



# Data Warehousing in *System z*

Carlos Guardia Rivas  
Executive IT Specialist - IBM Software Group



## Agenda

- *Datawarehousing trends and directions*
- *Business intelligence on System z*
- *System z datawarehousing architecture*
- *DB2 for z/OS datawarehousing functionalities*
- *ISAS 9600*
- *The IBM Business Intelligence complete picture*

## Data volume is exploding...



As of December 2008, the global monthly Internet traffic is estimated to be **5 to 8 exabytes**.

### Yet businesses still struggle

#### Lack of Insight

One in three managers frequently make critical decisions without the information they need



As of May 2009, the size of the World's total Digital content has been roughly estimated to be 500 billion gigabytes, or **500 exabytes**.

#### Inefficient Access

One in two don't have access to the information across their organization needed to do their jobs



By 2013, annual global IP traffic will reach **two-thirds of a zettabyte or 667 exabytes**.

#### Inability to Predict

Three in four business leaders say more predictive information would drive better decisions

Internet video will generate over **18 exabytes** per month in 2013


Source: Wikipedia 3/2010

Source: IBM Institute for Business Value

© 2010 IBM Corporation

The foremost issue facing enterprises today is the waste, inaccuracy and volume of missed opportunities that stem from information raging out of control

### Machine-generated versus authored data



#### Volume of Digital Data

- 57% CAGR for enterprise data through 2010
- Machine generated data : Sensors, RFID, GPS..

#### Variety of Information

- 80% of new data growth is unstructured
- Emails, images, audio, video..

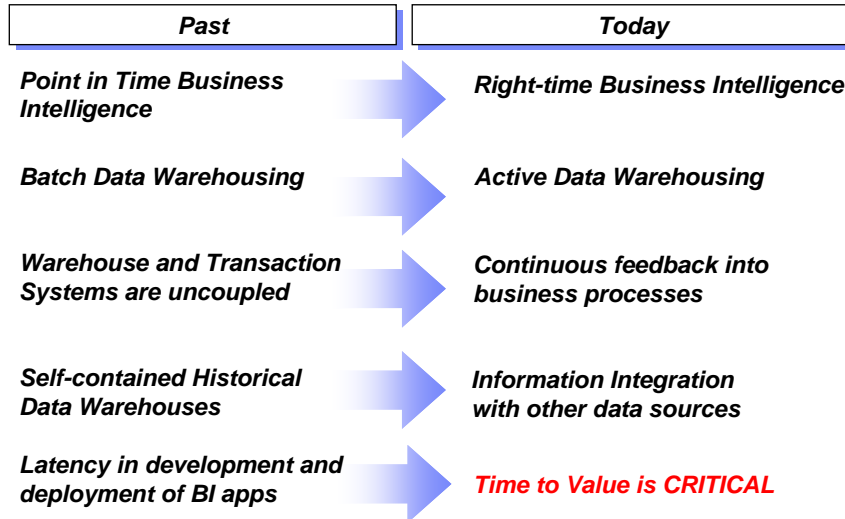
#### Velocity of Decision Making

- 42% of managers say they use wrong information at least once per week
- Rapidly changing business climate

Information is not being effectively captured, managed, analyzed and made broadly available to everyone who needs it

© 2010 IBM Corporation

## Business Intelligence Becoming Mission-Critical



## DataWarehousing trends

- High requirement for system, platform and data security
- Desire to optimize & leverage existing z infrastructure and skills
- Minimize complex and costly data movement
- Need help managing high growth of '3 Vs' – Volume, Variety, Velocity
- Increasing demands for sophisticated analysis with real time operational data — BI is becoming *mission critical*



### Dynamic Warehousing A New Approach to Leveraging Information

Information On Demand to Optimize Real-Time Processes



Dynamic Warehousing

OLAP & Data Mining to Understand Why and Recommend Future Action



Traditional Data Warehousing

Query & Reporting to Understand What Happened

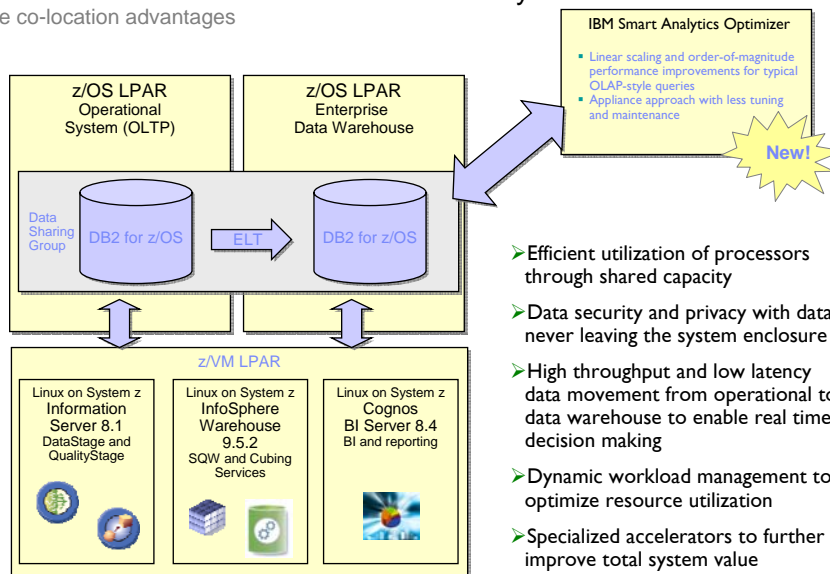


## Data Warehousing on DB2 for z/OS – What is driving this?

- **Many System z customers already use DB2 for z/OS for warehouse and BI**
  - IBM is responding to customer demand with new DB2 features, new software offerings and improved hardware performance and efficiency
- **The System z platform offers superior Total Cost of Ownership**
  - Customers want to leverage their System z infrastructure to do more with what they are already using successfully
    - TCO can be reduced through the utilization of existing processors, people, practices
    - TCO may also be achieved through a consolidation approach
- **New BI trends map well to the strengths of DB2 for z/OS and System z**
  - The distinction is blurring between warehouse and OLTP databases based on new trends such as Dynamic Warehouse and Operational BI, driving:
    - The need for increased reliability, availability, security, and compliance in a DWH DBMS
    - The need for very current warehouse data and/or collocation of warehouse and operational data
- **Specialty processors and the new z10 provide additional ways to optimize TCO**
  - zIIPs and IFLs are driving down hardware and software costs; DWH/BI can make excellent use of these processors, ultimately driving TCO advantages
  - The new processors are delivering excellent speeds


## The IBM Data Warehouse Architecture on System z

Unique co-location advantages




- Efficient utilization of processors through shared capacity
- Data security and privacy with data never leaving the system enclosure
- High throughput and low latency data movement from operational to data warehouse to enable real time decision making
- Dynamic workload management to optimize resource utilization
- Specialized accelerators to further improve total system value

## Topology (I)




1 LPAR / 1 subsystem




1 LPAR / 2 subsystems – data sharing

- Common zparms not desirable for mixed workloads
- Segregate buffer pools between OLTP and DW
- Virtual storage pressure
- Lower availability
- Tuning challenges
- Both OLTP and DW data available to each DB2 member, enabling drill down
- Possible to join OLTP and DW tables
- Simpler data movement for ETL
- Better availability

## Topology (II)



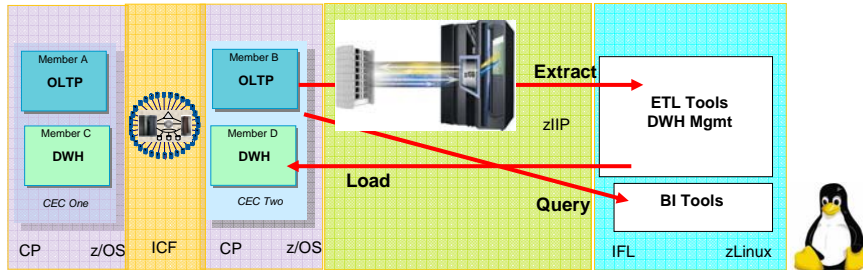
2 LPARs / 2 subsystems – data sharing



4 LPARs / 4 subsystems – data sharing

- Higher availability
- Shared processor capacity when LPARs located on same server
- Highest availability
- Work does not stop completely when a DB2 member/LPAR fails

## Specialty Processors in a DB2 for z/OS Warehouse Solution



**ICF** – Uniquely allows a Data Warehouse database to coexist with an OLTP database

**IFL** – Enables efficient data movement (secure, high-speed hipersockets)  
 – Lowers TCO through reduced hardware and software costs  
 – Enables use of zIIPs during extract and further reduces costs

**zIIP** – Further enables lower cost of Business Intelligence queries

## DB2 for z/OS features that support DWHing

### DB2 V8

- 64-Bit Addressability
- 2000 byte index keys
- MQT's
- Multi-Row Operations
- 225 way table joins
- In-Memory Workfiles
- Automatic Space Allocation
- **Non-Uniform Distribution Statistics on Non-Indexed Columns**
- **Parallel Sorting**
- **Data Partitioned Secondary Indexes**
- 2MB SQL Statements

### DB2 9

- **Partition by growth**
- **Index Compression**
- **Index on expression**
- **Dynamic index ANDing**
- RRF new row internal structure for VARCHAR
- Fast delete of all the rows in a partition
- Skipping uncommitted inserted/updated qualifying rows
- **Histogram Statistics**
- SQL enhancements
  - INTERSECT, EXCEPT, cultural sort, caseless comparisons, RANK, DENSE\_RANK, ...
- **DB2 VUE**




## Some DB2 10 Features for Data Warehousing

- **Moving Sum, Moving Average**
- **Enhanced query parallelism technology**
- **Workfile:**
  - spanned records
  - PBG support
  - in-memory enhancements
- **LOAD and INSERT performance**
- **Safe query optimization: assess “reliability” of access path choices**
- **More Access path stability**
- **IN list performance**
- **RID pool overflow to workfiles**
- **Index include columns**
- **Auto Stats**
- **Index list prefetch**
- **Advanced query acceleration techniques**

© 2010 IBM Corporation




## DB2 for z/OS for DWHing



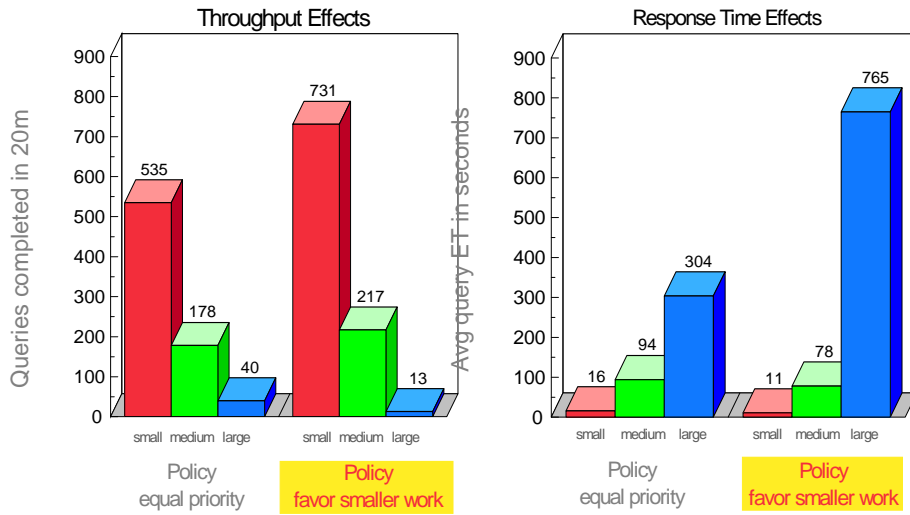
**Concurrent workload management**

**High compression ratios**



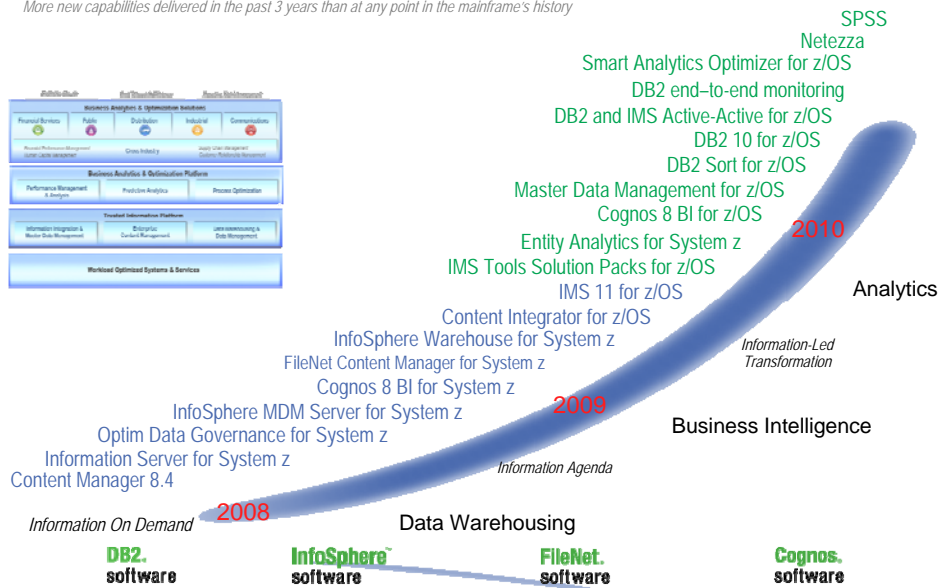
© 2010 IBM Corporation

## DB2 for z/OS workload management



## Information-Led Transformation for System z

More new capabilities delivered in the past 3 years than at any point in the mainframe's history







## IBM Smart Analytics System 9600 Unprecedented Value in Deploying New Business Analytics

### What is it?

The IBM Smart Analytics System 9600 is an integrated solution of hardware, software and services that enables customers to rapidly deploy cost effective game changing analytics across their business.



### How is it different?

- **Secure, Available Business Analytics**
  - Rapidly delivers analytic information to decision makers at the time of decision.
  - New environment for the availability, reliability and scalability necessary to stay aligned with the operational systems
- **Simplified administration**
  - Appliance-like delivery
  - Faster deployment at lower cost.
  - Leverages customers existing disaster recovery, and backup processes
- **Proven Operational Characteristics**
  - Extends the qualities of service of System z.
  - Reduces risk through extending System z manageability and security across the entire system.
- **High Value Operational BI**
  - Cost effective way to drive daily operational decisions

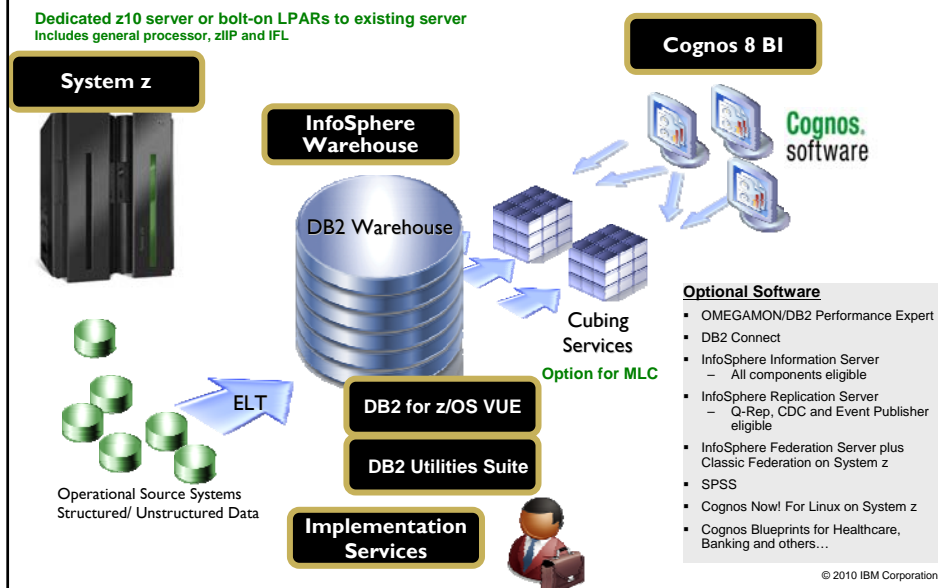
**Delivering business results in days, not months**

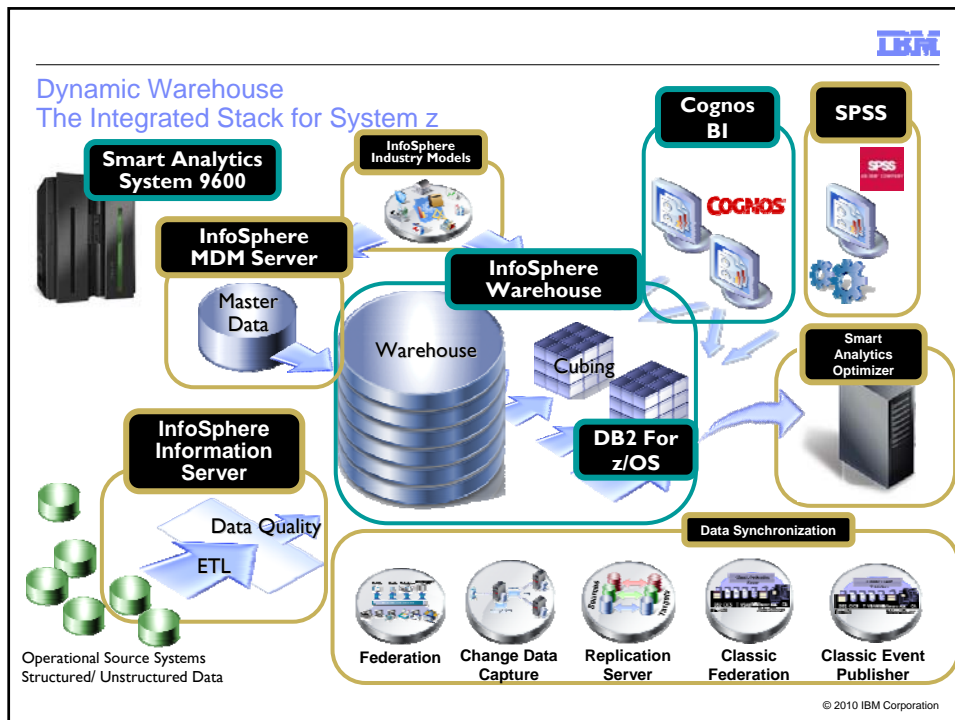
© 2010 IBM Corporation



## IBM Smart Analytic System 9600: High Value Dynamic Warehousing

Dedicated z10 server or bolt-on LPARs to existing server  
Includes general processor, zIIP and IFL





- 
- ### Benefits of a DB2 for z/OS Warehouse
- Blurring of distinction between warehousing applications and OLTP
    - Users demanding reliable access to reliable data from multiple sources
    - Global competitiveness requires 24 x 7
    - Evolving demand for real time or near real time
  - Operational and datwarehouse data together means
    - Reduced complexity
    - Reduced cost
    - Shared processes, tools, procedures
    - Streamlined compliance and security
  - Significant DW capabilities in V10 make DB2 competitive
    - DB2 VUE
  - zIIP specialty engine allows for IT optimization
  - Unprecedented query performance with SAO appliance
  - Bundled offering with ISAS 9600
  - Better leverage System z skills and investment
  - No other solution can match the qualities of service of z
- © 2010 IBM Corporation