



End-to-end
Information on Demand

InfoSphereTM software

Trusted Information



Information Management software

IBM Balanced Warehouse

Henry Cook
Henry_Cook@uk.ibm.com
+44 7802 931 964

IBM Balanced Warehouse speeds time to value by giving you everything you need in one easy to install package

1. **Everything you need for warehousing**

- Data Modelling
- Data Flow / Manipulation Tools
- Easy Data Mining
- Text analytics
- Cubing Services

2. **Easy to install and administer**

- Pre-configured, pre-tested
- 'Query Ready'
- Low TCO

3. **High Performance and Scalable**

- Full implementation of 'Shared Nothing' Massively Parallel Processing (MPP)
- Mixed Workload Support

4. **Affordable**

- Low Cost / Tb
- Extreme Compression
- Predictable costs

5. **Proven**

- Highly Robust and Reliable

The IBM Balanced Warehouse is part of IBM's Information On Demand Strategy, It provides the middle layer of and end to end capability

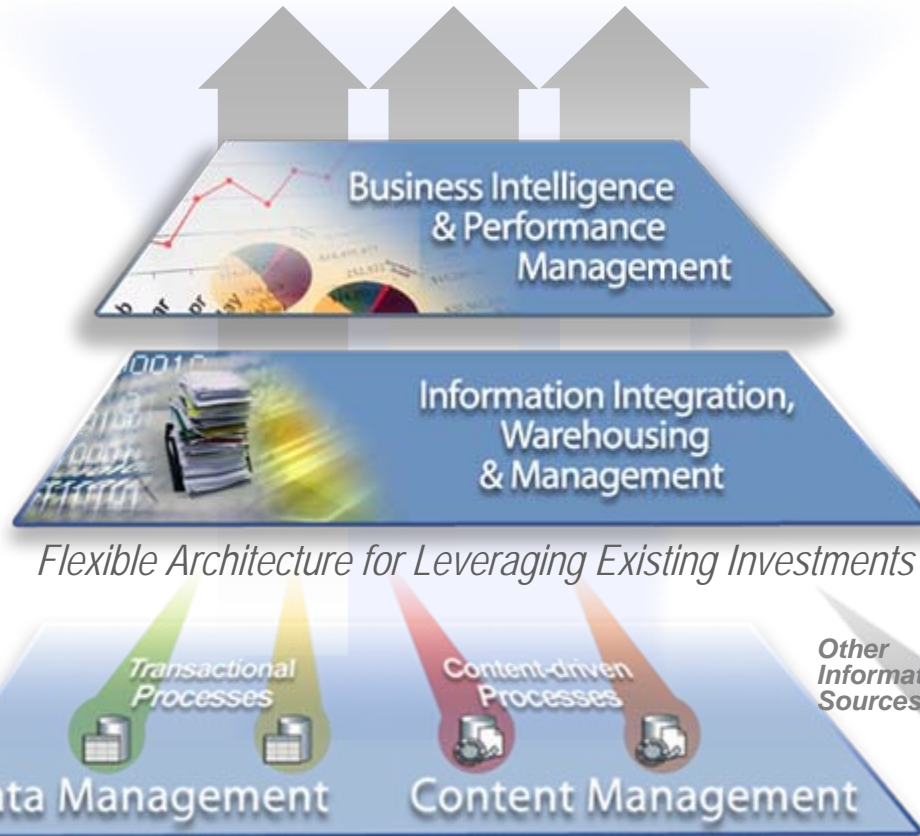
Customer & Product Profitability *Financial Risk Insight* *Workforce Optimization* *Dynamic Supply Chain* *Multi-Channel Marketing*

Business Optimization

Better Business Outcomes



End-to-End Capabilities



Plan, understand and optimize business performance

Establish and maintain an accurate, trusted view of information


Manage data and content over its lifetime and use as part of processes

The Results Are In!

“IBM Offers The EDW Market’s Most Comprehensive Appliance Product Family”

InfoSphere Warehouse ranked high in the leadership pack based on our current offering, strategy, & market presence

- ✓ *“IBM Offers The EDW Market’s Most Comprehensive Appliance Product Family”*
- ✓ *“IBM’s InfoSphere Balanced Warehouse product family offers best-of-breed functionality”*
- ✓ *“IBM has the broadest range of EDW appliance packaging options”*
- ✓ *“By continuing to expand, deepen, and integrate its Information On Demand (IOD) portfolio, IBM keeps its EDW offerings and strategy in the industry forefront”*
- ✓ *“Strong financials and a steady stream of customer wins demonstrate IBM’s continued market leadership”*



The Forrester Wave is copyrighted by Forrester Research, Inc. Forrester and Forrester Wave are trademarks of Forrester Research, Inc. The Forrester Wave is a graphical representation of Forrester's call on a market and is plotted using a detailed spreadsheet with exposed scores, weightings, and comments. Forrester does not endorse any vendor, product, or service depicted in the Forrester Wave. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change.

IBM Experience with Data Warehousing

Ranked as a leader in Gartner 2009 Data Warehouse Magic Quadrant, and for last 10 years

11 of the top 12 banks
 7 of the top 8 auto manufacturers
 5 of the top 6 insurance companies
 4 of the top 6 general merchandisers
 4 of the top 5 specialty retailers
 3 of the top 4 food & drug stores



Example: Gamesys

- Coping with High Growth
- Ease of Install
- Low Administration
- Performance and Scalability
- High Availability
- Faster provision of information
- Customer Analysis

Information Management



Gamesys finds a safe bet with its balanced warehouse solution from IBM

Overview

■ Business challenge

Despite its rapid growth, in an increasingly competitive market, Gamesys needed to work harder at retaining its existing customers. This required deeper analysis of reliable business information but the company's analytics and reporting solution was struggling to cope.

■ Solution

The existing data platform was replaced by an IBM InfoSphere™ Balanced Warehouse solution using IBM's Retail Data Warehouse Industry Model. Reporting is provided through Cognos®. The new solution is already delivering quick and easy access to reliable business information and will enable simplified analysis of customers across geographic locations.

■ Key benefits

- Detailed customer information enables improved product offerings and marketing
- Business reports available within 30 minutes rather than 24 hours
- Minimised exposure to fraud due to realtime monitoring of gaming data
- Increased speed to market for new products
- Scalable solution accommodates business growth
- Platform to maximise the value of business information.



Founded in 2001, the Gamesys Group has rapidly built a profitable online entertainment business providing Instant Win, Bingo and Casino games to a worldwide audience. The Group handled over £1 billion of cash bets in 2008 and is one of the fastest growing businesses in the e-gaming sector.

Gamesys was using a spreadsheet-based solution for its business analytics and reporting. Although initially perfectly adequate, the solution was struggling to cope with the increase in data generated by the company. Furthermore, as a young business, Gamesys is still developing its online functionality and this was inhibited because the business logic for reporting was embedded in the spreadsheets. Any change made to the main data model supporting live

operations would affect the reports coming from the spreadsheets, possibly causing reporting discrepancies.

With legislation adapting slowly to online gaming, Gamesys has to operate from multiple geographic locations to remain within the law. The Group's main data environment resides in Malta and business information was transferred overnight to a Management Information System in London. This delay in the availability of data posed a substantial business risk. Players spotting vulnerability in a newly launched game can exploit it to their own advantage, potentially exposing Gamesys to extremely large financial losses. Swift access to information was crucial to reduce the length of time that vulnerabilities go unnoticed.

The Infosphere Balanced Warehouse™

Real time, ready to go data warehousing solutions

*The **Infosphere Balanced Warehouse** is the complete data warehousing solution comprised of pre-tested, scalable and fully-integrated system components of DB2 Warehouse, Server and Storage.*



Balanced Warehouse

IBM Balanced Warehouses include **everything** required to serve as a foundation for your business intelligence solutions. Some of the included highlights are:


- Robust IBM InfoSphere Warehouse software
- The latest in IBM Server technology
- High performance IBM storage
- Comprehensive total solution support

Infosphere Balanced Warehouse

IBM's Data Warehousing Appliance



Balanced Warehouse



Simplicity

- Predefined configurations for reduced complexity
- One number to contact for complete solution support

Flexibility for growth


- Add modules to address increasing demands
- Multiple entry points for different needs <1Tb to 100's Tb
- Reliable, nonproprietary hardware for reusability

Optimized performance


- Preconfigured and certified for guaranteed performance
- Based on best practices for reduced risk
- Industry leading low TCO

Massively Parallel Processing (MPP) spreads the processing across many processing 'nodes' each is a powerful server

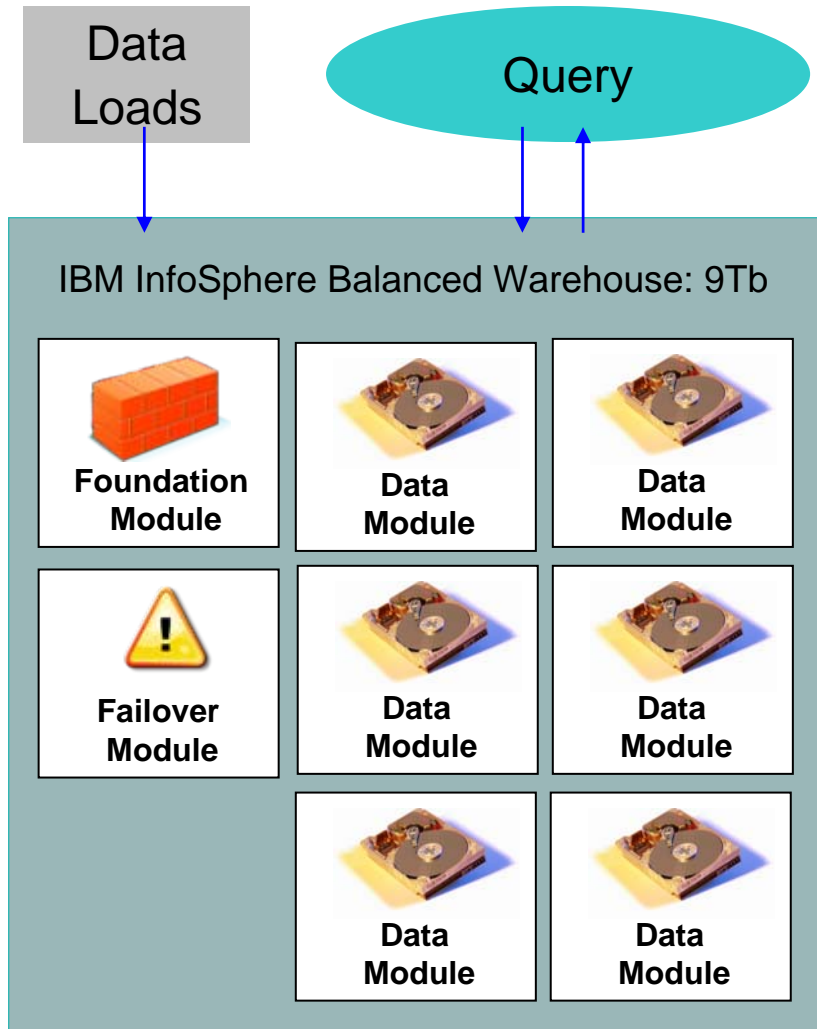
IBM[®] Balanced Warehouse™



The data for a query is split across all the nodes and a query runs in parallel across all the pieces

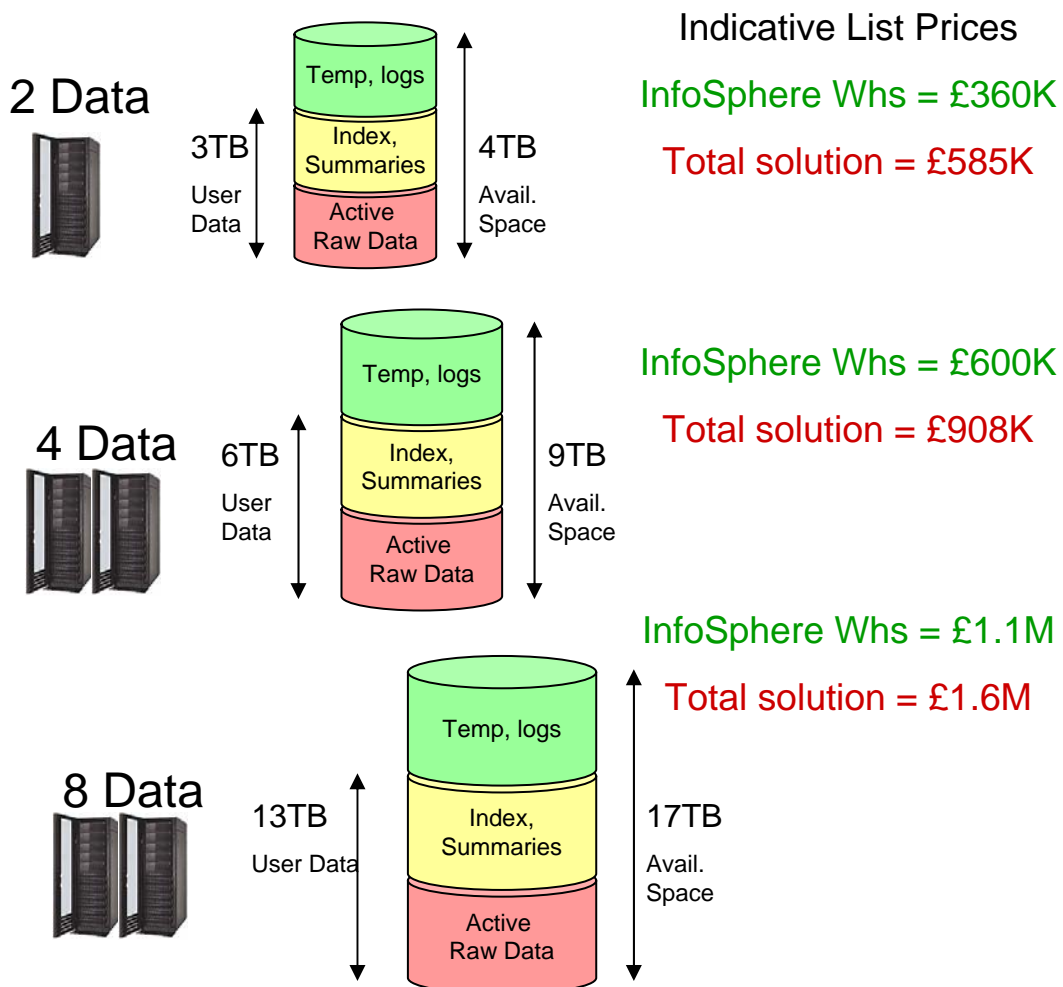
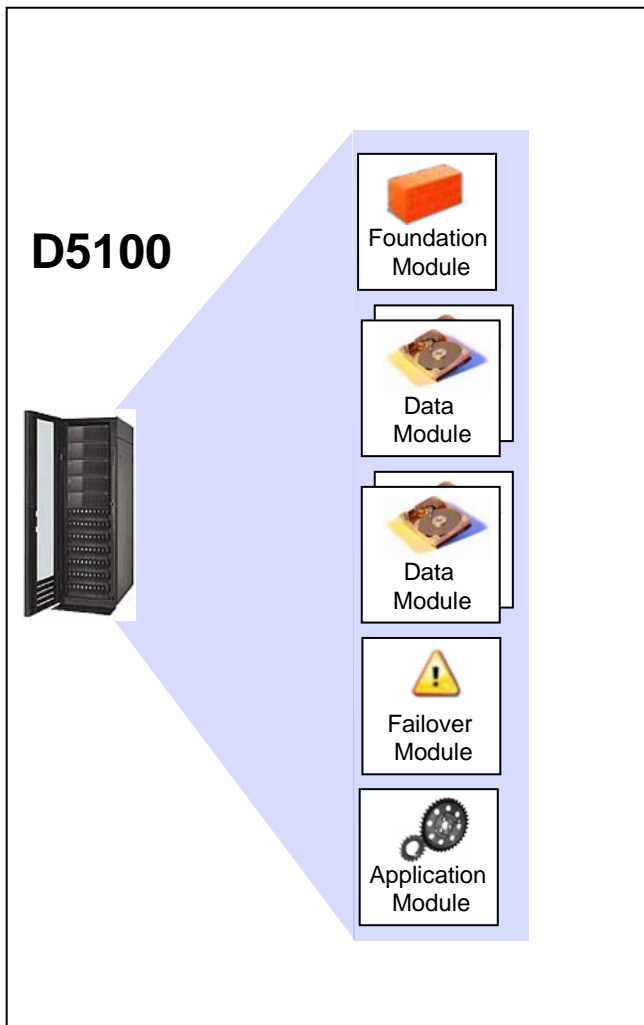


Example InfoSphere Balanced Warehouse™ configuration



- **D-Class**
6 x Data +
1 Foundation +
1 Failover
- **Data building block of**
1.5 Tb approx capacity
- **9 Tb capacity**
- **24 Data CPU cores**

For larger scaling the IBM Balanced Warehouse D-Class range is available, a sample of which are shown here



IBM Balanced Warehouse C-Class Preset offerings fully featured yet affordable appliances

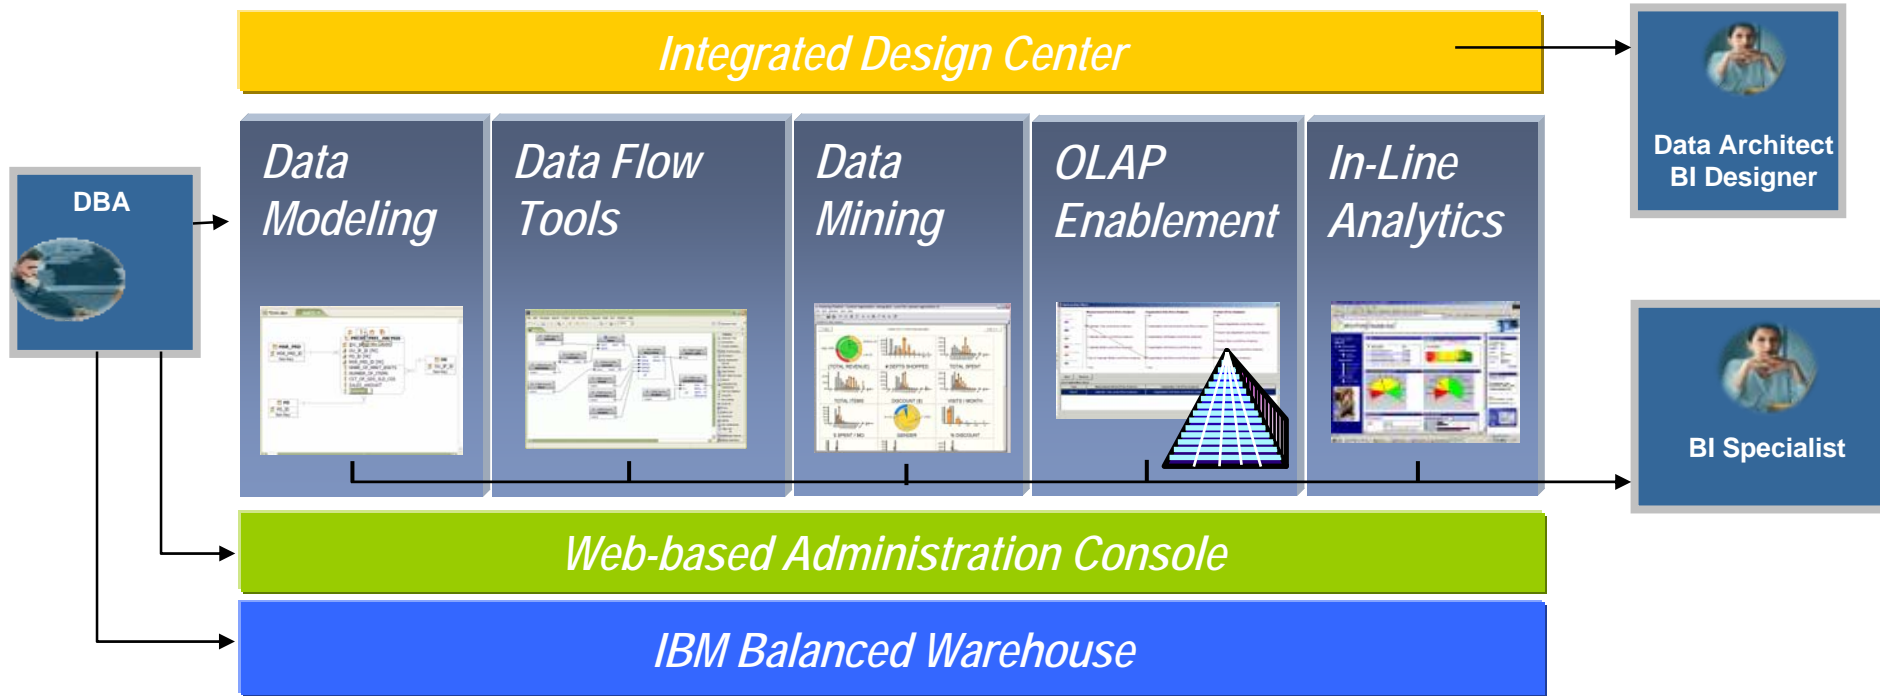
C Class	Small	Medium	Large	XL
Model	C3000 	3000 	C4000 	C4000
Cores	1 x Quad	2 x Quad	4 x Quad	4 x Quad
InfoSphere Warehouse	£4K-£26K	£4K-£52K	£262K	£262K
Hardware	£12K	£21.5K	£54K	£86K

List Prices shown for hardware / software

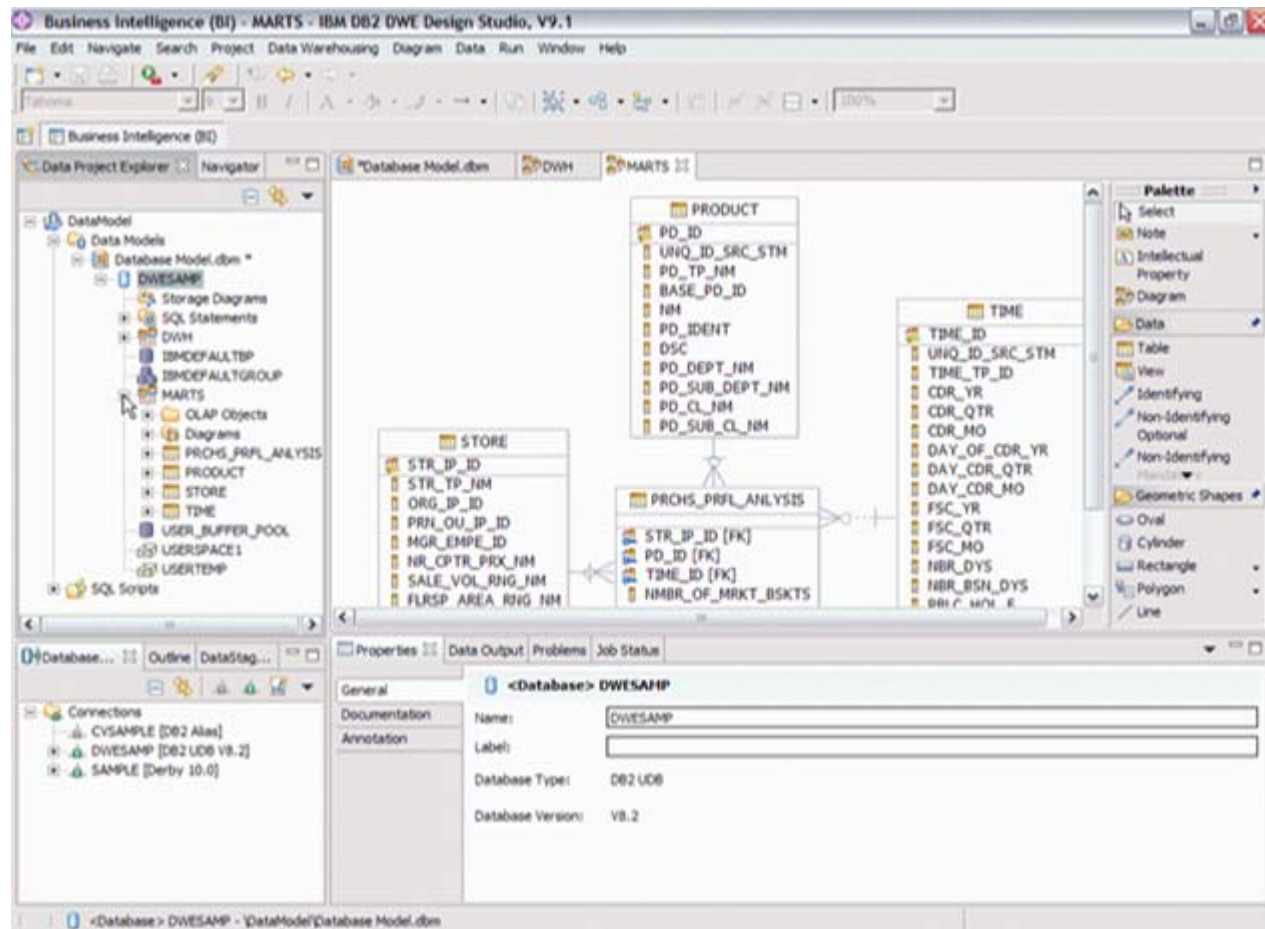
C3000 is rack and IW Intermediate edition (User or PVU), C4000 is Enterprise (PVU)

Build and Support not included

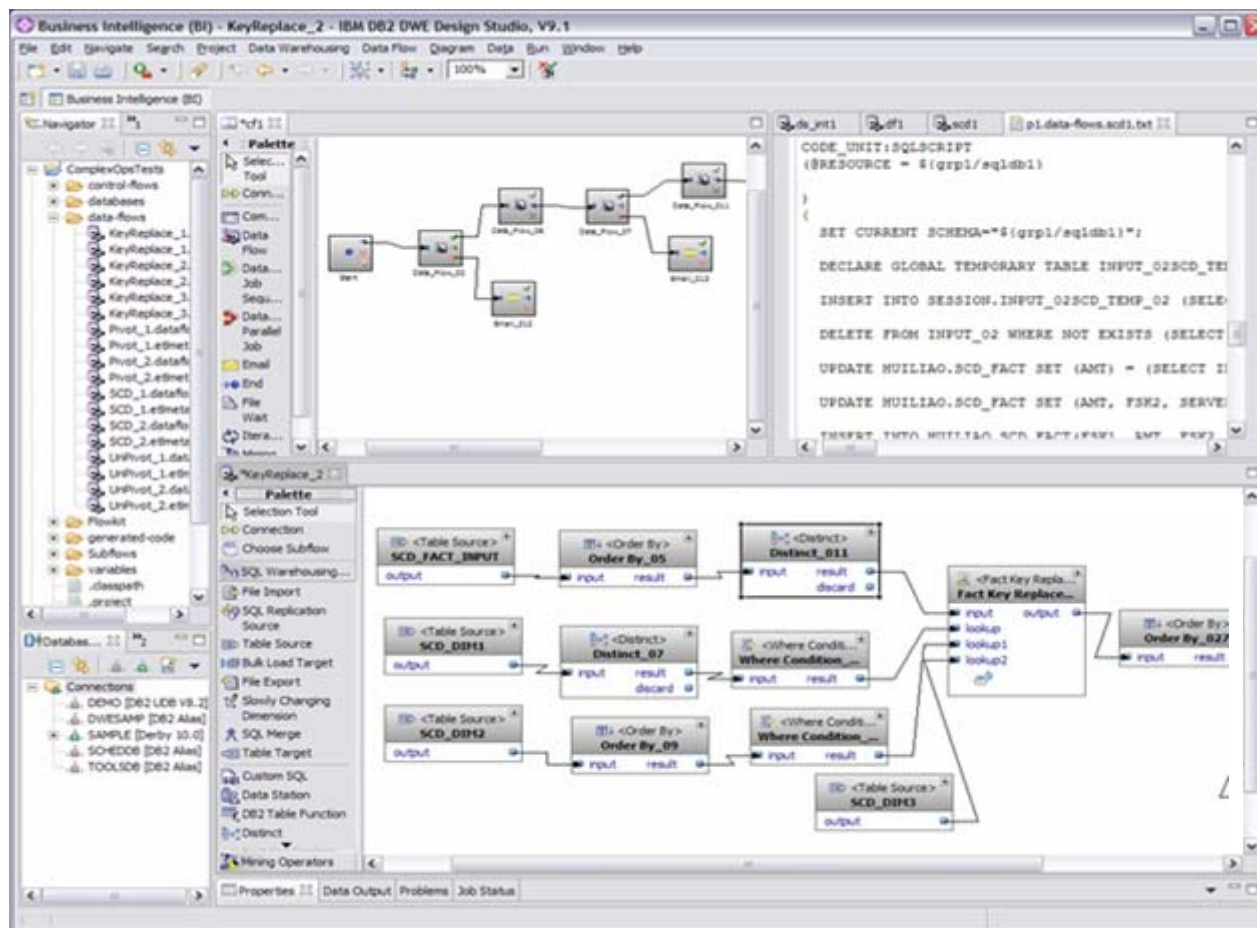
An Integrated Development Environment combines all of the tooling necessary to implement Data Warehouse projects, the Balanced Warehouse is not just a Database Engine



A Data Modelling Tool (ex Rational) is included for logical and physical modelling of your data

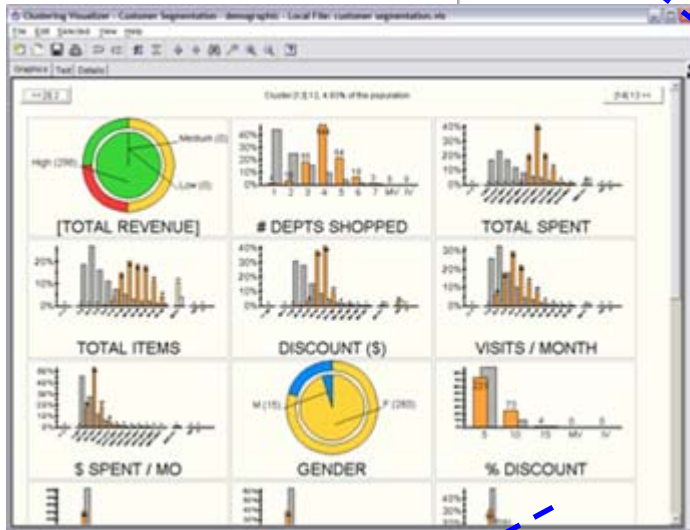


A simple 'data flow' tool is included it can be used directly or interfaced to the InfoSphere Data Stage tool



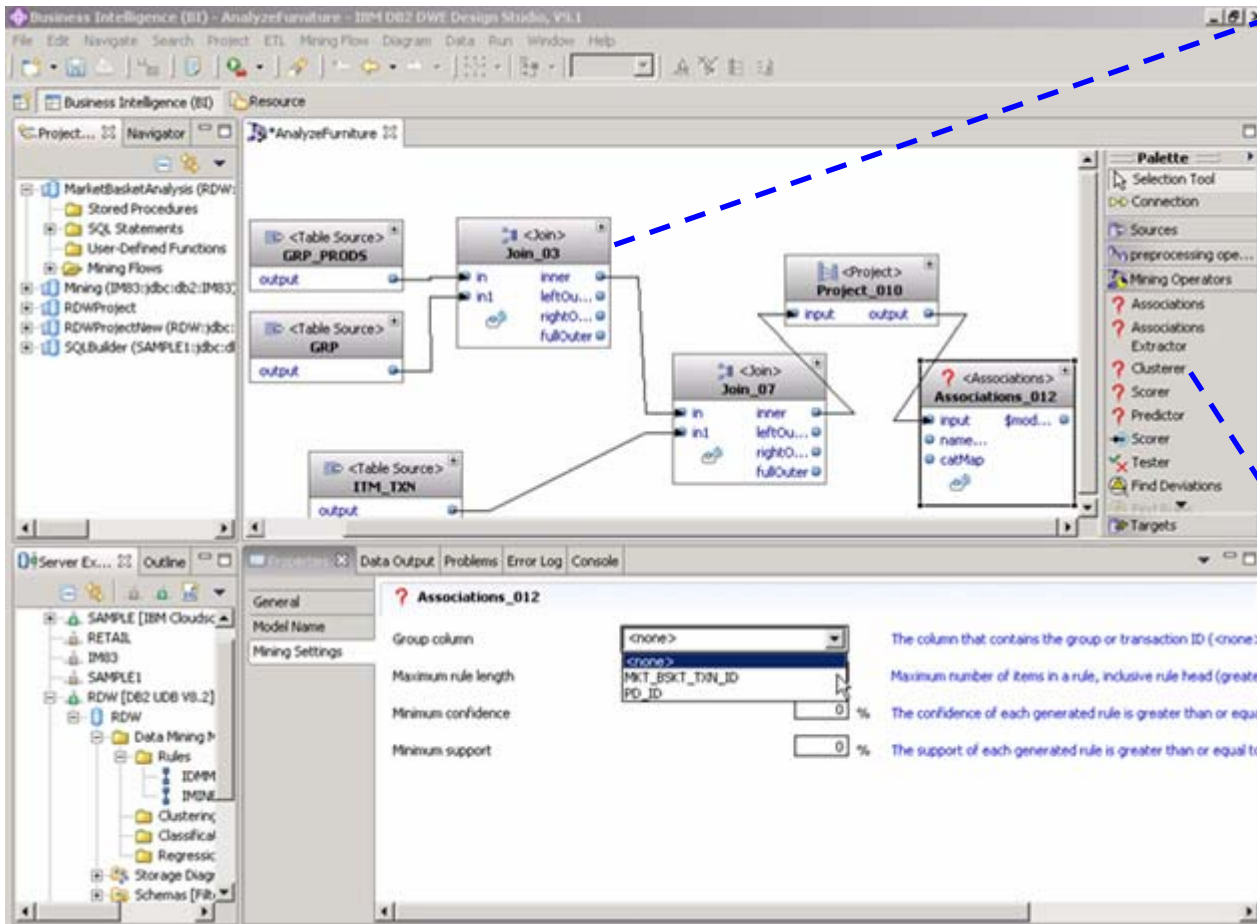
Data Mining visualisation tools are included making it easy to produce and publish sophisticated data mining results

Can be embedded directly into applications and Web pages



Out-of-the-box visualization tools

Integrated data mining for customer segmentation, market basket analysis etc is included. Mining can be invoked from query tools and sophisticated mining used without the need for heavyweight statistical skills

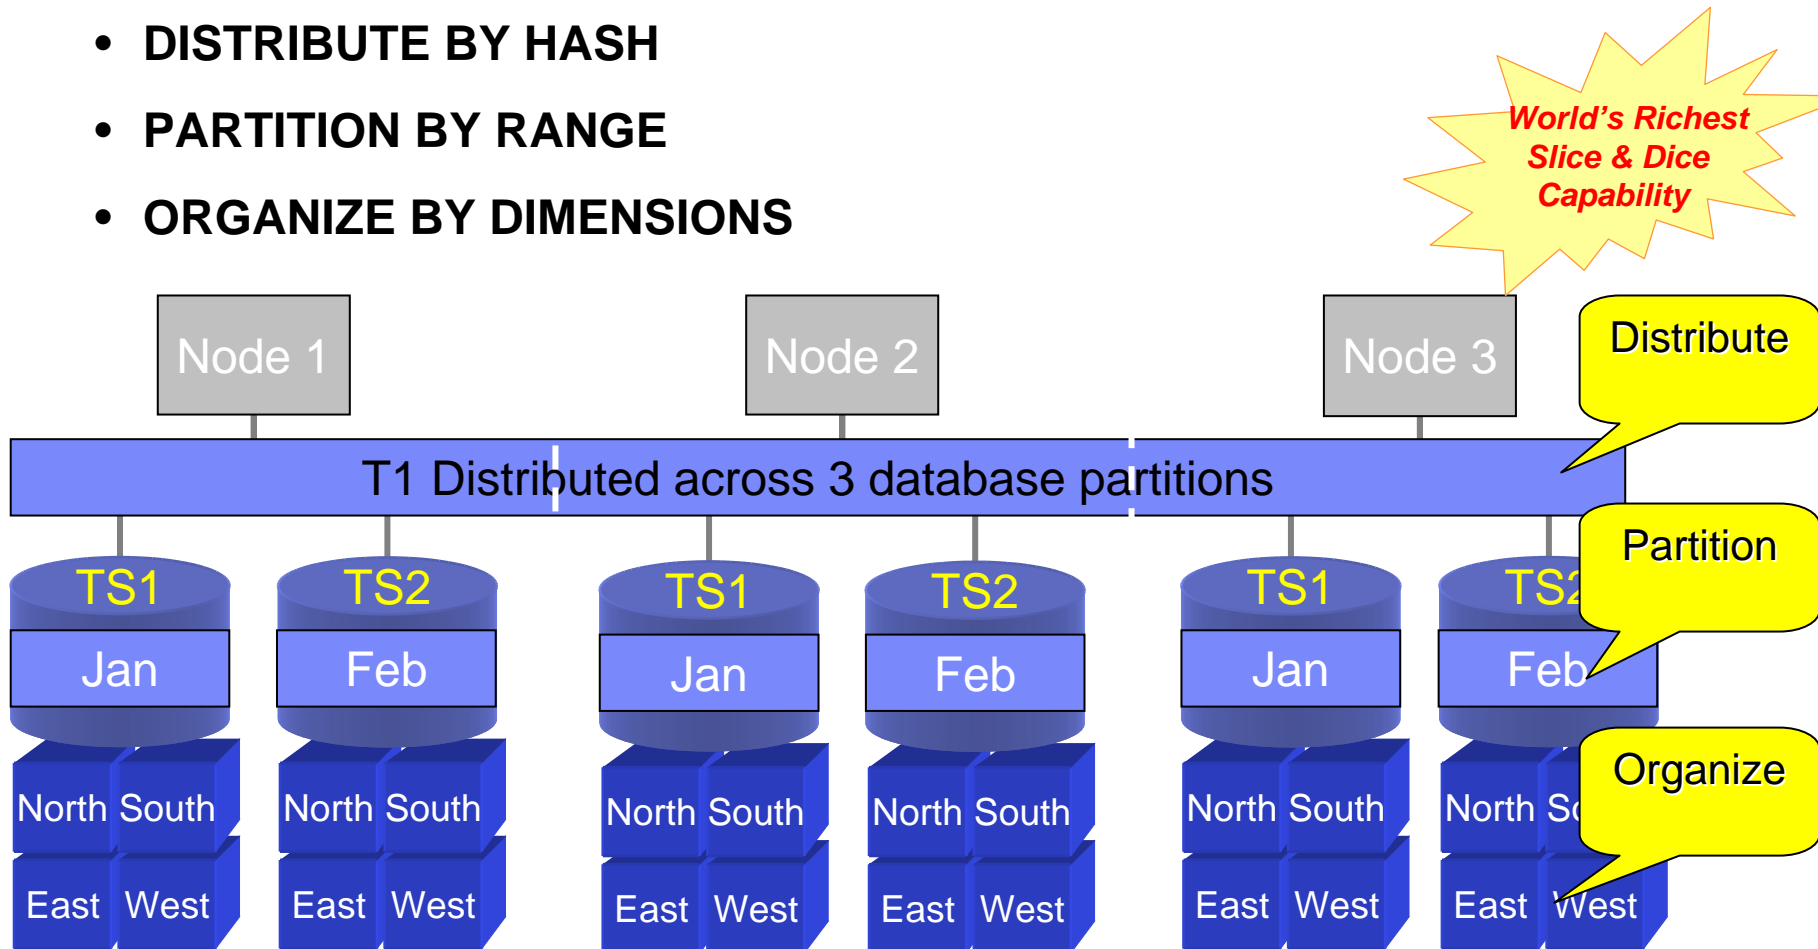


Filter required data directly in the warehouse
 Get the subset of products that you are interested in performing market basket analysis on. Integrated data movement and transformation capabilities allow you to do to this in line within mining processes.

Drag-and-drop interface
 Seamlessly add specific analytics and mining operations into a data flow and specify the attributes in the pane below

In addition data access can be optimised by storing according to the dimensions along which it is accessed. This is automatic and transparent for the user

- **DISTRIBUTE BY HASH**
- **PARTITION BY RANGE**
- **ORGANIZE BY DIMENSIONS**



Large Wireless Carrier

Understanding customers in real time


Business Challenge

- 360 Degree Customer View
- Unified Customer Contact Information
- Churn prediction

Solution

- Warehouse w/ Near Real-time Feeds
 - Load over 1B Call Records/Day (up to 1.6B)
 - 10 Billion transactions per day
 - 32 TB Raw Data
- 1,000s of Concurrent Users
 - 7,000 Customer Care Users
 - Up to 37000 queries/day
- IBM DW, SAS, 16x8 P5 pSeries

Continuous Data Load → **Call Data Records**



Business Benefits



- Fraud Detection < 4 hours
- Campaign Responses Up 66-300%
- Margin per Customer up 20%

Technology Benefits

- Scale & Performance
 - Users
 - Volatility
 - Data

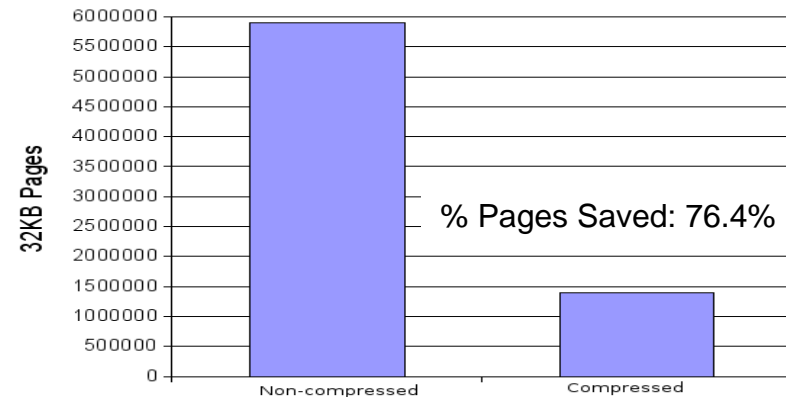


Data can be automatically compressed by 50+% to save storage, yet with little or no effect on performance – a massive saving




Tablescan	No Compression	Row Compression
# of Rows	59 986 052	59 986 052
# of Pages	1 023 162	426 292
MB	7 993	3 330
Run time (secs)	90,993	43,268
MB/sec	87,847	76,972
Total user cpu time (sec,ms)	48,91	94,49
Time waited for prefetch	302 878	67 950
user cpu time/row	0,00000082	0,00000158
user cpu time/page	0,000048	0,000222
rows/page	59	1

T1 Compression - 179.9GB Initial Size



Materialised Query Tables can be defined to automatically optimise queries. Transparent to the user they can be added or removed to maximise overall throughput

- **Tables of Pre-Summarized Data**
- **Cost-Based Optimizer Routing**
- **Maintained by the system**
- **Full and Partial Matching**
- **Transparent to Applications / Users**
 - System detects their presence and uses them if appropriate – no programming needed
 - Easy to add / remove from system




Cubing Services: OLAP Analytics and Open Access

Other Reporting Tools

Cognos

Excel




MDX


MDX

MDX

SQL /
ROLAP



InfoSphere Balanced Warehouse



New York Police Department

Timely and accurate information delivered to officers at the scene helps fight crime

Challenge

- Derive insight from billions of records and reach detectives within minutes, not days or weeks.
- Correlate arrests, complaints, summonses, homicides, shootings, locations, 911 calls

Solution

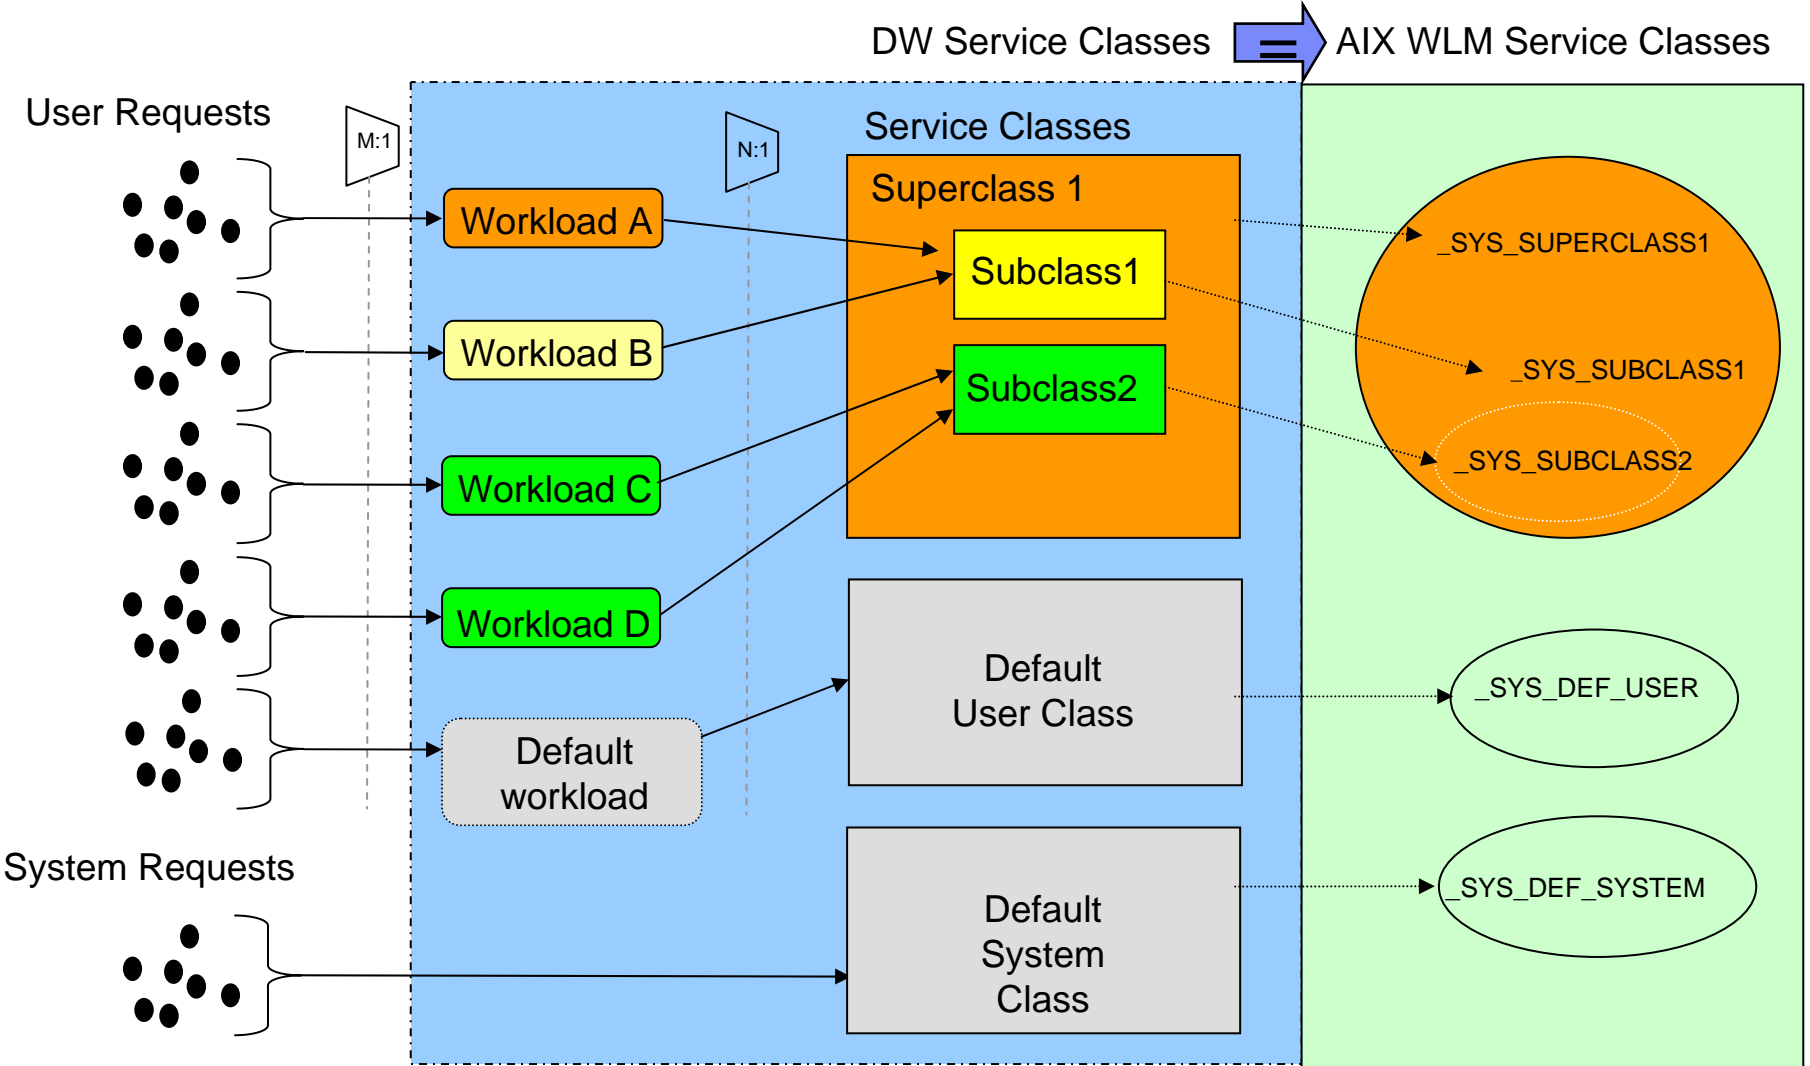
- IBM and Business Partner Cognos created a Crime Information Warehouse solution that provides the technology needed to integrate departmental data across different systems, resulting in a single source for actionable information that better addresses policing needs.
- NYPD can now analyze and integrate all information assets across structured and unstructured data sources in real-time



Business Benefits

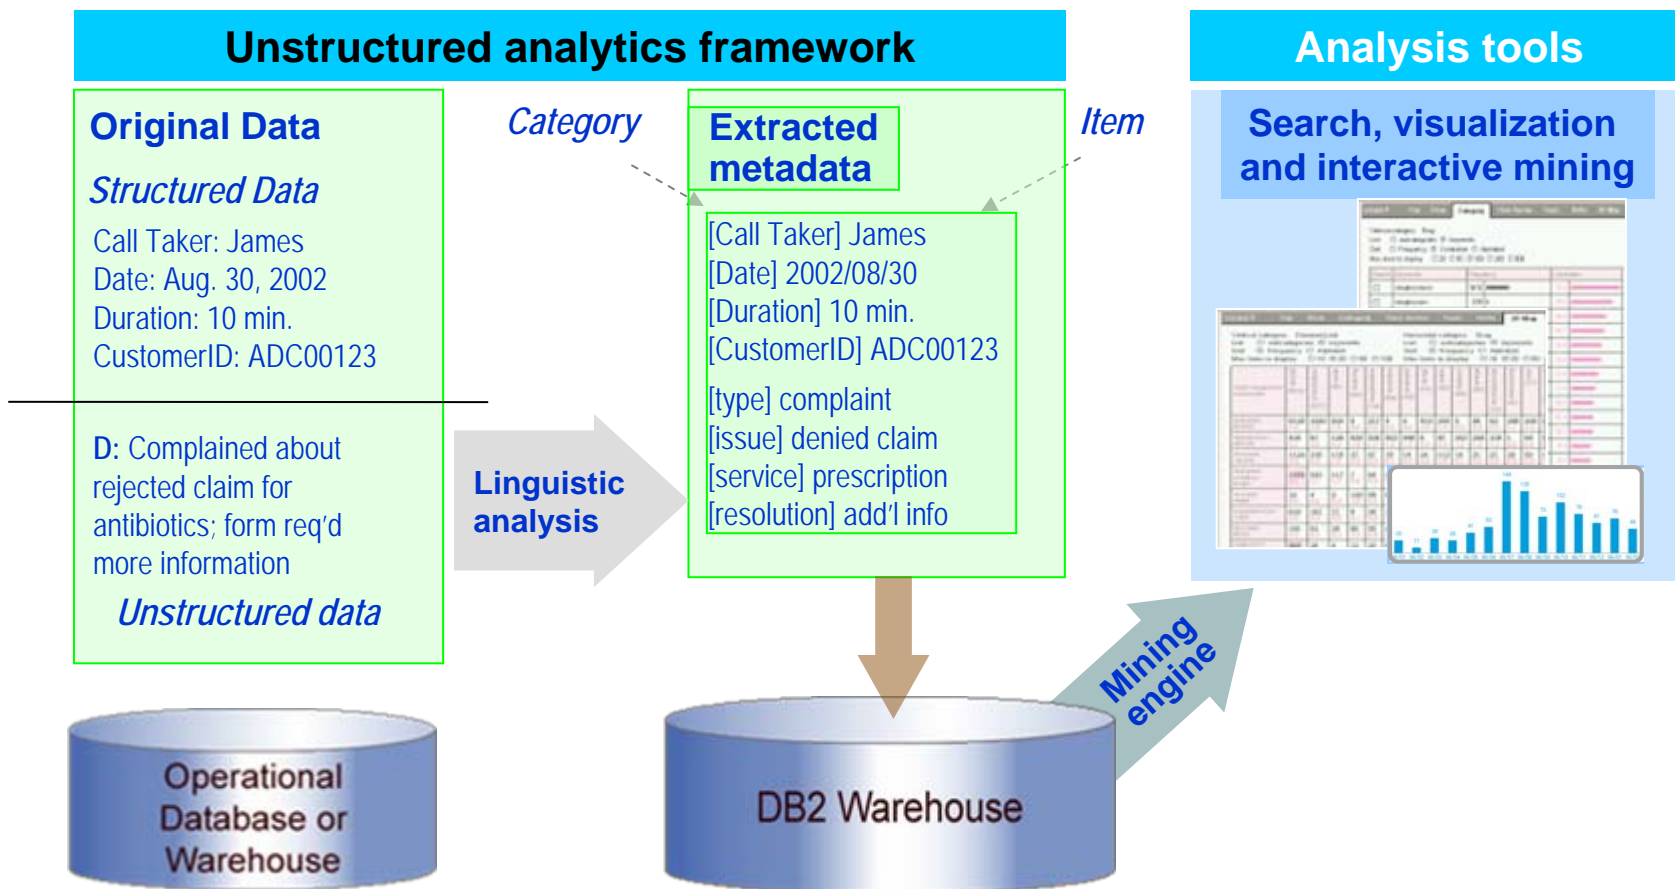
- All forms of information delivered in a timely fashion – realtime delivery
- Insightful relationships revealed through advanced analysis
- Timeliness means the difference between life & death
- Information reaches detectives within minutes, instead of days or weeks
- Integration of large number of complex data sources, both structured and unstructured

Workload Management is a particular strength of the Balanced Warehouse, multiple complex workloads can be defined and managed.



IBM OmniFind Analytics Edition provides analysis of unstructured text information

- Rich analysis interface for combining structured and unstructured data
- Combines search, text analytics and data visualization



Summary: IBM Balanced Warehouse speeds time to value by giving you everything you need in one easy to install package

1. Everything you need for warehousing

- Data Modelling
- Data Flow / Manipulation Tools
- Easy Data Mining
- Text analytics
- Cubing Services

2. Easy to install and administer

- Pre-configured, pre-tested
- 'Query Ready'
- Low TCO

3. High Performance and Scalable

- Full implementation of 'Shared Nothing' Massively Parallel Processing (MPP)
- Mixed Workload Support

4. Affordable

- Low Cost / Tb
- Extreme Compression
- Predictable costs

5. Proven

- Highly Robust and Reliable



Business Intelligence

Paul Levey, CSI

Address bar: http://lae9646/cognos8/cgi-bin/cognosisapi.dll?b_action=xts.run&m=portal/welcome/welcome.xts&homeQuery=b_action%3d

File Edit View Favorites Tools Help

Google Search [Icons: Mail, Search, Print, etc.] Sign In

Welcome to IBM Cognos 8 [Home, RSS, Print, Page, Tools, Help icons]

IBM[®] COGNOS[®] 8

My Content



[My home](#)



[IBM Cognos content](#)

Administration



[Administer IBM Cognos content](#)

My Actions



[Query my data](#)



[Analyze my business](#)



[Create professional reports](#)



[Manage my events](#)

Show this page in the future

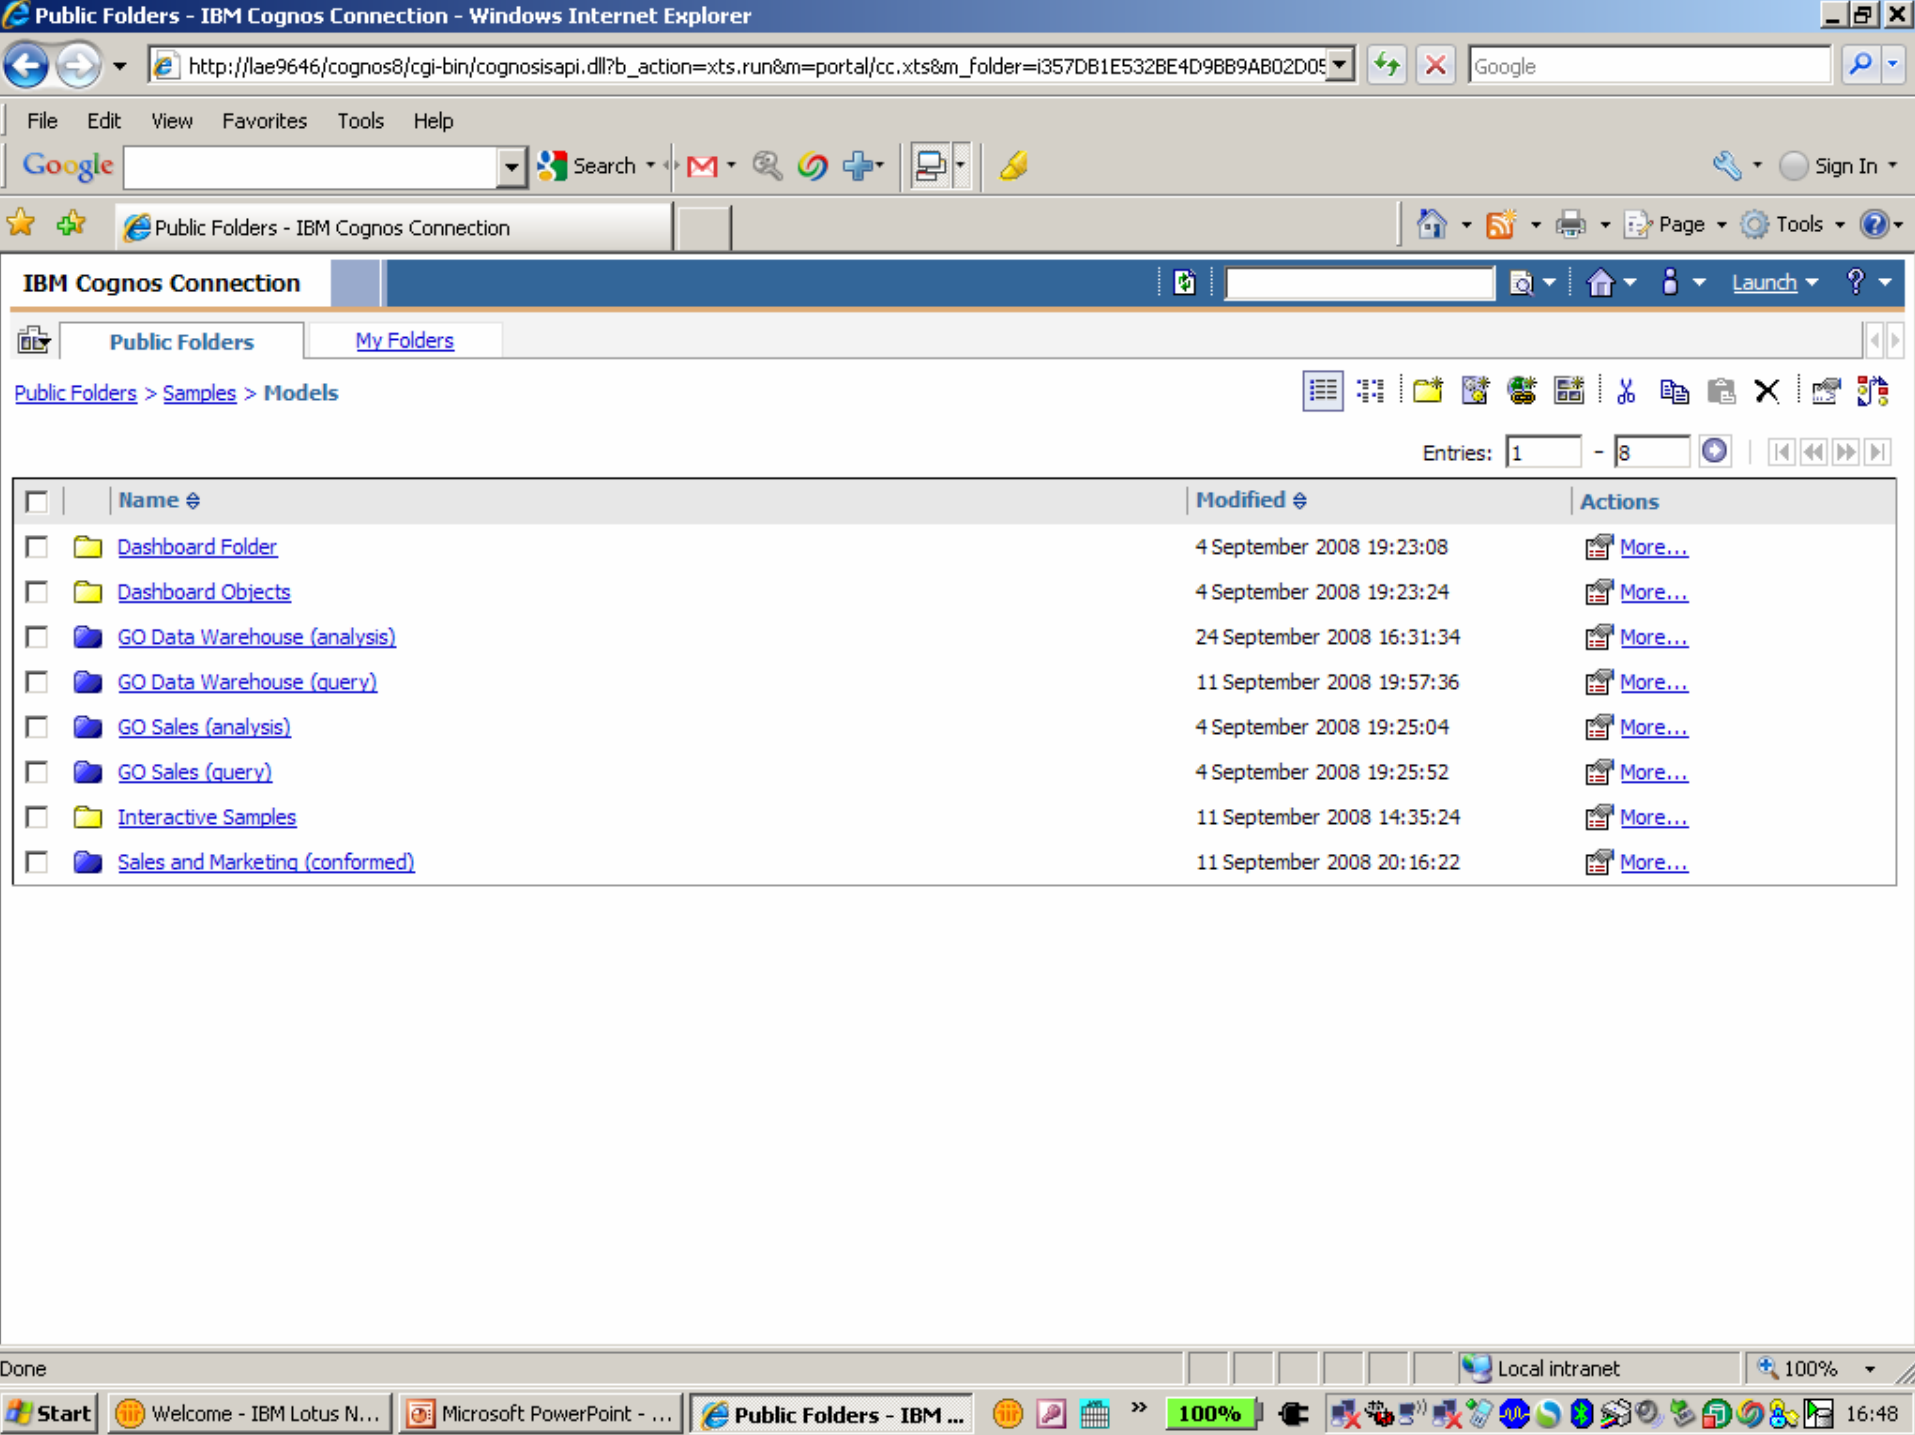
[Quick Tour](#)



Done

Local intranet 100%

Start Welcome - IBM Lotus N... Microsoft PowerPoint - ... Welcome to IBM Co... 100% 16:47



http://lae9646/cognos8/cgi-bin/cognosisapi.dll?b_action=xts.run&m=portal/cc.xts&m_folder=i357DB1E532BE4D9BB9AB02D05

Google

File Edit View Favorites Tools Help

Google

Search

Sign In

Public Folders - IBM Cognos Connection

Page Tools

IBM Cognos Connection

Public Folders My Folders

Public Folders > Samples > Models

Entries: 1 - 8


<input type="checkbox"/>	Name	Modified	Actions
<input type="checkbox"/>	Dashboard Folder	4 September 2008 19:23:08	More...
<input type="checkbox"/>	Dashboard Objects	4 September 2008 19:23:24	More...
<input type="checkbox"/>	GO Data Warehouse (analysis)	24 September 2008 16:31:34	More...
<input type="checkbox"/>	GO Data Warehouse (query)	11 September 2008 19:57:36	More...
<input type="checkbox"/>	GO Sales (analysis)	4 September 2008 19:25:04	More...
<input type="checkbox"/>	GO Sales (query)	4 September 2008 19:25:52	More...
<input type="checkbox"/>	Interactive Samples	11 September 2008 14:35:24	More...
<input type="checkbox"/>	Sales and Marketing (conformed)	11 September 2008 20:16:22	More...



Employee Satisfaction


By Department 2006

Survey topic scores by department



Customer Service average score is -12.1% compared to the company average.

Survey topic scores, targets and industry standard



Employee rankings and terminations by department

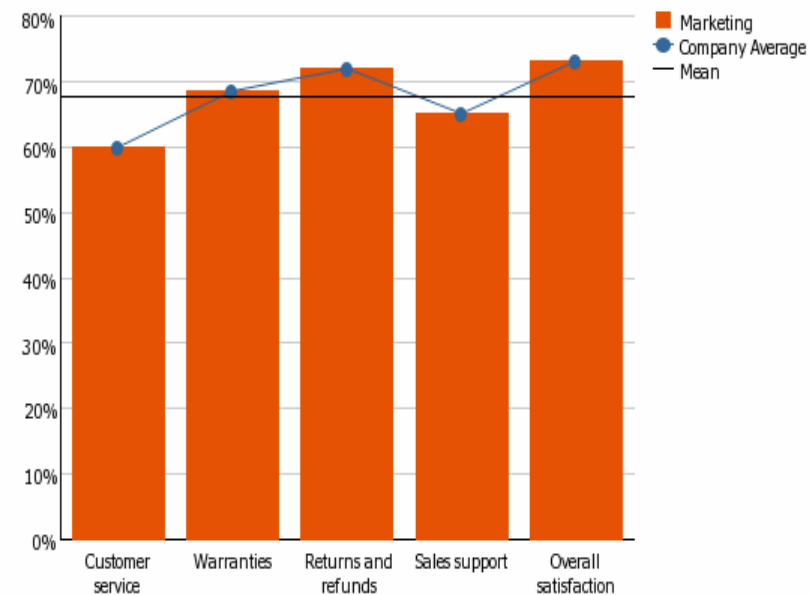
	Employee ranking					Terminations
	Poor	Satisfactory	Good	Very Good	Excellent	
Human Resources		2.9%	2.9%	2.9%	2.9%	2.9%
Finance	0.9%	0.9%	0.9%	0.9%	0.9%	1.8%
Marketing	3.0%	3.0%	3.0%	3.0%	3.0%	9.1%
Sales (Corporate)	2.9%	2.9%	2.9%	2.9%		8.8%
Procurement		3.8%	3.8%	3.8%	3.8%	0.0%
Information Services and Technology	2.5%	2.5%	2.5%	2.5%		2.5%
Production and Distribution	1.3%	1.3%	1.3%	1.3%	1.3%	7.9%



Customer Satisfaction

Asia Pacific Customers 2007

Survey Results



Return reason	Percentage of Customers	Return quantity
<u>Defective product</u>	5.66%	1,923
<u>Incomplete product</u>	32.35%	11,001
<u>Unsatisfactory product</u>	18.99%	6,456
<u>Wrong product ordered</u>	8.45%	2,873
<u>Wrong product shipped</u>	34.55%	11,749
Summary	100%	34,002

4 Golf only	Retailer type	Retailer topic score	Incomplete product	Return quantity	% of total retailer returns	% of all returns by return reason
4 Golf only	Golf Shop	72.12%	Incomplete product	23	100.00%	100.00%
4 Golf only				23		
Summary				23		



Insertable Objects

- Sales and Marketing (cube)
 - Time
 - Time
 - Current Month
 - Last Month
 - QTD
 - Prior QTD
 - QTD Grouped
 - YTD
 - Prior YTD
 - YTD Grouped
 - Products
 - Product brand
 - Retailers**
 - Promotions
 - Order method
 - Measures

Rows: Promotions Retailers Columns: Time Context filter:

Revenue		2004	2005	2006	2007	Time
Regular sale	Americas	280,177,081.19	339,630,425.66	437,478,810.25	313,811,914.03	1,371,098,231.13
	Asia Pacific	215,794,202.54	276,799,668.62	360,648,449.58	263,953,378.07	1,117,195,698.81
	Northern Europe	95,607,229.66	126,242,270.15	159,691,722.54	120,317,245.10	501,858,467.45
	Central Europe	193,496,733.49	252,449,058.66	309,699,176.06	231,222,764.33	986,867,732.54
	Southern Europe	86,672,525.48	115,085,178.65	154,160,096.94	113,628,190.69	469,545,991.76
	Retailers	871,747,772.36	1,110,206,601.74	1,421,678,255.37	1,042,933,492.22	4,446,566,121.69
	TrailChef Campaign	Americas	1,898,434.95	1,395,731.50	1,616,043.25	1,434,767.78
Asia Pacific		1,767,538.62	1,541,291.50	1,940,905.23	1,352,122.72	6,601,858.07
Promotions.Promotions.Promotions.TrailChef Campaign			586,057.86	700,762.27	693,679.20	2,749,095.58
Central Europe		1,472,328.40	1,186,629.21	1,340,977.42	1,218,354.57	5,218,289.60
Southern Europe		572,106.32	591,640.98	782,862.67	598,638.41	2,545,248.38
Retailers		6,479,004.54	5,301,351.05	6,381,550.84	5,297,562.68	23,459,469.11
EverGlow Campaign		Americas	1,171,713.73	472,839.15	274,757.26	1,109,661.03
	Asia Pacific	1,068,262.89	438,945.50	261,977.23	1,187,679.16	2,956,864.78
	Northern Europe	301,729.97	172,893.00	107,420.83	423,205.23	1,005,249.03
	Central Europe	868,652.70	349,070.43	210,600.53	939,571.45	2,367,895.11
	Southern Europe	370,867.91	157,164.21	130,176.80	482,731.11	1,140,940.03
	Retailers	3,781,227.20	1,590,912.29	984,932.65	4,142,847.98	10,499,920.12
	Course Pro Campaign	Americas	2,866,965.20	2,859,734.14	7,147,784.79	7,585,365.99
Asia Pacific		2,874,969.88	2,893,508.48	6,846,195.68	7,226,264.15	19,840,938.19

- Current Activities
- Past Activities
- Upcoming Activities
- System
- Schedules

Filter

Background activities
 Interactive activities


Run by:
 Any
[Select a user...](#)

Status:
 Any

Priority:
 Any

[Advanced options](#)

[Reset to default](#)



<input type="checkbox"/>	Name	Request time	Run by	Status
No entries.				

Last refresh time: 8 July 2009 12:10:55

IBM Cognos Administration

Status Security Configuration

- Users, Groups, and Roles
- Capabilities
- User Interface Profiles

Capability

Entries: 1 - 15 [Refresh] [Navigation]

Name ▲	Modified ⇅
Adaptive Analytics ▼	19 June 2009 14:49:17
Administration ▼	19 June 2009 14:49:17
Analysis Studio ▼	19 June 2009 14:49:17
Cognos Viewer ▼	19 June 2009 14:49:17
Controller Studio ▼	19 June 2009 14:49:17
Data Manager ▼	19 June 2009 14:49:17
Detailed Errors ▼	19 June 2009 14:49:17
Drill Through Assistant ▼	19 June 2009 14:49:17
Event Studio ▼	19 June 2009 14:49:17
Glossary ▼	19 June 2009 14:49:17
Hide Entries ▼	19 June 2009 14:49:17
Lineage ▼	19 June 2009 14:49:17
Metric Studio ▼	19 June 2009 14:49:17
Planning Contributor ▼	19 June 2009 14:49:17
PowerPlay Studio ▼	19 June 2009 14:49:17

Last refresh time: 8 July 2009 12:58:20