for-SMBs?asrc=EM\\_NLN\\_16106429&track=NL-544&ad=859763](http://searchbusinessanalytics.techtarget.com/news/2240113703/Copin-g-with-midmarket-BI-project-management-challenges-Tips-for-SMBs?asrc=EM_NLN_16106429&track=NL-544&ad=859763)

# Analytics and Business Intelligence

**Business Intelligence**

**Predictive Analytics**




Source: The Data Warehousing Institute, Smart Companies in the 21<sup>st</sup> Century, July 2003

# DB2 Web Query for i – Intuitive, Insightful, Extensible

*Fast, easy access to business information assets for query, reporting and analysis*

- Easily spot trends or exceptions in data with real time reports
- Give Executives the means to track how the business is performing through intuitive Key Performance Indicator dashboards
- Create self-service reporting environment that eliminates dependency on I/T
- Provide data to spreadsheet aficionados painlessly and in real time
- Execute and distribute reports in many different formats - on demand or scheduled – via e-mail or saved for later view
- Integrate reports into existing applications for seamless access to data





## Importance of Metadata

### ▪ **Impact of it's absence:**

- the 1999 NASA Mars Climate Observer mission failed because of a metadata problem...
- thrust calculation data was provided in the US measurement scale of pounds/square foot, but was interpreted as metric numbers representing Newtons/second.
- the *data* was correct. The *understanding* and *usage* of it was not.
  
- Result: the orbiter went missing (it probably crashed on Mars) and a \$300M mission failed!



## The Difference between a GUI based query tool and Business Intelligence is a META DATA Layer which DB2 Web Query deploys

**The Meta Data Layer allows you to:**

- **Document your database (which is probably not documented now)**
- **Standardize on the meaning of data in one place (single version of truth)**
- **Expand your community of report authors because they don't have to know the intricacies of the data or SQL**
- **Integrate DB2 functions and stored procedures into your reports (very powerful and flexible)**
- **Define relationships between tables/files, and between data elements to hide complexities of the database and provide advanced analytics like auto drill down capabilities (with no programming required)**

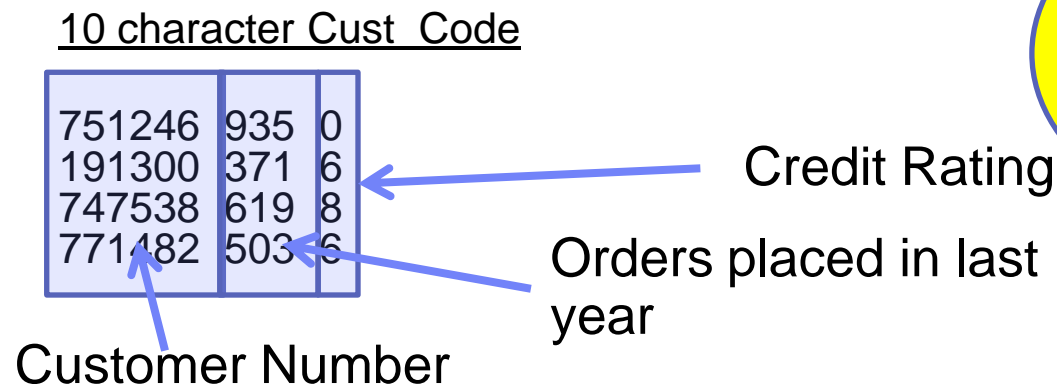
# Leveraging a Meta Data Layer to Shield Complexities of Database

- **Standardize Field/Column Formats**
  - Ex, use commas, set currency symbol, suppress leading zeros
- **Standardize/Decompose Date Fields**
  - Ex: Integer defined as MMDDYYYY
- **Create Filters**
  - Ex: Define a set of countries as “Europe”
- **Define JOINS and Dimensional Reporting**
- **Create Business Views**
  - Organize Columns/fields for easier report development
- **Define Stored Procedures, Views, UDFs to Meta Data Layer for some REALLY COOL function**
  - Row/Column Level Security
  - Text Search Engine
  - Bring in weather report, local news and more

The screenshot displays the IBM Developer Studio interface. The top window shows a file explorer for the 'baseapp' directory, listing various files such as 'cen\_alerts.acx', 'cen\_alerts.mas', 'cen\_currate.acx', 'cen\_currate.mas', 'cen\_hr.acx', 'cen\_hr.mas', 'cen\_inventory.acx', 'cen\_inventory.mas', 'cen\_orders.acx', 'cen\_orders.mas', 'cen\_plant.acx', 'cen\_plant.mas', and 'cen\_stores.acx'. The bottom window shows the 'Join Properties' dialog, which is used for defining database joins. It features two columns: 'Left Source Columns' and 'Right Source Columns'. The 'Left Source Columns' table lists columns from the 'ORDERS' table, including ORDERNUMBER, PRODUCTNUMBER, ORDERDATE, STORECODE, PLANTCODE, SALESREP, QUANTITY, LINETOTAL, COSTOFGOODSSOLD, RETURNS, and WARRANTYEXP. The 'Right Source Columns' table lists columns from the 'GEN\_INVENTORY' table, including PRODUCTNUMBER, PRODUCTTYPE, PRODUCTCATEGORY, PRODUCTNAME, MODEL, QUANTITYINSTOCK, PRICE, and COST. The 'Join Type' is set to 'Inner Join', and the 'Join Conditions' section shows the expression 'ORDERS.PRODUCTNUMBER = GEN\_INVENTORY.PRODUCTNUMBER'. A Venn diagram on the left illustrates the inner join operation.

## The Power of Meta Data: Example

- We have a FIELD defined that has multiple sub-fields embedded
- Our report wants to show CREDIT RATING – based on the value of that field, we want CREDIT RATING to be set to “GOOD”, “BAD”, or “AVERAGE”
- What we need to do is pull out the 10<sup>th</sup> digit of that CUST\_CODE field and apply a rule.
- BUT WE DON'T WANT TO HAVE TO DO THIS FOR EVERY REPORT THAT INCLUDES CREDIT\_RATING



For credit rating >7 is “Good”,  
4-6 is “Average”,  
<4 is “Bad”

# The Power of Meta Data: Example

- Create a new field in meta data called **CREDIT\_RATING**
- Define the rule, and TEST

**Define Calculator**

Name: Credit\_Rating    Format: 7    Title:

Expression    Relational Expression

```

1 IF SUBSTR(10, STORESUBCODES, 10, 10, 1, 'A1') GT '7'
2 THEN 'GOOD'
3 ELSE IF SUBSTR(10, STORESUBCODES, 10, 10, 1, 'A1') GT '3'
4 THEN 'AVERAGE'
5 ELSE 'BAD'
    
```

Fields/Variables    Functi

Name

- f PATTERN(len
- f POSIT(source
- f RJUST(length
- f SOUNDEX(len
- f SPELLNM(out
- f SPELLNUM(ou
- f SQUEEZ(len
- f STRIP(len
- f STRREP(len
- f SUBSTR(len
- f TRIM(trim\_w
- f UPCASE(len

Character - DBCS

Character - Vari

Data Source and

Date - Legacy

Date - Standard

Date-Time

Format Conversi

Legacy

Numeric

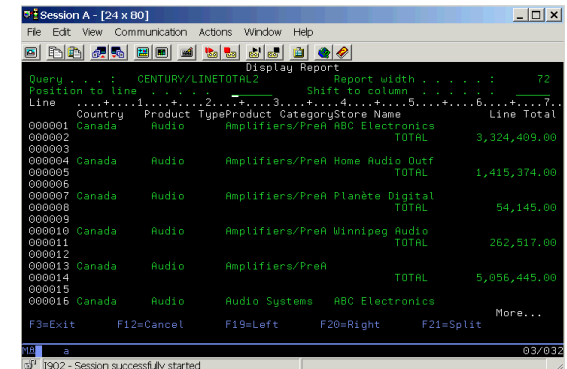
System

	Credit_Rating	STORESUBCODES
1	BAD	7512469350
2	AVERAGE	1913003716
3	GOOD	7475386198
4	AVERAGE	7714825036
5	BAD	3488300461
6	GOOD	4486836459
7	BAD	1916147250
8	AVERAGE	8697604804
9	GOOD	3157018379
10	AVERAGE	4472614927
11	GOOD	7709948048
12	BAD	3359468361
13	GOOD	5804993948
14	AVERAGE	3379383827
15	AVERAGE	1150593726
16	BAD	2836049271
17	AVERAGE	7577415605
18	AVERAGE	2614825015
19	AVERAGE	6591592014
20	AVERAGE	7993792105
21	AVERAGE	7683604814
22	GOOD	3159388148
23	AVERAGE	0713735837
24	BAD	4363389371
25	BAD	0359271593

## IBM i client issue: Query/400 reporting nightmares!

**Many i clients have hundreds if not thousands of Query/400 definitions that have proliferated over time**

- Redundancy of reports is commonplace
- Maintenance of reports is a nightmare
- Dependency on report “authors” is problematic
- Query/400 CANNOT use the latest DB2 for i Query processing technologies
- Extraneous steps taken for simple problems (getting data to a spreadsheet, or doing advanced data calculations)



```

Session A - [24 x 80]
File Edit View Communication Actions Window Help
Query . . . : CENTURY/LINETOTAL Report width . . . : 72
Position to line . . . . . Shift to column . . . . .
Line . . . . . 1 . . . . . 2 . . . . . 3 . . . . . 4 . . . . . 5 . . . . . 6 . . . . . 7 . . . . .
Country Product Type Product Category Store Name Line Total
000001 Canada Audio Amplifiers/PreA ABC Electronics TOTAL 3,324,409.00
000002
000003
000004 Canada Audio Amplifiers/PreA Home Audio Outf TOTAL 1,415,374.00
000005
000006
000007 Canada Audio Amplifiers/PreA Planète Digital TOTAL 54,145.00
000008
000009
000010 Canada Audio Amplifiers/PreA Winnipeg Audio TOTAL 262,517.00
000011
000012
000013 Canada Audio Amplifiers/PreA TOTAL 5,056,445.00
000014
000015
000016 Canada Audio Audio Systems ABC Electronics More...
F3=Exit F12=Cancel F19=Left F20=Right F21=Split
03/032
i302 - Session successfully started
  
```



# Why Modernize

- Turn DATA into INTELLIGENCE
- Delivered at the right time, in the right formats
- Take advantage of SQL Capabilities
- Document the database, and standardize on data meaning (single version of the truth)
- Get I/T out of the Report Writing Business
  - With self service reporting solutions
- Eliminate redundancy
- Adapt more quickly to changes in the operational systems/databases
- Deploy reports in a variety of highly intuitive ways
  - Mobile, Excel, Dashboards, OLAP
- Improve the perception of the “AS/400”

Parameters

Country: France, Germany, Spain, United States  
 Product Type: No Selection, Audio, Camcorders, Cameras

Run, Reset, Clear Output, Run in a new window

**Revenue by Country and Product Type**

Product Type	Country		
	France	Spain	TOTAL
Audio	16817832.00	16458158.00	33275990.00
Camcorders	19806592.00	21066201.00	40872793.00
Cameras			
<b>TOTAL</b>			

Data at 19/06/2012

**3e - PDF Revenue Summary by Product Category**

Product Category	Product Type	Revenue	Gross_Profit
Amplifiers/PreAmps/Tuners	Audio	\$42,374,428.00	\$16,634,858
Audio Systems	Audio	\$122,345,680.00	\$40,062,860
CD Players and Recorders	Audio	\$53,847,459.00	\$16,008,999
Digital Cameras	Cameras	\$184,103,667.00	\$50,774,837
Digital Camcorders	Camcorders	\$13,614,953.00	\$7,102,353
DVD	Video	\$329,872,045.00	\$81,103,145
DVD Camcorders	Camcorders	\$379,376,637.00	\$79,003,287
Handheld and PDA	Office	\$18,533,190.00	\$4,465,770
MiniDV Camcorders	Camcorders	\$51,539,451.00	\$17,411,091
MP3	Audio	\$43,491,588.00	\$17,052,928
Organizers	Office	\$11,712,495.00	\$6,755,190
Receivers	Audio	\$35,907,113.00	\$12,909,113
Speakers	Audio	\$84,717,053.00	\$60,036,063
TV	Video	\$168,799,539.00	\$18,027,839
VCR	Video	\$21,688,621.00	\$5,417,671
<b>TOTAL</b>		<b>\$1,561,923,919.00</b>	<b>\$432,766,004</b>

**Product Sales Information Dashboard**

Average Margin by Period, On-Time Delivery overall, Order backlog by Period, On-Time Delivery by Product Type

Revenue and Returns by Type

Product Type	2011		2012	
	Returns	Revenue	Returns	Revenue
Audio	\$ 41,351	\$ 190,920,684.00	\$ 46,604	\$ 191,762,637.00
Camcorders	16,748	216,517,006.00	20,366	228,004,039.00
Cameras	18,624	90,791,413.00	20,366	93,362,254.00
Office	12,546	14,658,888.00	14,153	15,568,797.00
Video	19,701	623,444,717.00	22,189	726,919,403.00
<b>TOTAL</b>	<b>\$109,570</b>	<b>\$746,302,698.00</b>	<b>\$123,668</b>	<b>\$776,628,131.00</b>

Revenue/Margin breakout by Product Type

Revenue: Audio, Camcorders, Cameras, Office, Video  
 Margin: Audio, Camcorders, Cameras, Office, Video

The Greatest BI Tool in the world is useless if it doesn't perform !

“The most widespread technical problem reported by practitioners was slow query performance.”



- **Survey of over 2000 companies that have implemented Business Intelligence Applications**
  - The BI Survey 8 – Nigel Pendse

## Infrastructure Matters....

- **Scalability**

- The average data warehouse doubles in size every 18 months

- **Optimized for Workload**

- Complex Querying and Aggregation is very different from OLTP workloads and requires different tuning

- **Resilient**

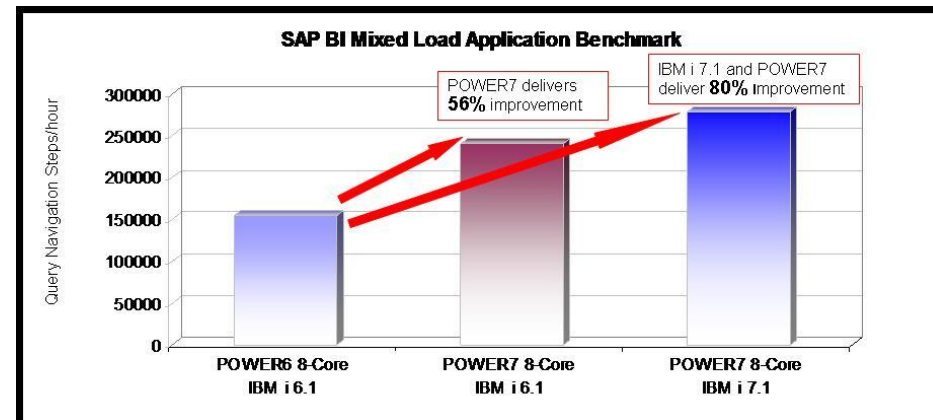
- Analytics are as important in terms of availability as line of business applications

- **Minimize Risk**

- Leverage what you do best to minimize risk
- Security policies; HA/DR infrastructure; Operational policies and procedures

## DB2 for i Enablers for Analytics – Built into DB2 for i

- **SQL Query Engine (SQE)**
  - Real Time Statistics
  - Automated Index Creation and Query Adjustment
- **Database Parallelism**
  - Leverage Multiple Cores
- **Materialized Query Tables**
  - Aggregate Processing
- **Encoded Vector Indexing**
  - IBM Patented Query Acceleration Technology
- **Management Tools – System i Navigator**
  - SQL Plan Cache, Monitor Analysis of Queries
  - Index Advice



\*See detailed certified benchmark results at

<http://www.sap.com/solutions/benchmark/bimxl.epx>

# DB2 Symmetric Multiprocessing (feature of IBM i )

## ■ SELECTING

- Index scan or probe
- Table scan or probe via bitmap or RRN list
- Table scan

## ■ JOINING

- Index scan or probe
- Hash


## ■ GROUPING

- Index scan or probe
- Hash

## ■ ORDERING

- Index scan or probe
- Sort

- Creating temporary indexes for joining, grouping or ordering is SMP enabled



# SQL Query Engine

- Advanced query optimization, query execution engine, and management tools

- Part of DB2 since V5R3
  - Enhanced with each subsequent release
- Leverages more DB2 performance and management facilities
  - Encoded Vector Indexes, Materialized Query Tables
  - SQL Plan Cache (part of IBM i Navigator)



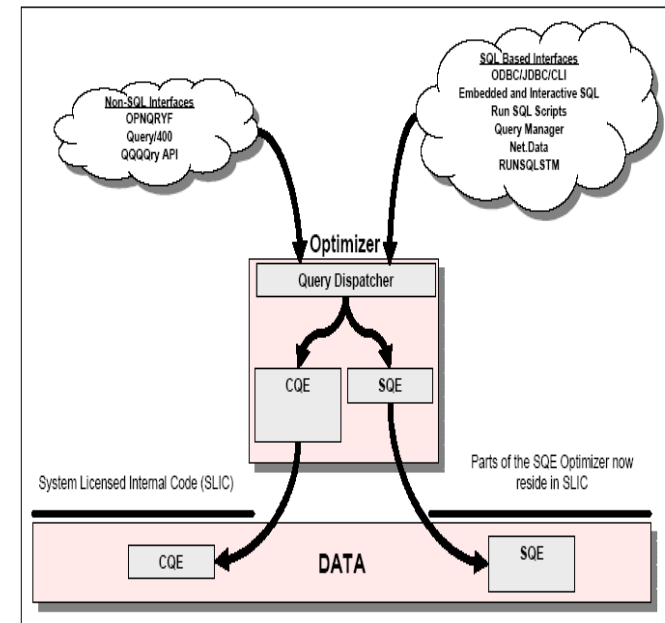
- 7.1 Enhancements

- Adaptive Query Processing (AQP)
- Support for Logical Files

- NOTE: CQE, or “Classic” Query Engine, is also part of DB2 to support non SQL Standard Interfaces for accessing DB2**

- Query/400
- Some ISV Applications
- Limited ability to leverage


**Power is performance redefined**





# Materialized Query Tables (MQT)

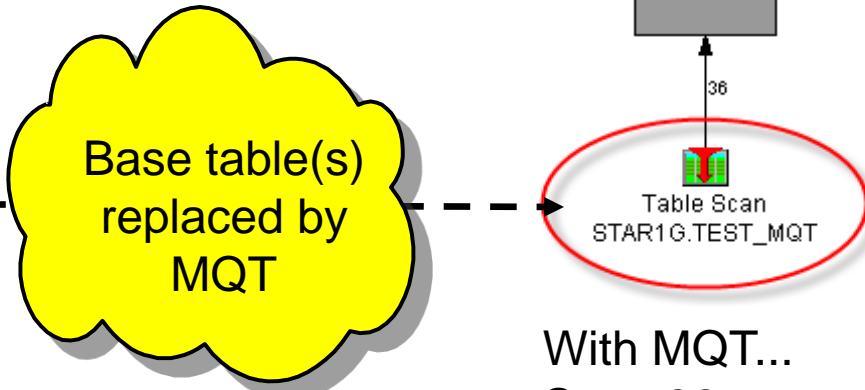
Before...



Without MQT...  
Scan and aggregate  
6,000,000 rows

```
SELECT year, quarter, month,
       SUM(revenue_w_tax) AS srevenue_w_tax,
       SUM(revenue_wo_tax) AS srevenue_wo_tax,
       SUM(profit_w_tax) AS sprofit_w_tax,
       SUM(profit_wo_tax) AS sprofit_wo_tax,
       SUM(quantity) AS squantity,
       COUNT(*) as number_items_per_group
FROM   ITEM_FACT
GROUP BY year, quarter, month;
```

After...



With MQT...  
Scan 36 rows

Highlight  
MQT  
In V5R4

V6R1: Improved Refresh Performance!

# Encoded Vector Index (EVI) Aggregates (7.1)

Symbol Table				
Key Value	Code	Count	Include Sum()	Include Sum()
Arizona	1	5000	1500	2005
Arkansas	2	7300	3200	450
...				
Wisconsin	49	340	575	1200
Wyoming	50	2760	210	0

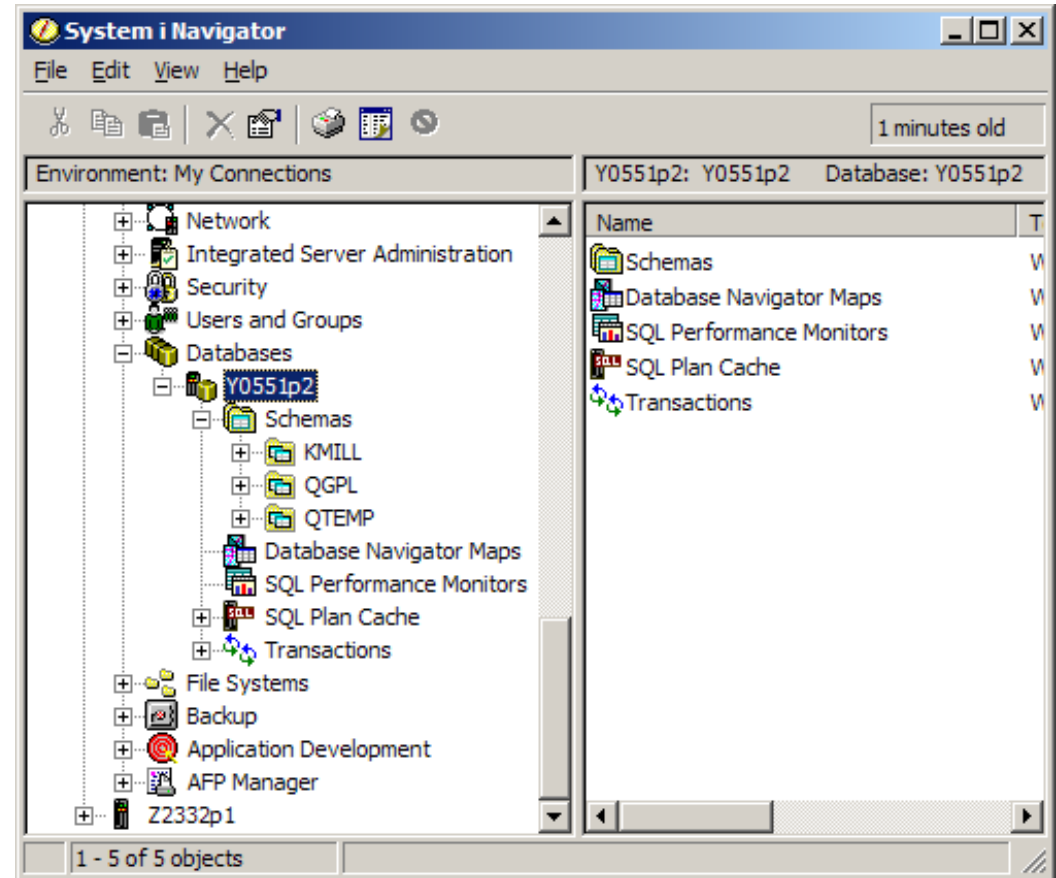
Optional (7.1)

Vector	RRN
1	1
17	2
5	3
9	4
2	5
7	6
50	7
49	8
5	9
...	...


- Symbol table contains information for each distinct key value
  - Each key value is assigned a unique code (key compression)
  - Code is 1, 2, or 4 bytes depending on number of distinct key values
- Rather than a bit array for each distinct key value, use one array of codes

# Establish Best Practices for Managing Query Performance

- **DB2 for i Administration**
  - Part of i Navigator
  - Performance Analysis Tools
    - Database Monitors
    - SQL Plan Cache
    - Visual Explain
    - Index Advice
    - Index Evaluation
    - Reporting Functions
      - Save to spreadsheet
      - Filters
  - Documentation Tools
    - Database Navigator



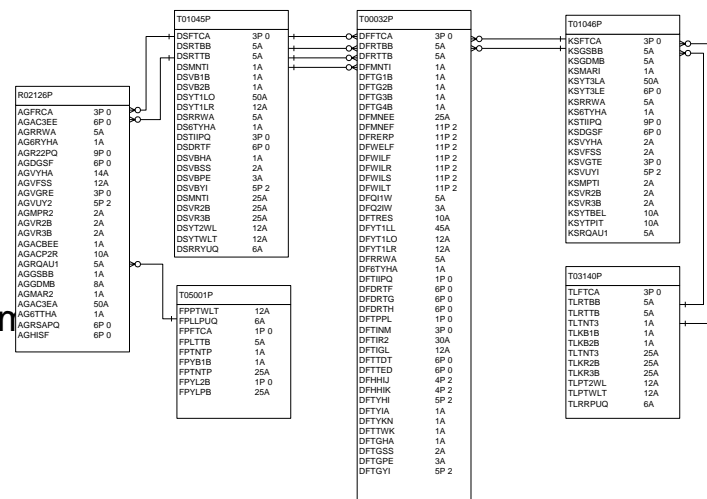
# What is the right INFRASTRUCTURE to support IBM i BI?

Operational Reporting	<i>Operational Data Store</i>	Data Warehouse
<ul style="list-style-type: none"> <li>• Deployed On Existing i server</li> <li>• Simple reporting and BI</li> <li>• No data replication or transformations</li> <li>• Mixed workloads (OLTP and heavy query)</li> </ul>  <p>DB2 Web Query</p>	<ul style="list-style-type: none"> <li>• Optimized environment for operational reporting</li> <li>• Separate BI from OLTP workload</li> <li>• Multi-purpose 2<sup>nd</sup> System</li> <li>• Simple, low cost data replication</li> <li>• <u>Foundation</u> for data warehouse</li> </ul> <p>IBM i for Business Intelligence</p>	<ul style="list-style-type: none"> <li>• Extend value of IBM i for Business Intelligence to Data Warehouse</li> <li>• Add an ETL tool for data transformation</li> <li>• CDC (transport) and an ETL (transform) can provide near real time analytics</li> <li>• Fully leverage advanced DB2 i technology</li> </ul> <p>IBM i for Business Intelligence + ETL and Services</p>
<p><b>Mixed Workload</b> <i>Diminished Efficiencies</i></p>	<p><b>Workload Optimized</b> <i>Simplify - Accelerate Value - Reduce Cost</i></p>	
<p>DB2 Web Query Standard Edition Getting Started Services Memory and CPW !</p>	<p>IBM i for BI Solution Additional IBM i Licenses for other uses</p>	<p>IBM i for BI Solution Extract/Transform/Load (ETL) Tool Additional Services</p>



# Option 1: Querying Production Databases

- **Isolate workloads as best you can**
  - Create separate subsystem for query jobs
    - DB2 Web Query installs into its own subsystem
  - Keep memory pool LARGE with reduced number of active jobs
- **Indexing and aggregate strategy plays a major role in query performance**
  - MQTs, EVIs, and traditional indexes
    - Read white papers on MQTs and Index and Statistics Advice
      - <http://www-03.ibm.com/servers/enable/site/bi/strategy/index.html>
- **Leverage Query Governors**
  - QQRYSIMLMT
- **Consider Report Delivery**
  - Do you need “real time” vs. batch?
  - Control queries with parameterized reports
  - Use DB2 Web Query “Active” technologies
    - Once the report is run you are no longer impacting system
- **Let DB2 do as much work as possible**
  - DB2 Web Query tips and techniques (e.g., joins)



## Option 2: Isolation through an *Operational Data Store*

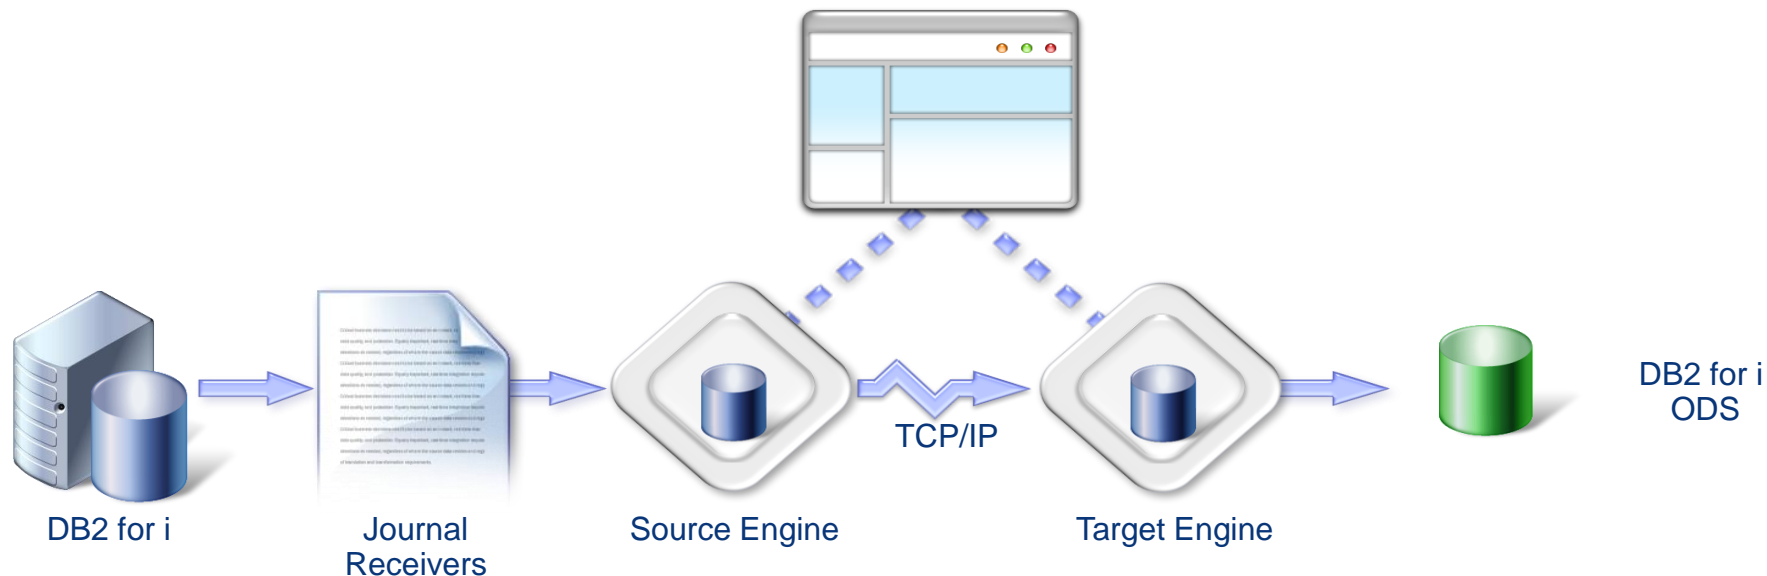
- **What is an Operational Data Store (ODS)?**
  - A COPY of the operational (transaction) data base most often used for reporting purposes
  - ISOLATED from production workloads
  - Kept up to date based on requirements
    - Near real time?
  - The data model stays pretty much the same as production, with possibly some minor changes
    - Could have subset of fields/ columns
    - Could contain more historical data than production systems
    - Could contain minimal data transformations
    - Could encrypt/mask certain data elements
  - Enhanced for performance
    - Aggregations of the data
    - Indexing
  - Security model might be different
  - Platform is TUNED for reporting purposes
  - Get to IBM i 7.1 and POWER7 for this workload

## Example: Populating the ODS with IBM's Infosphere CDC

### ▪ Changed Data Capture (CDC) Replication Software


- Based on journaling
  - Remote or Local
- Requires both before and after images to be store in journal receiver
- Techniques to do INITIAL loads with sync points
- Java GUI to set up and monitor (no programming required)
- Minimal transformations

Monitoring and Configuration




# New IBM i for Business Intelligence Packaged Solution

Production System



Extract, Transport  
and Load Data  
using **InfoSphere  
Software**



IBM i for BI



## IBM i for Business Intelligence

- Installation Services Included in Price
- Includes DB2 Web Query Standard Edition
- Replicate production database to 2<sup>nd</sup> system
- Isolate query workloads (tune, optimize)
- LOW Entry price point that you can build on
- Offered in 3 Sizes: Small/Medium/Large

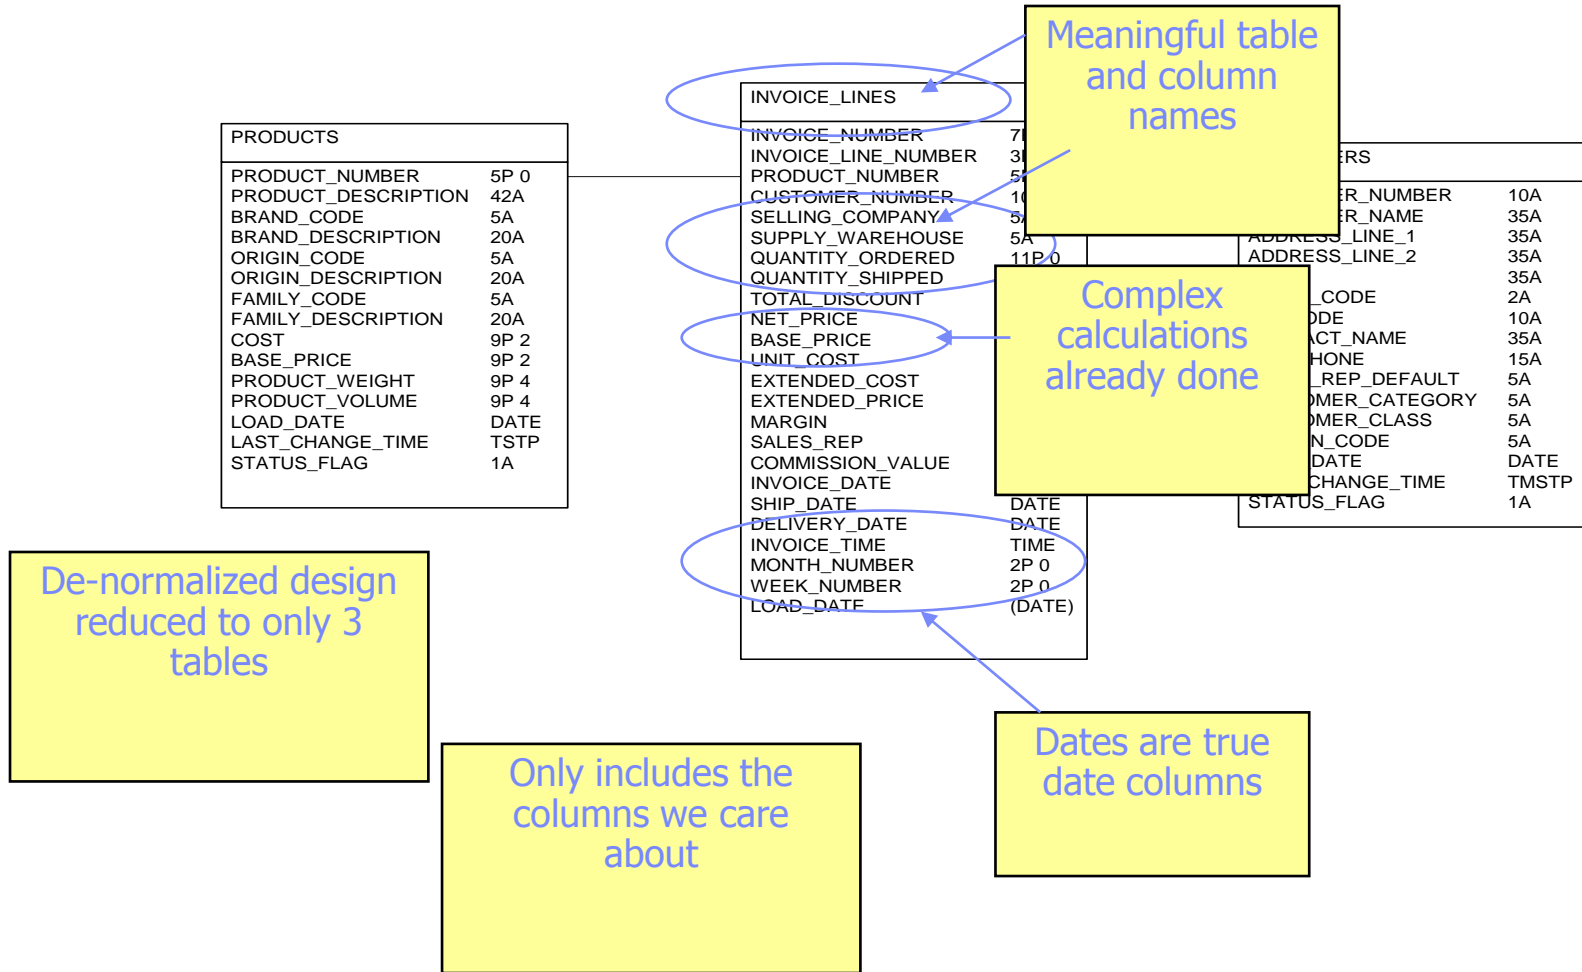


## Option 3: The Data Warehouse

- Adds Extract, Transform, and Load



# The Data Warehouse Model



Data Warehouse Schema

# Transformation Example: Surrogate Keys

Customer File - US	
CUSTNO	CUSTNAME
1001	John Smith
1002	Mary Jones
1003	Chris Anderson
1004	David Perry

Customer File - Canada	
CUSTNO	CUSTNAME
1001	Harry Potter
1002	Jeremy Carr
1003	Penny Hayes
1004	Debbie Thornton



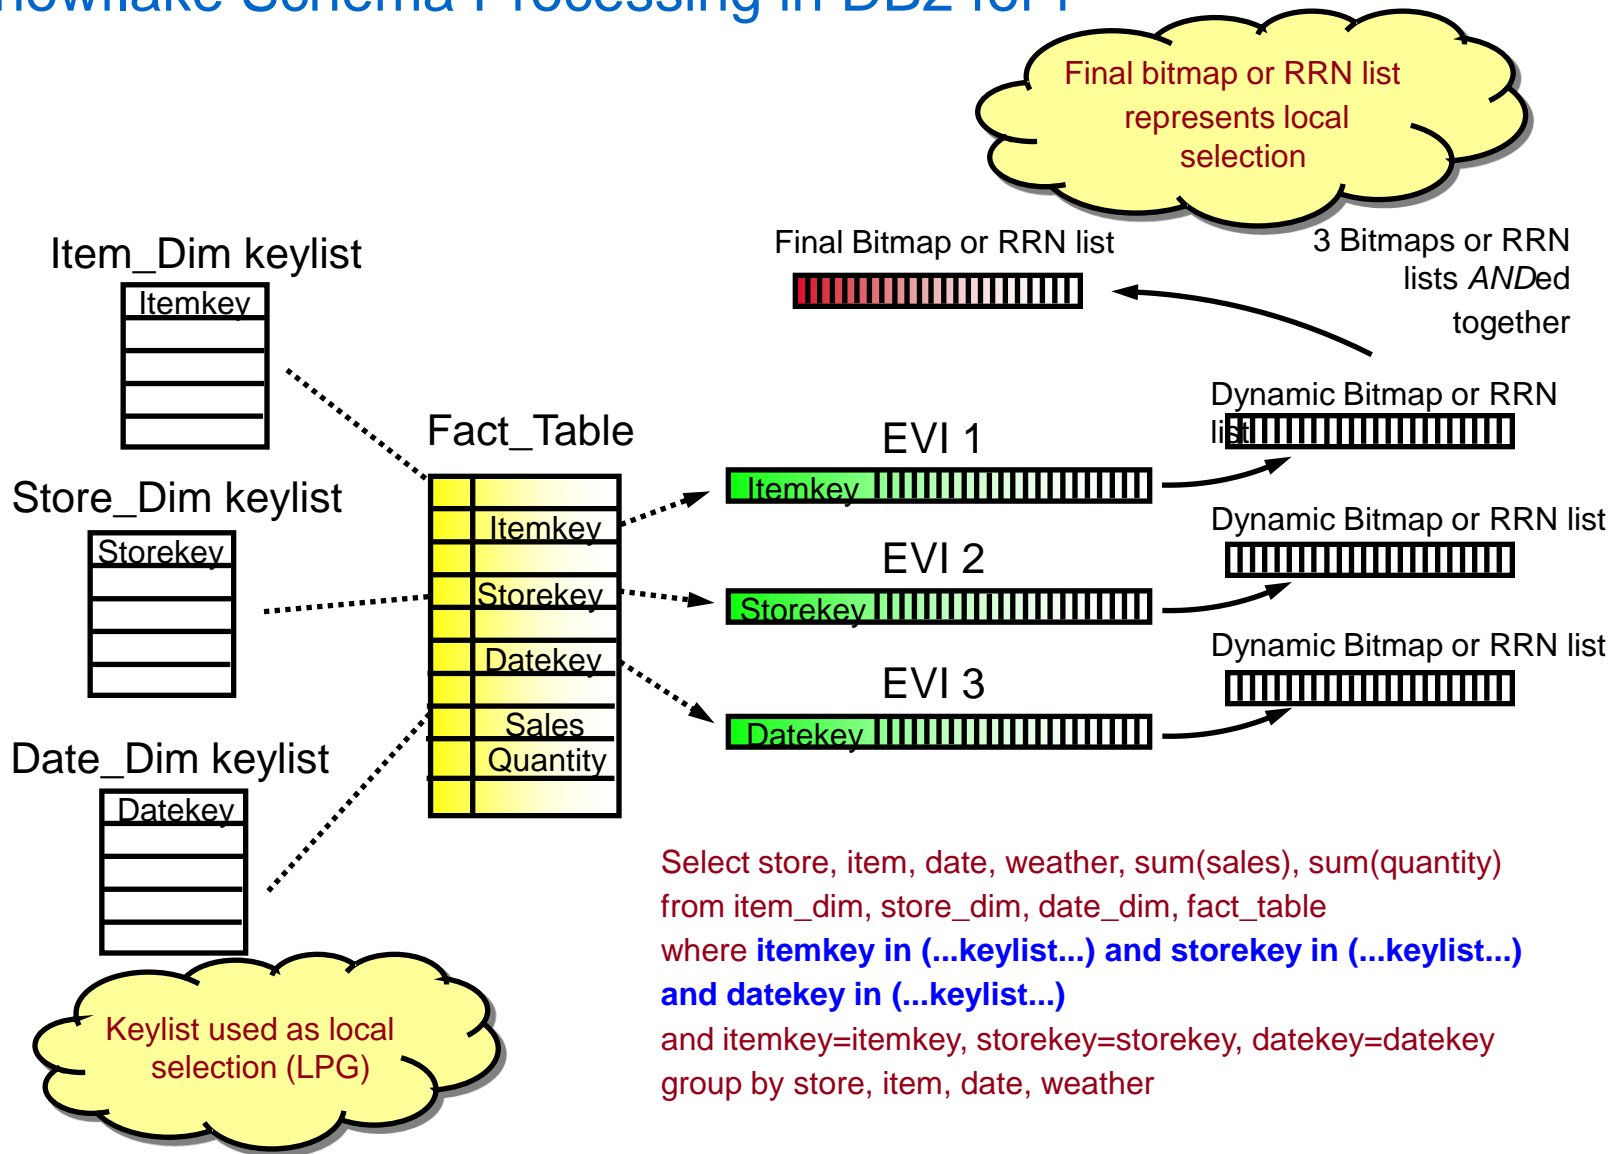
Surrogate key is a sequential number with no correlation to replaced value(s)

Customer File - Data Warehouse			
CUSTNUMBER	CUSTNAME	REGION	OLDNUM
1	John Smith	US	1001
2	Mary Jones	US	1002
3	Chris Anderson	US	1003
4	David Perry	US	1004
5	Harry Potter	CANADA	1001
6	Jeremy Carr	CANADA	1002
7	Penny Hayes	CANADA	1003
8	Debbie Thornton	CANADA	1004

PK

Secondary Index

# Star / Snowflake Schema Processing in DB2 for i



## ETL Alternatives


- **Do it yourself**
  - Custom coding
  - Consider use of SQL vs. RPG record level processing
  
- **IBM i based (DB2 Web Query Meta Data Integration)**
  - Information Builder's Data Migrator
    - [www.ibi.com](http://www.ibi.com)
  - Coglin Mill's Rodin DB2 Web Query Edition
    - [www.thinkrodin.com](http://www.thinkrodin.com)
  
- **High End (AIX LPAR)**
  - IBM InfoSphere Data Stage
    - Strong source and target support
    - Parallelism built into the load processes
    - Many data transformations built in

# Expanding the IBM i for BI into Data Warehousing




Production System

Purchase additional  
ETL Tools/Services

IBM i for BI



Extract and Transport Data using **InfoSphere Software**



Transform and Load Data into DB2 Data Warehouse

## IBM i for Business Intelligence

- ETL process TRANSFORMS and cleanses data
  - Bulk load or refresh
  - Scheduled or continuous updates
- Restructuring of the data improves analytics  
(for example: create a customer profiling database)
- Isolate query workloads (tune, optimize)

## Lastly, a word about Cognos as it relates to IBM i

- **All of the previously mentioned DB2 for i query optimization issues apply !**
- **Cognos does not run in IBM i, however, can access DB2 for i**
- **Cognos CAN run in an Linux on Power or AIX Partition**
  - For BEST PRACTICES for running Cognos in an AIX partition, refer to:
    - [http://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/whitepaper/aix/v6r1\\_cognos/methods](http://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/whitepaper/aix/v6r1_cognos/methods)
  - Exploiting PowerVM in a Cognos environment
    - <http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg247842.html?Open>

# Questions & Answers

