Leveraging DB2 for z/OS Data Sharing for Your Data Warehouse

Session Number 1351A

Gary Crupi, IBM

IBM Software
Information On Demand 2011

Topics for Today

- When is System z interesting and what is the analytics solution on System z?
- When would data sharing be used?
- What are the implications of this setup?
- Q&A

When is System z interesting for date warehousing and business intelligence / analytics?

- **Start with an existing investment in System z infrastructure**
  - HW/SW/Applications, Processes and People

- **Add the need to see value in one or more of the following areas:**
  - Proximity of the operational production data for security, low latency or application integration
  - Quality of Service provided by System z
    - Continuous Availability, DR Ready, Secure, Responsive Mixed Workload Management, Instrumentation
  - Unique opportunity to mitigate risk and deliver the right information at the right time in a granular, incremental fashion
    - And integrating Dev/Test/QA in a single server

- **With requirements for one of several use cases:**
  - Enterprise Data Warehouse
  - Data distribution hub
  - High speed OLAP data mart
  - Data Mart Consolidation
  - Operational data store for online analytics
  - Competitor data archiving

IBM Smart Analytics System 9700 / 9710*

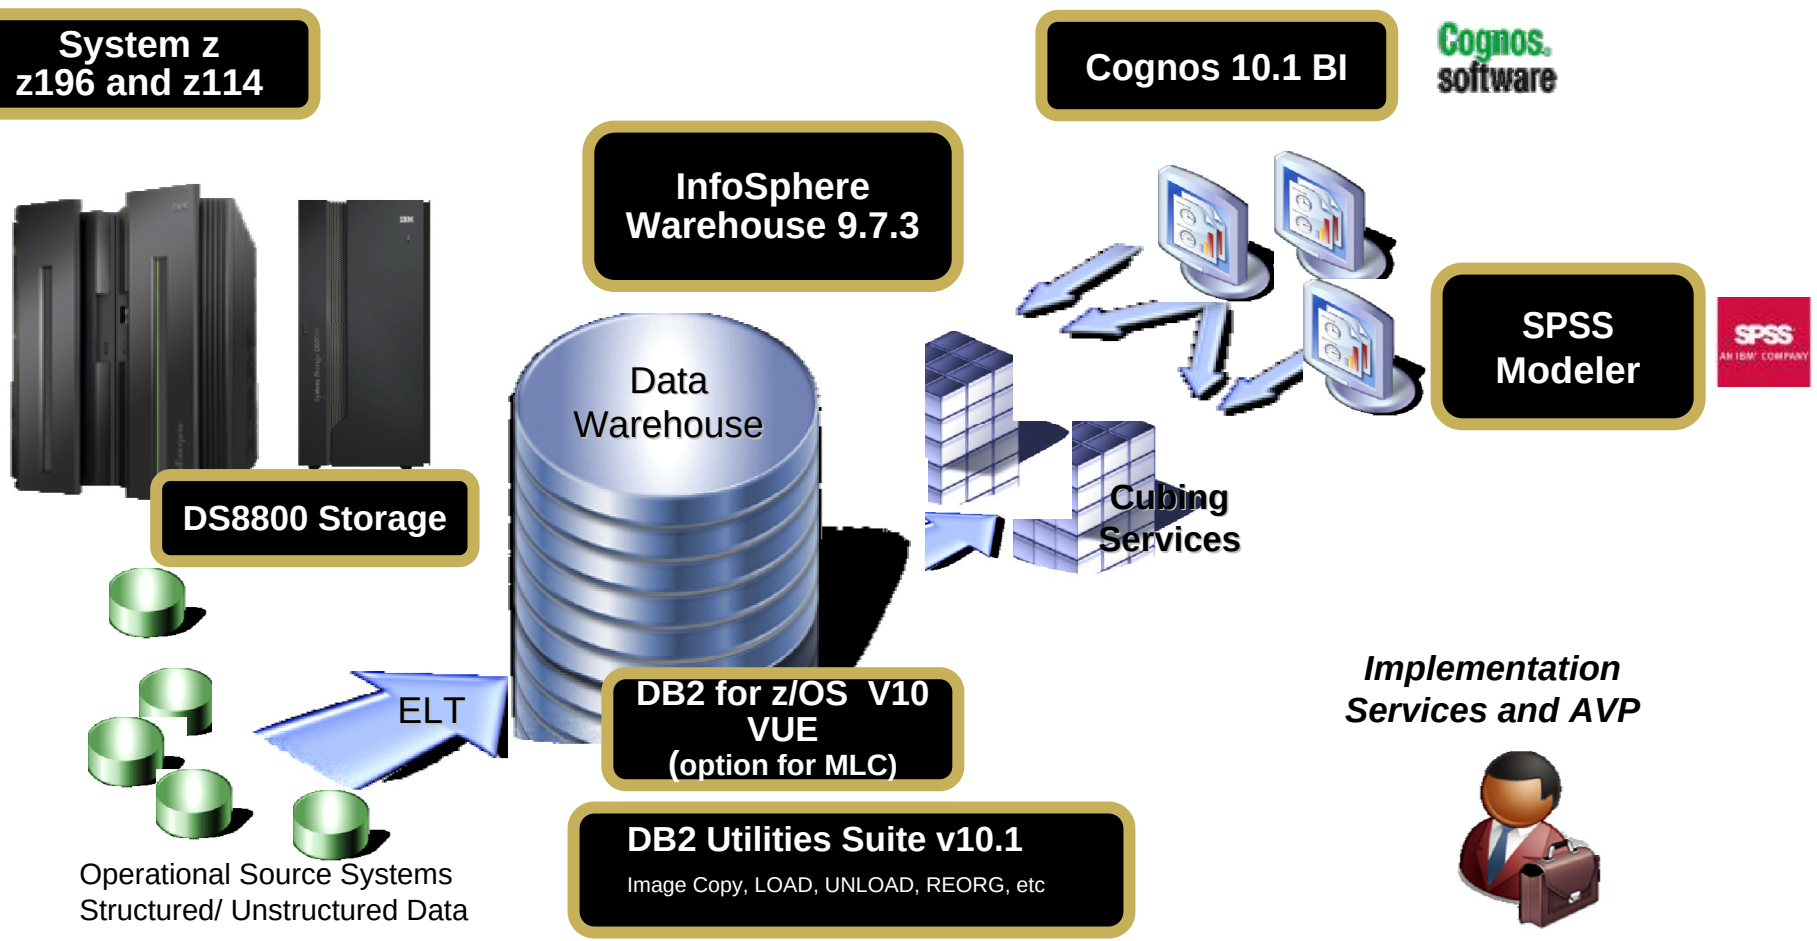

***Standard and Foundation Configuration available for the 9700, Foundation Configuration only for 9710 (build on additional components as needed)**

IBM Smart Analytics System 9700 / 9710 and DB2 Analytics Accelerator

The ultimate hybrid, workload optimized analytics server

- **Applies System z level of security, availability, integrity and reliability (in all the categories the highest in industry) to ALL your data (transactional and analytical)**
- **A true hybrid DBMS: providing applications a single interface to both transactional and analytical data**
- **The fastest performance thanks to specialized query engines: DB2 for transactional and Netezza for analytical access pattern to data**
- **TCO reduction through the effects of consolidation: single security policies, uniform database management, same skills, same tools, same procedures (instead of separate approaches in each of these categories).**
- **Delivering dedicated capacity at a competitive price that matches the value received**

Scenarios that can leverage data sharing

Transactional Analytics – Workload Distribution
High Concurrency, High Volume, High Availability Required

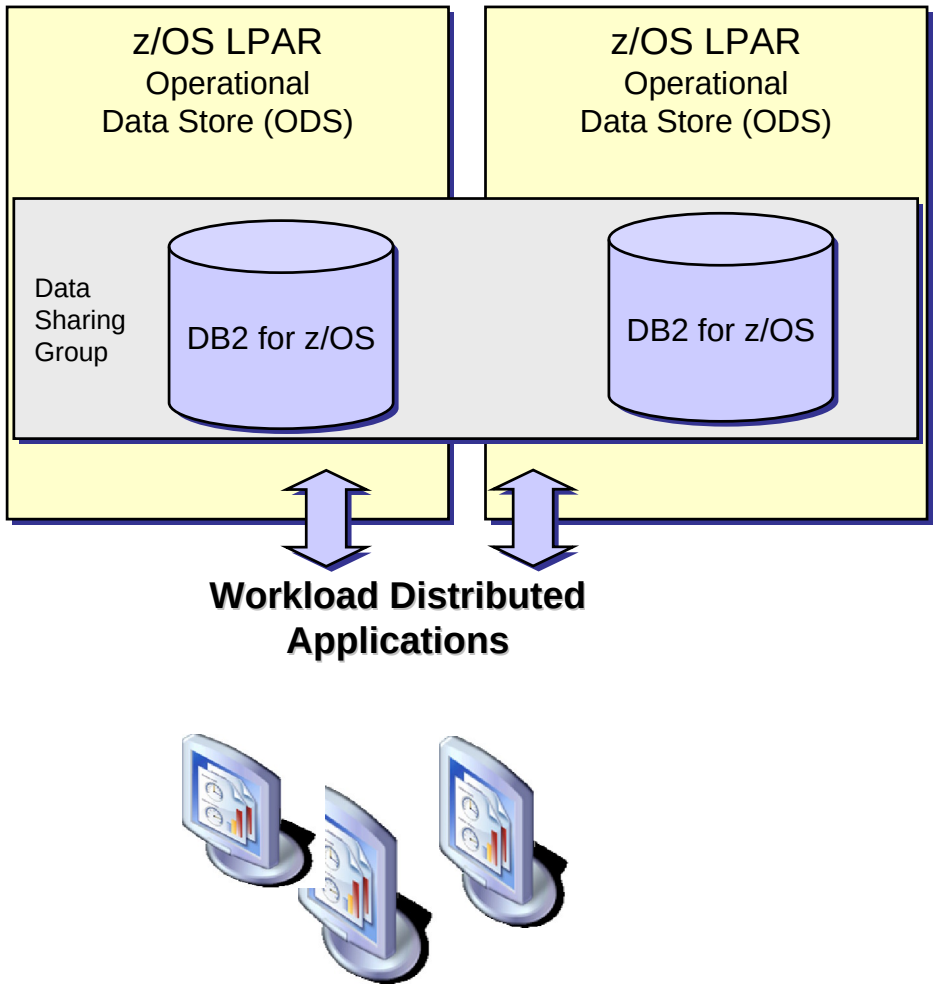

Typical Scenario

- Critical OLTP System
  - Desire to Isolate non-direct revenue producing activity
  - Web access must be 24X7 (emphasis on QoS)
  - User population measured in hundreds of thousands
  - Estimated 10K concurrent web users (trans/sec)

Solution

- Data Sharing for Web Facing ODS
  - Highly Available Operational Data Store
  - OLTP Replication to ODS provides data availability in seconds
  - Maintain OLTP system responsiveness, and availability.
  - Leverage IBM Smart Analytics System and 90%+ specialty processor offload
  - Introduce IBM DB2 Analytics Accelerator for business analysis

Highly Available Analytics Required – Large variety of consumers

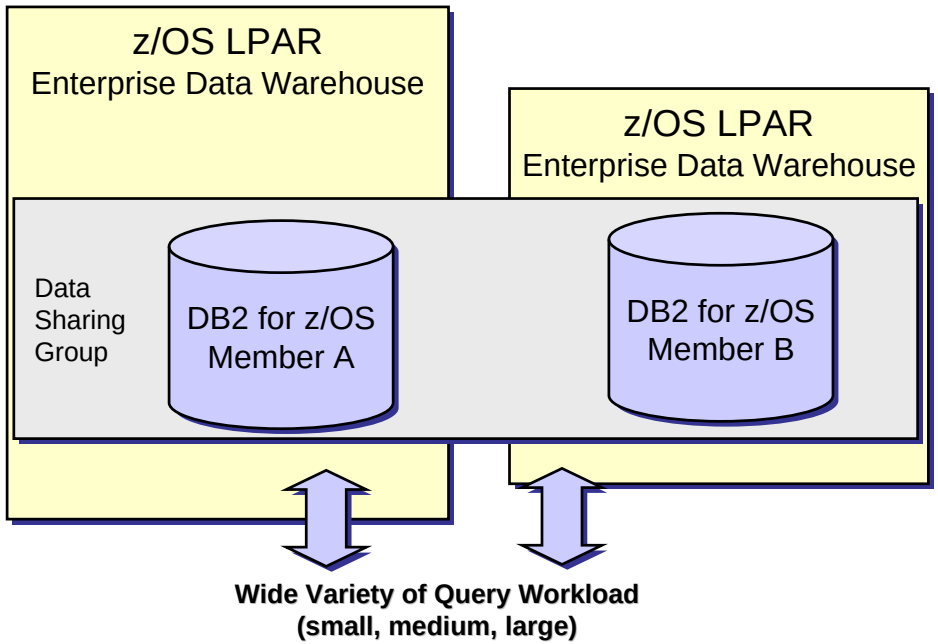

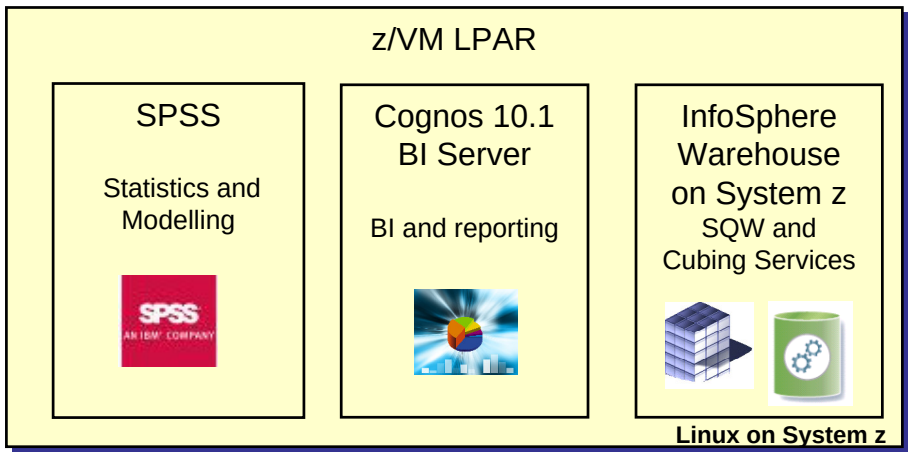

Typical Scenario

- Business requirement to build a data warehouse environment that is highly available even during planned and unplanned outages.
- Mixture of queries ranging from simple to extremely complex
- Providing the infrastructure to service the entire workload in duplicate systems is cost prohibitive, but query and other SLAs must be maintained even during outage periods.

Solution

- Data Sharing with complex query member, metadata/administration member, and capacity backup
- Routing is normally achieved via location aliases
- Leverage DB2 10 for z/OS Dynamic Location Aliases and Automation
  - Other alternatives available
- http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.dshare/src/tpc/db2z_managedynamiclocationalias.htm

Fit The Solution To The Workload – Access to all Data

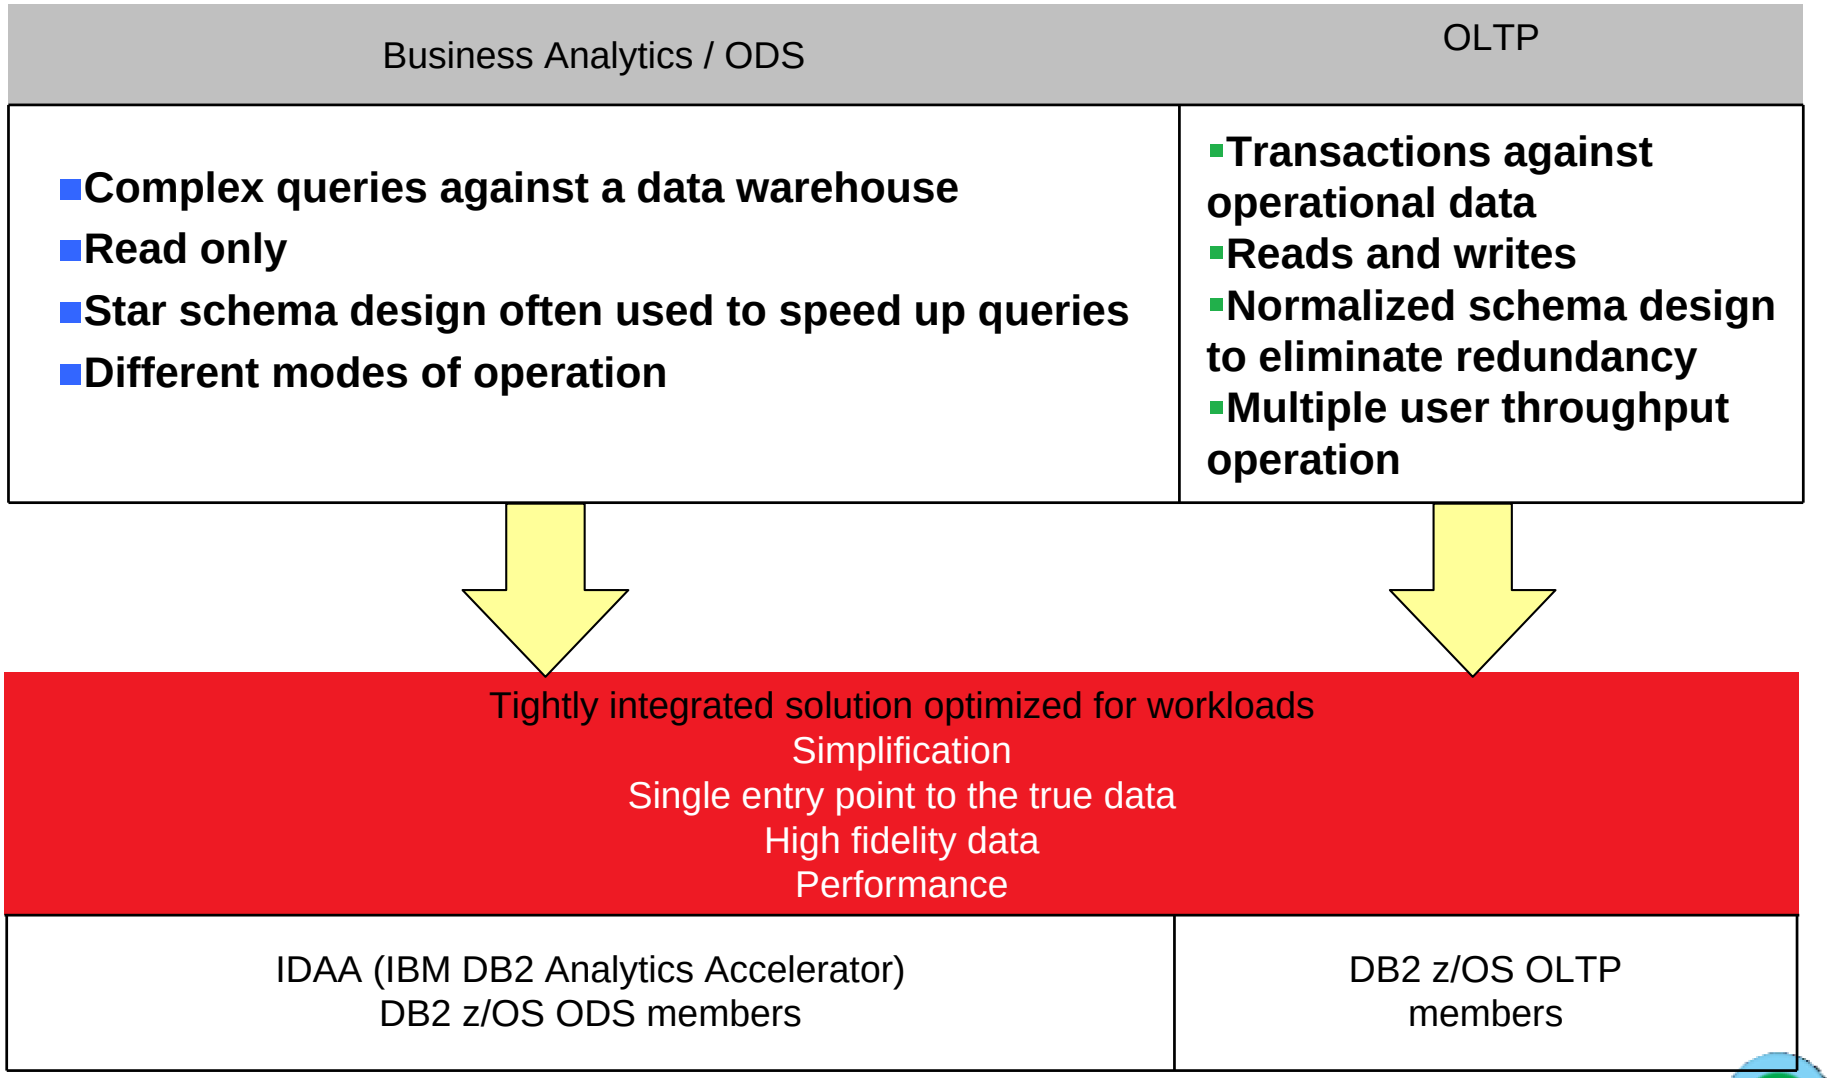

OnLine Transactional and Analytics Processing (OLTAP)

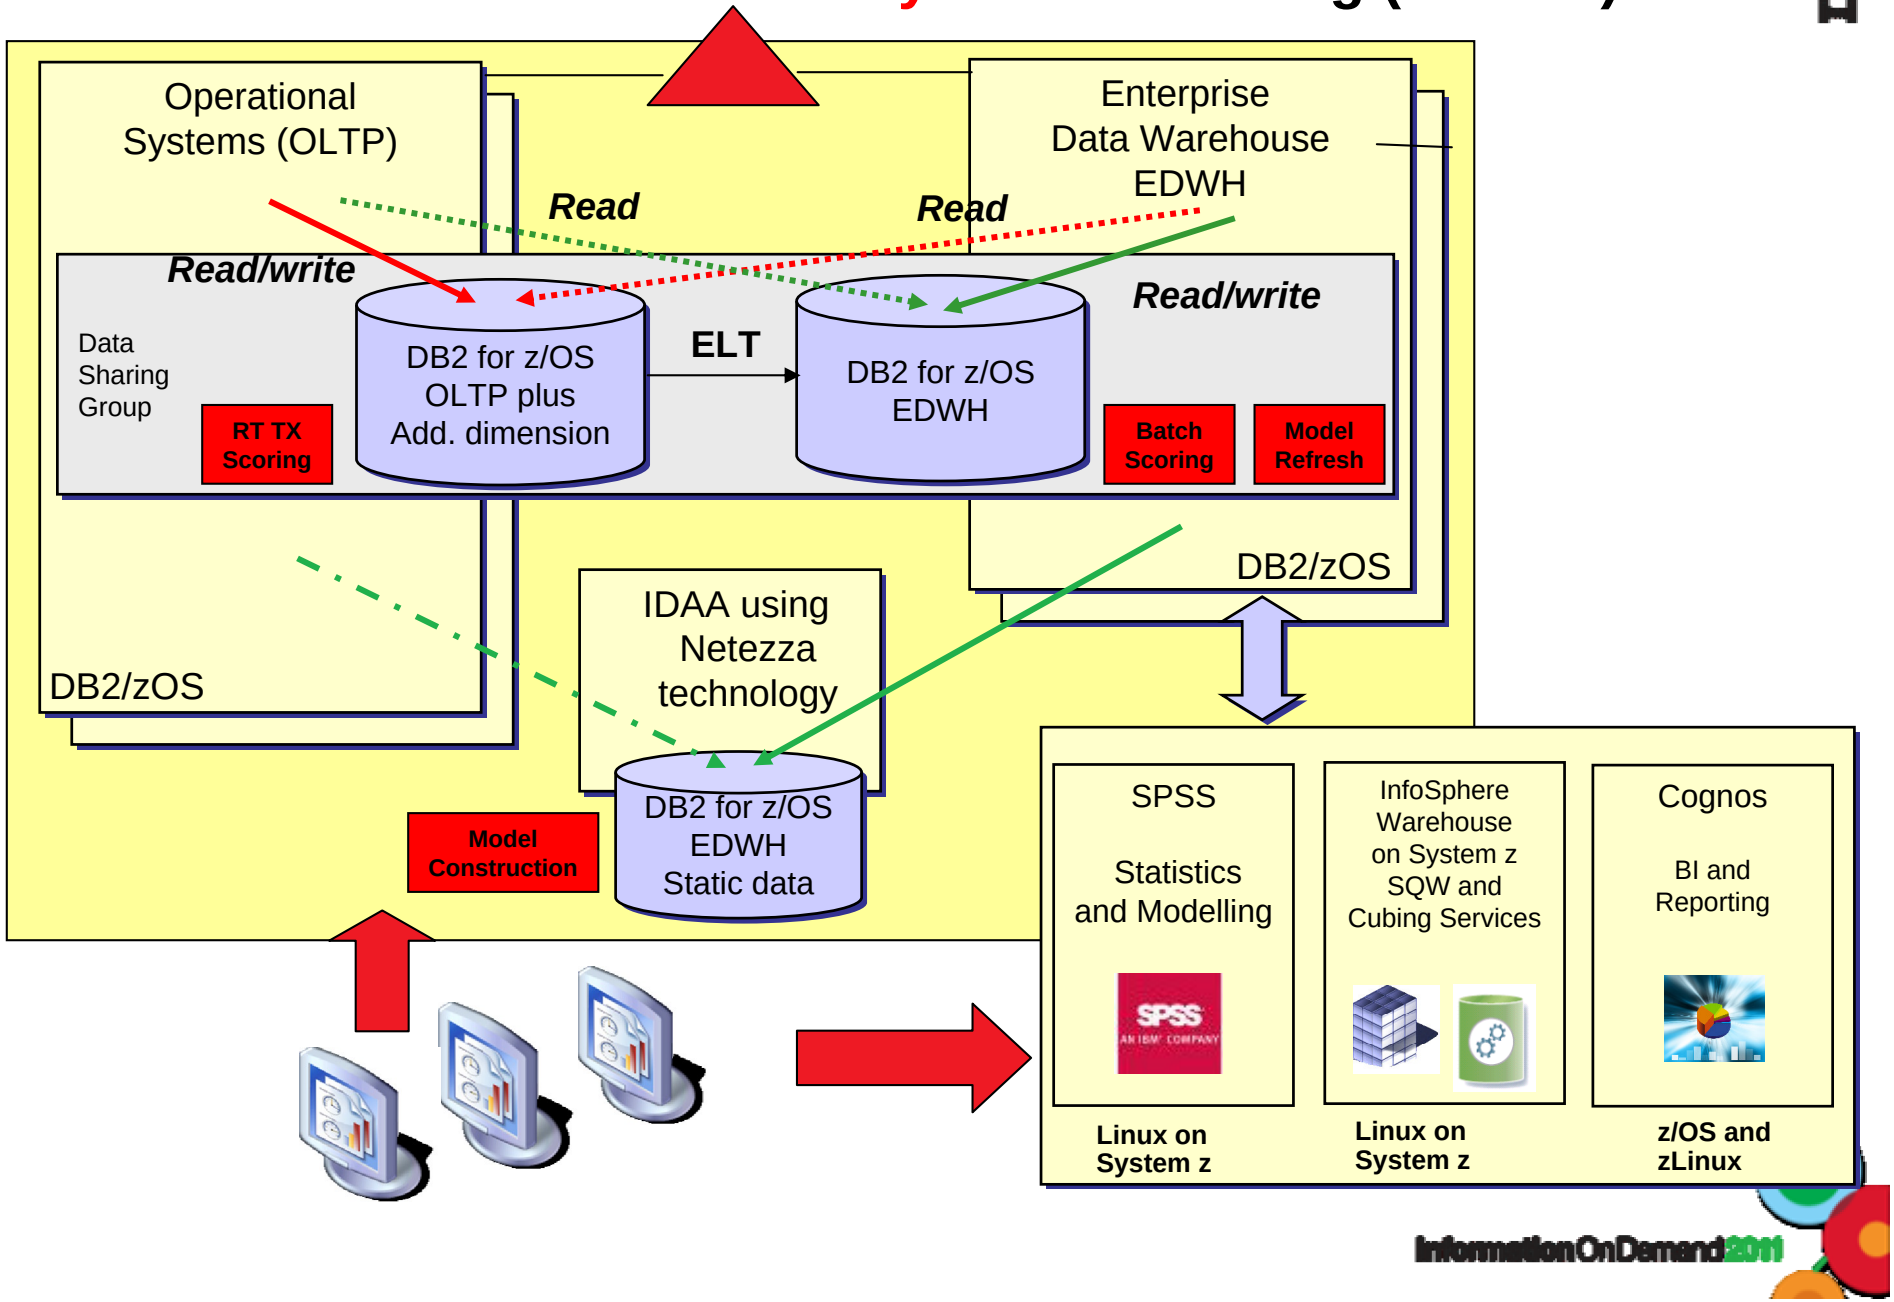

Unique Value for all of your Data Serving Requirements

- Simplification of Business Analytics Lifecycle Architecture and access to data for real-time predictive analytics without impact to OLTP
- Deliver more critical analytics for the enterprise with shorter time to value
- Benefit from well-established System z core strengths:
  - Leverage existing environment, high availability, backup and governance procedures as well as skills
  - Leverage architecture of isolated systems that can access all data (Data Sharing)
  - Efficient data movement (ELT instead of ETL)
    - Low Latency movement of data with Replication Server
- Benefit from recently delivered or soon coming functionality
  - Core System z Business Analytics Software Portfolio
  - System z196 technology in support of compute-intensive processing
  - IDAA with integrated Netezza technology
  - DB2 10 for z/OS enhancements

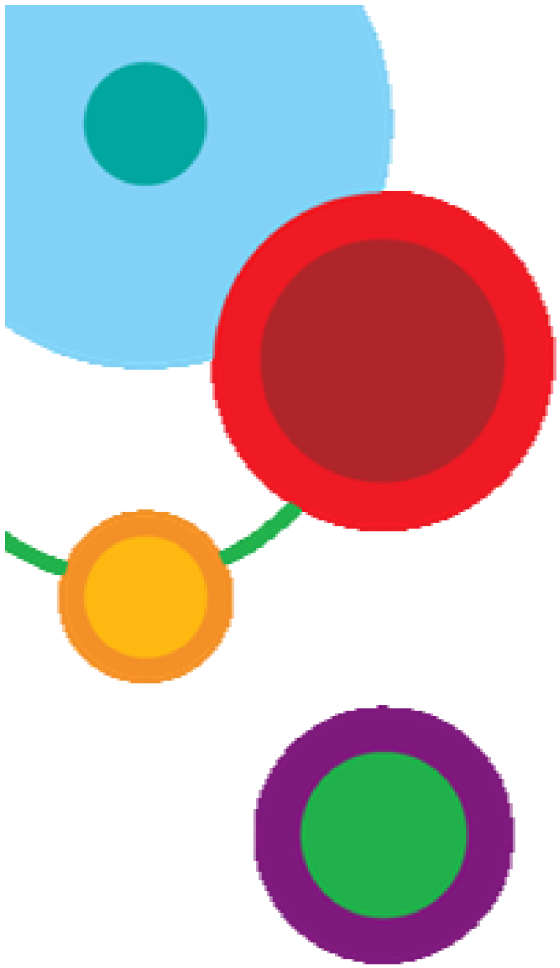

What should be considered when using a data sharing implementation in this manner?

Topology

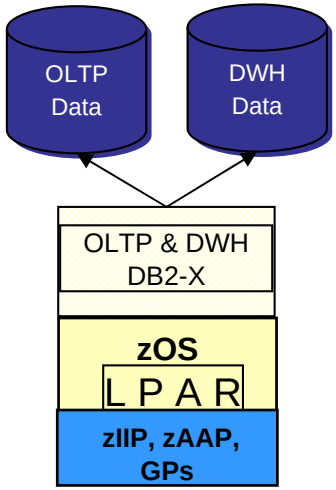

1 LPAR / 1 subsystem

- Common zParms may not desirable for mixed workloads
- Segregate buffer pools between OLTP and DW
- Virtual storage pressure
- Lower availability
- Tuning challenges

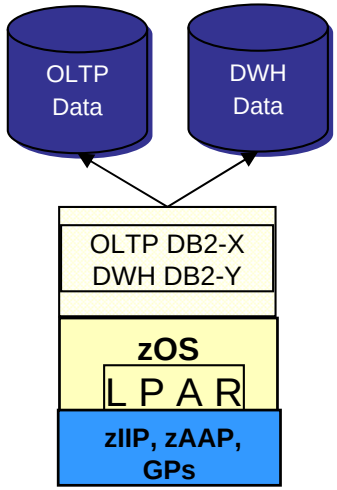

1 LPAR / 2 subsystems – data sharing

- 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
- Virtual storage pressure relief
- Better availability

Topology

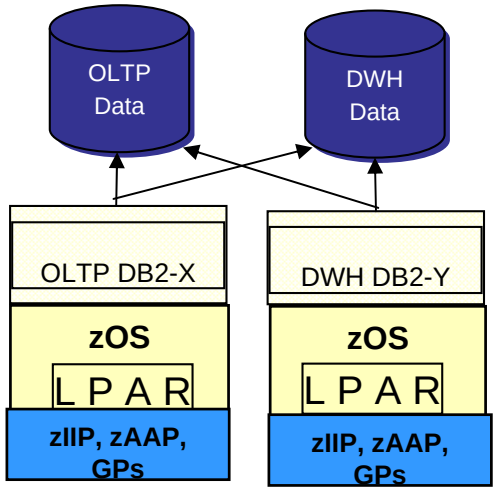

2 LPARs / 2 subsystems – data sharing

- Same benefits as prior topology, plus…
- Higher availability
- Shared processor capacity when LPARs located on same server

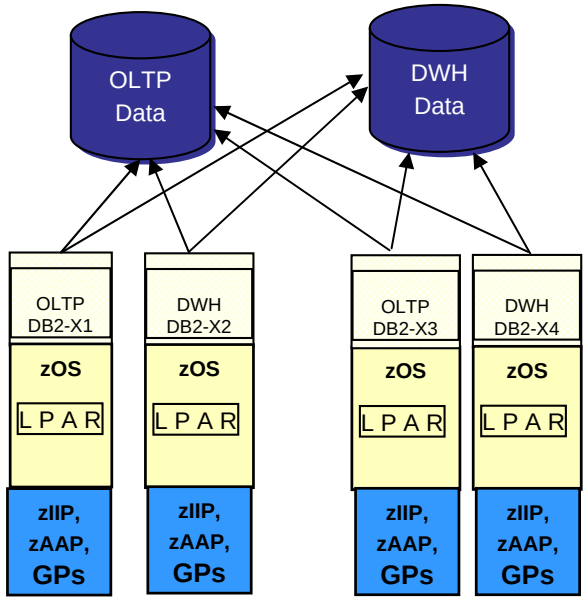

4 LPARs / 4 subsystems – single data sharing group

- Same benefits as prior topology, plus…
- Highest availability
- Work does not stop completely in case of a DB2 member planned or unplanned outage

Example – Credit Card Self Service

- **Concept: Build strategic differentiation by providing customers access to their own data over the web**

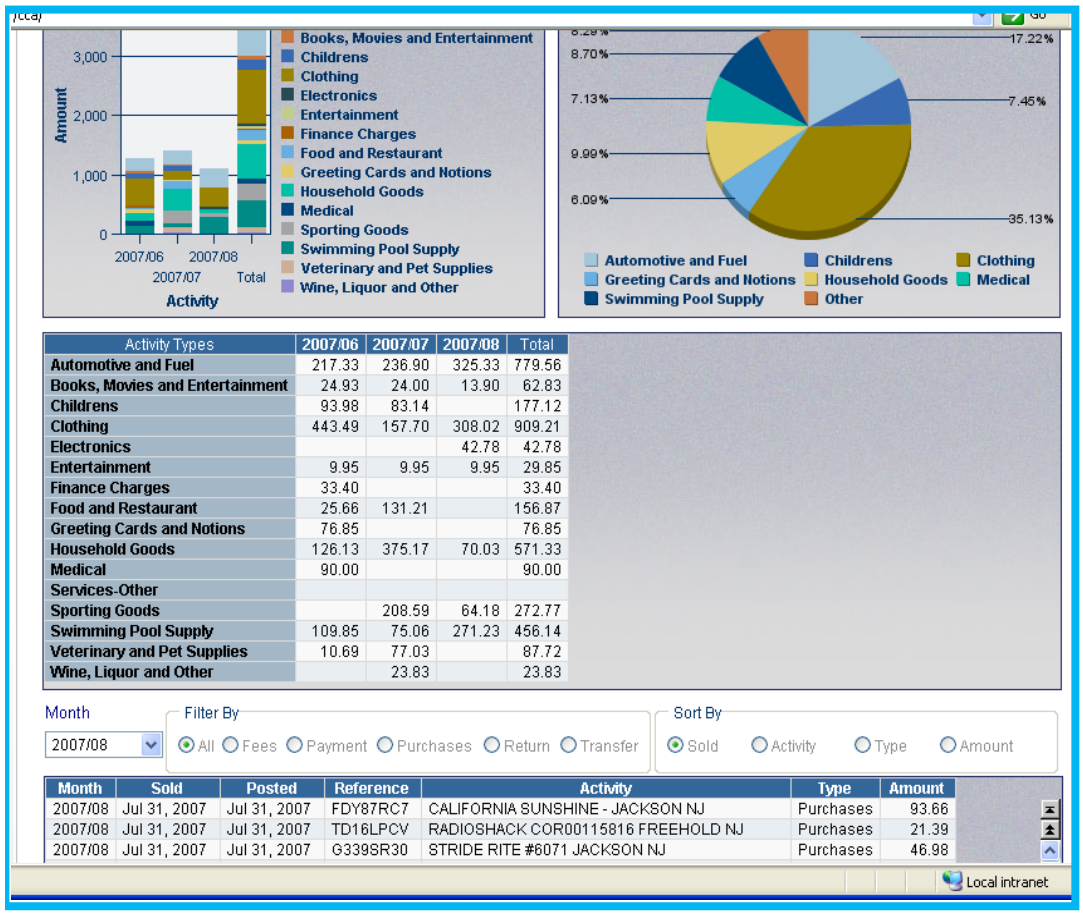

- **Go beyond the transaction level and deliver information like "categorical expense reporting" via grids and trend graphs**
- **Since the information is not "portable" this drives further "stickiness"**
- **Establish differentiation from competitors**

Adding the Query Member… Things to Consider

- **zParms set best for complex query processing**
- **Locking Considerations**
  - If queries are executed with UR from the Query Member, then there is no impact to the OLTP Member workload.
- **Objects are already shared / GBP Dependent**
  - Query Member access does not create additional GBP access from OLTP Member… you can't become "more dependent" than you already are
  - PCLOSET – time interval DB2 checks to downgrade to non GPB Dep (Default 10)… increase it to keep a steady state (avoid surprise) of GBP DEP, otherwise impacts OLTP (hiccup effect)

Adding the Query Member… Things to Consider

- **Objects are intentionally not GPB Dependent**
  - UR on Query Member does not change rules of GBP Dependency
  - Will become dependent and experience data sharing overhead

- **Populate DW Schema**
  - Off hours, throughout day, small chunks via Q-Rep on single data sharing group (log data sets, not buffers)
  - Or, via the OLTP member to avoid incurring GBP Dependency

- **Accessing both types of data**
  - Depending on the application, you can chose the side to invoke the application
    - DW data versus OLTP data… DW data for last month, OLTP for fraud alert in last few hours… process could direct users in different ways

Study Architecture

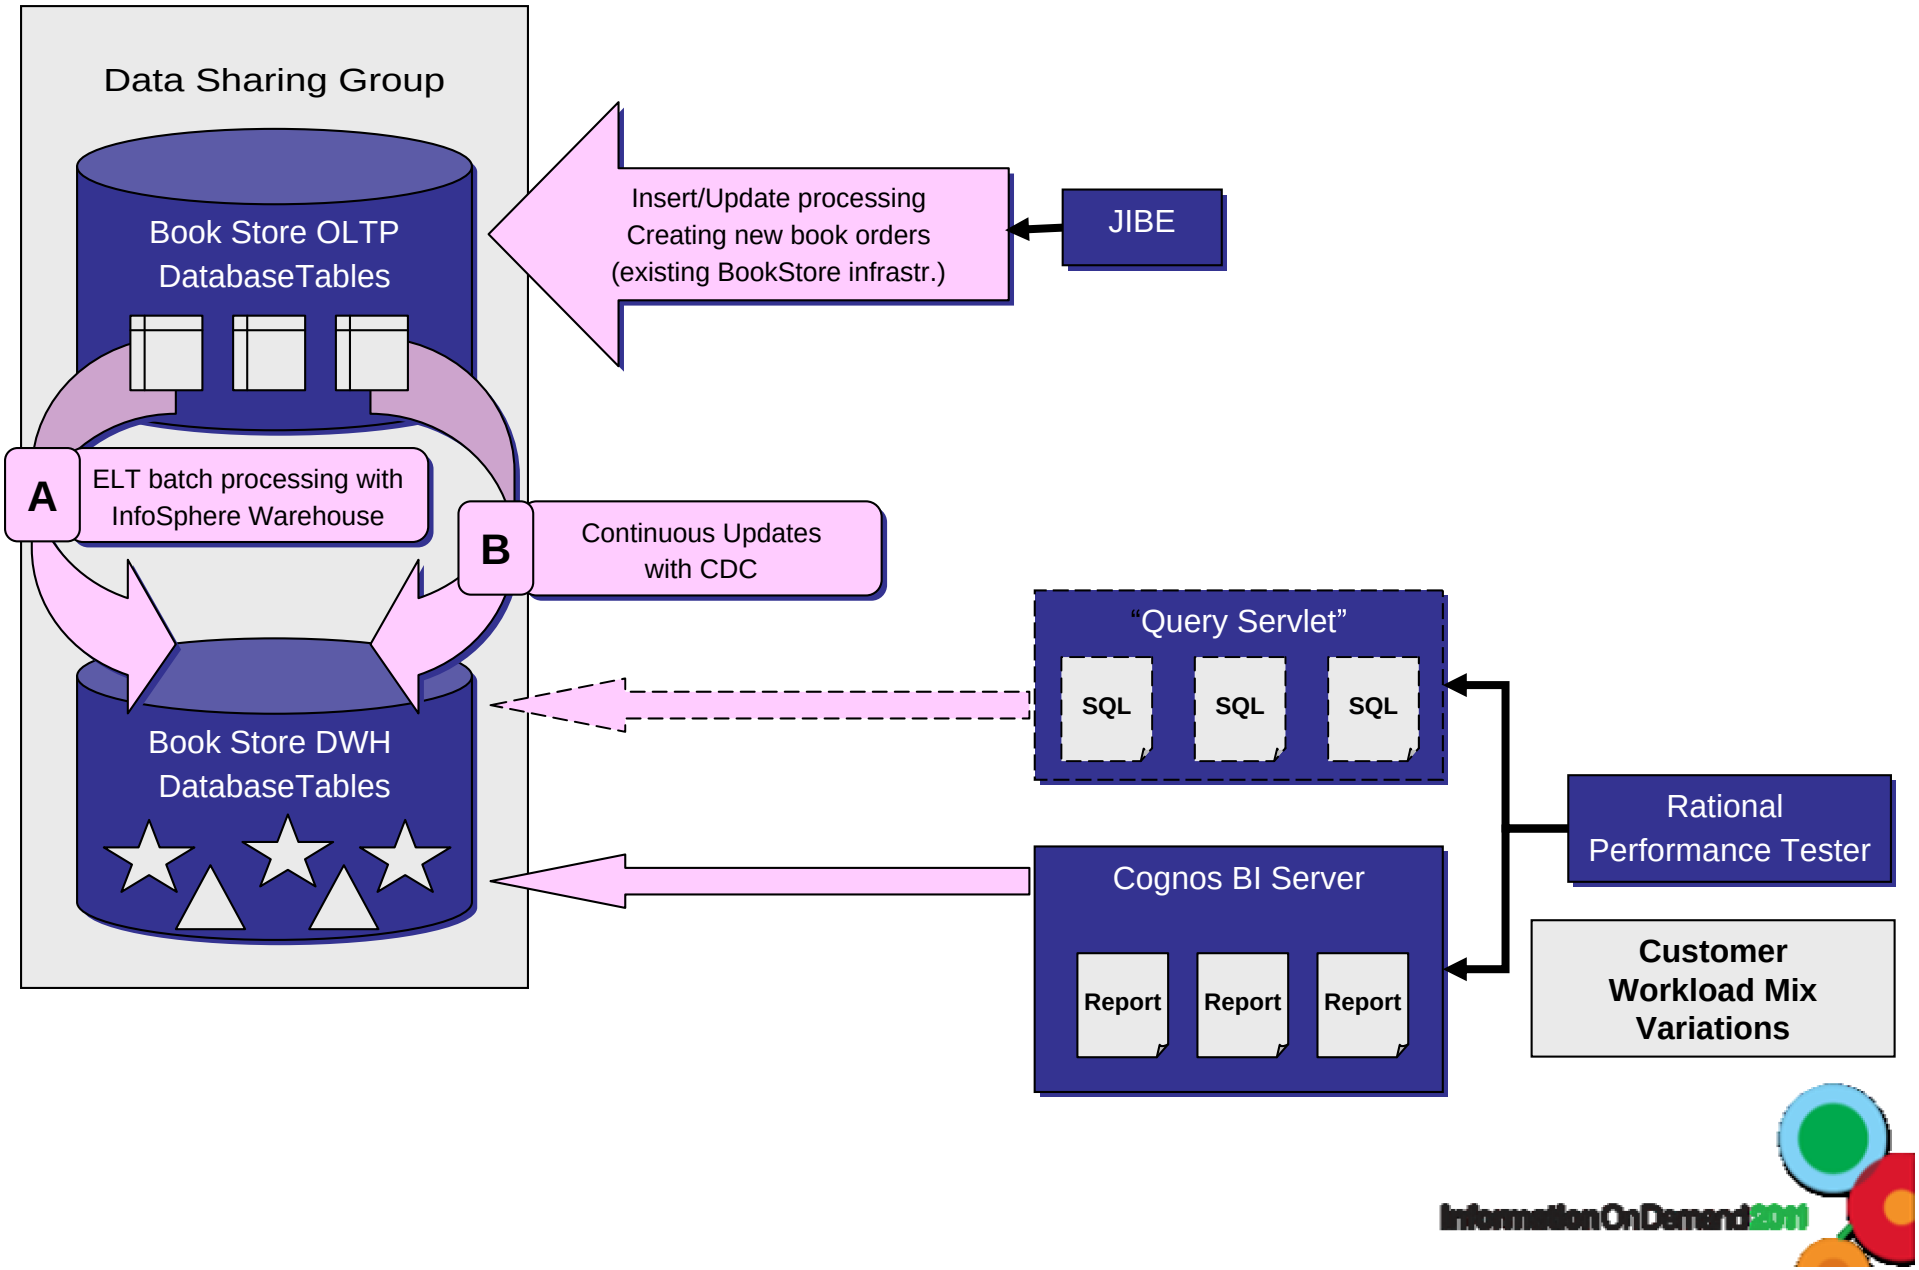

Workload Distribution

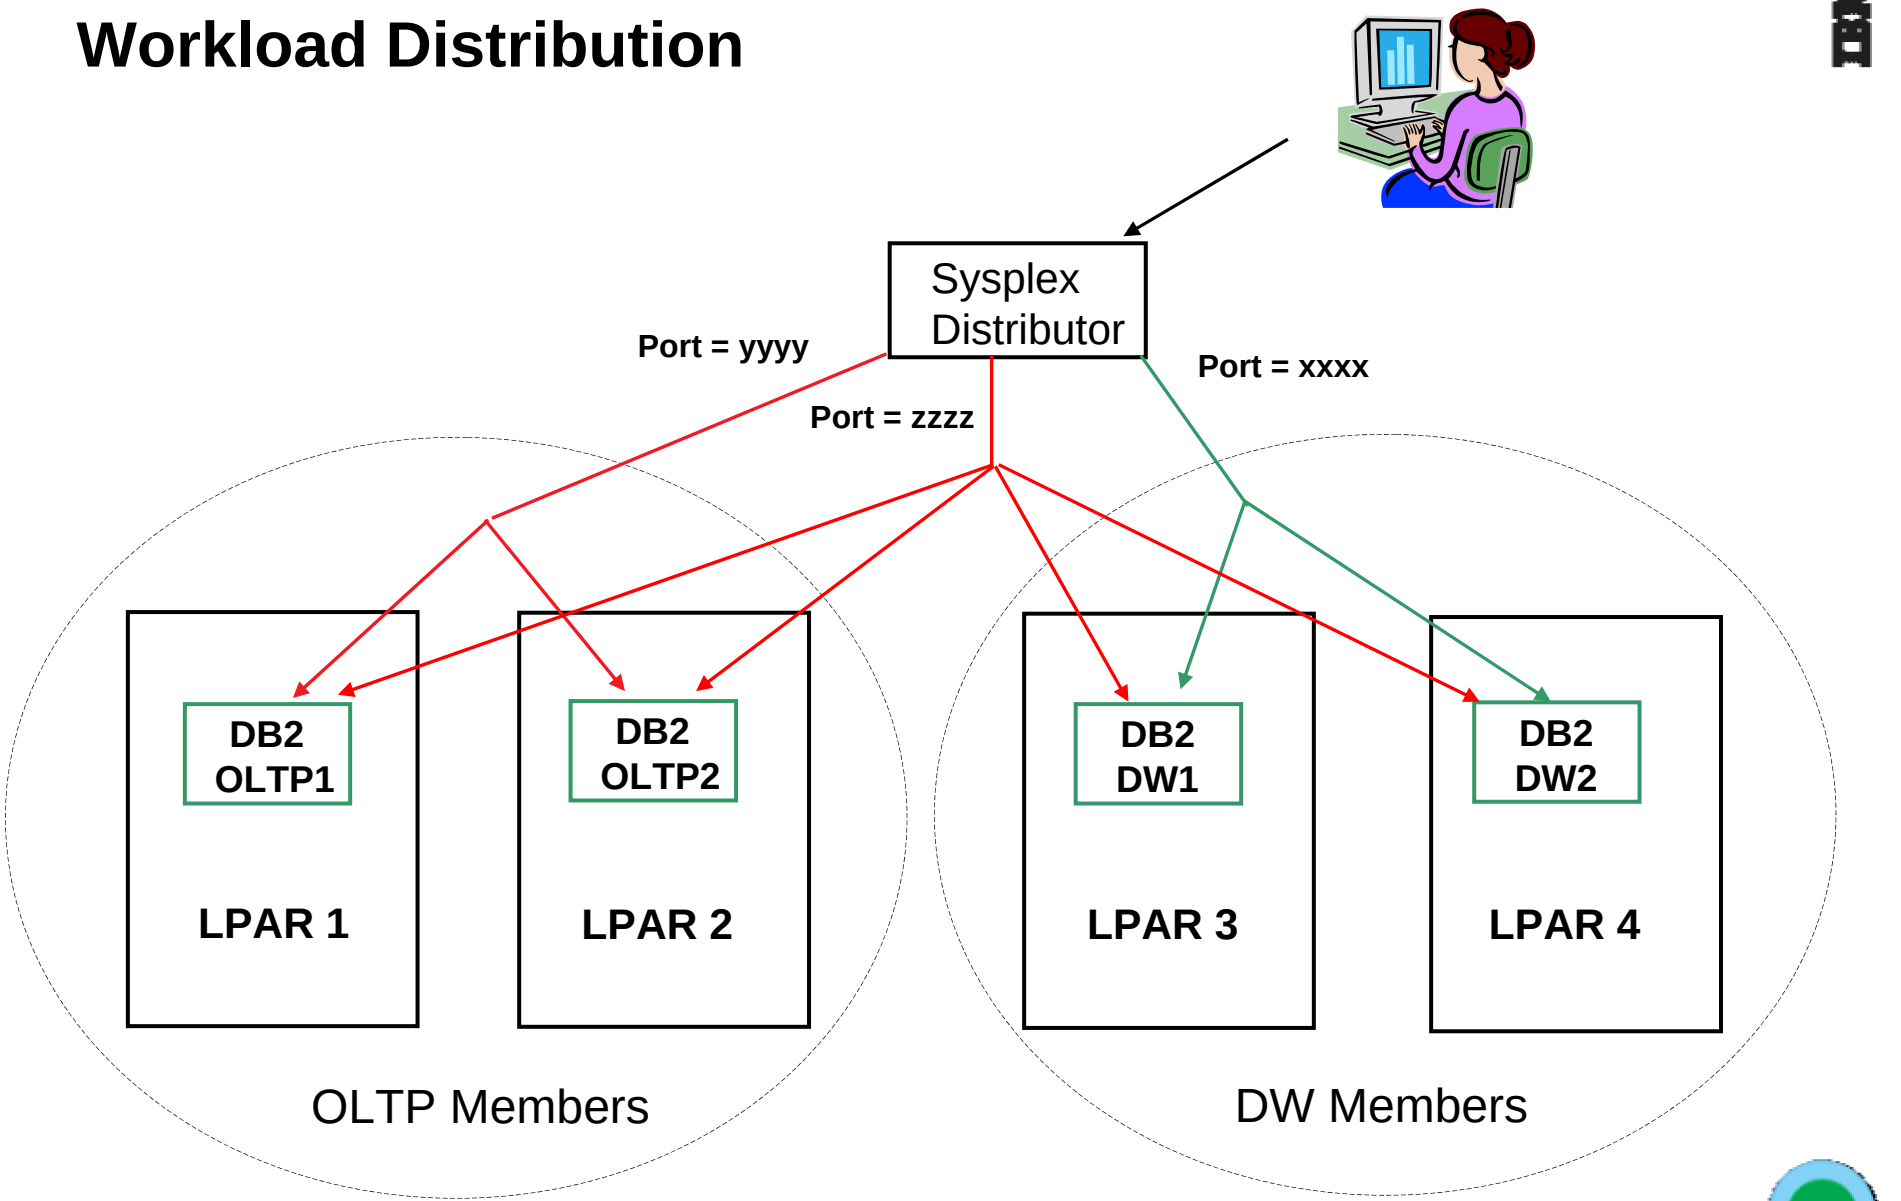

Workload Management

- **Gartner Magic Quadrant Feb, 2010**
  - "Through 2012, **mixed workload** performance will remain the single most important performance issue in data warehousing"
  - "As a direct effect of the complex mixed workload, with continuous loading and the increase in automated transactions from the functional analytics in OLTP, the transactional OLTP DBMSs may be able to erode the performance edge that was formerly attributed to specialized data warehouse DBMS solutions"
- **Bloor Research - 2008**
  - "Data warehousing environments are increasingly required to host these short running queries in addition to their traditional role of supporting analytics and business intelligence. These different types of queries have different requirements: it is the function of mixed query workload management to ensure that the latter do no impinge on the performance requirements of the former"
- **Catterall Consulting - 2010**
  - "Mixed workload performance. This was identified as the number one performance issue for data warehouses. Mixed BI workloads, in which fast-running, OLTP-like queries vie with complex, data-intensive analytic processes for system resources, are becoming common as so-called 'operational BI' gains prominence"

System z WLM provides:

- *Consistent response times for shorter consumption work where end-user is awaiting response*
- *Favor the company's Important $$$ Knowledge workers / Analysts*
- *…while still opening up the warehouse to larger community without fear for monopolization*
- *Manage each workload to it's specific goal*
- *Ability to dynamically change policy for time-based workloads. Ex. ETL / Refresh / Batch Reporting*

Workload Management

$$$$
Service
Center
Reps

Real-time
Dashboards

OLAP

Operational BI

Predictive Analysis

ETL / Refresh

Batch
Reporting

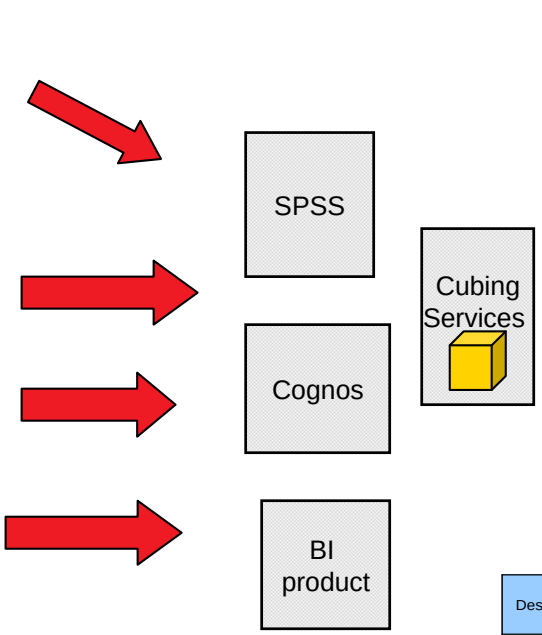

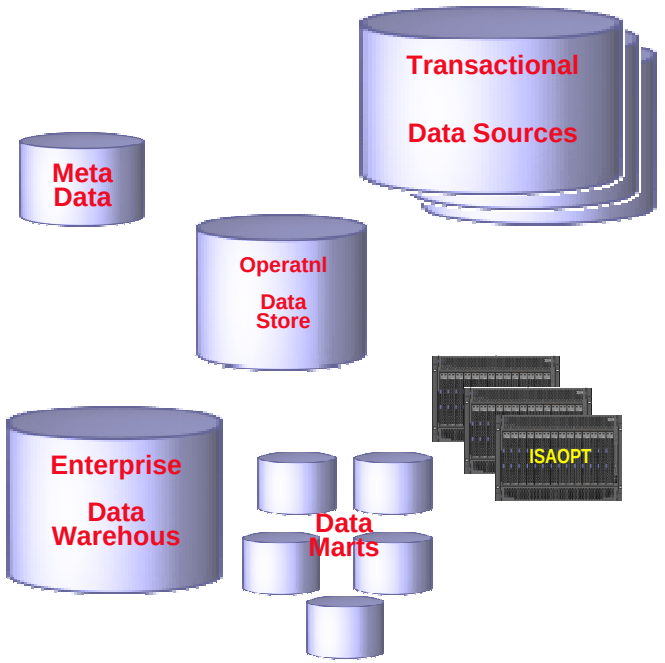

WLM Service Policy

Description	Service Class	Prd	Imp	Duration (SU)	Type	Goal
High Priority BI DDF Queries	DWDDFHI	1	2	25,000	Resp Time	90% < 3 sec
		2	2	100,000	Resp Time	80% < 15 sec
		3	3	1,000,000	Velocity	30
		4	4		Velocity	10
Standard BI DDF Queries	DWDDFST	1	3	25,000	Resp Time	90% < 3 sec
		2	4	500,000	Velocity	10
		3	5		Discretionary	
Scheduled Reports	DWSCHED	1	4	2,000,000	Velocity	30
		2	5		Velocity	10
Refresh	DWREFSH		4		Velocity	10

WLM Classification – Set Client Information

DDF Work Classification Attributes

Attribute	Type	Description
Accounting Information	AI	**Set by Client Info API**: client accounting string
Correlation Information	CI	driver program name by default but application can set
Collection Name	CN	Collection name of the first SQL package accessed by the DRDA requester in the unit of work
Connection Type	CT	Always 'DIST ' for DDF server threads
Package Name	PK	Name of the first DB2 package accessed
Plan Name	PN	Always 'DISTSERV' for DDF server threads accessed via DRDA requesters
Procedure Name	PR	Name of the procedure called if first request in unit of work
Process Name	PC	**Set by Client Info API :** client transaction name
Subsystem Collection Name	SSC	Usually the DB2 data sharing group name
Subsystem Instance	SI	DB2 server's MVS subsystem name
Sysplex Name	PX	Name assigned to sysplex at IPL
Userid	UI	DDF server thread's primary AUTHID
Subsystem Parameter	SPM	**Set by Client Info API** : assigned the concatenation of client userid/workstation name

WLM Classification – Set Client Information

Call the Stored Procedure to propagate the values….

- call sysproc.wlm_set_client_info (client_userid, client_wrkstation, client_application, client_application_accounting)

The flow / different transformations…..

sqleseti - Set client information	JDBC Client Info Properties	SQL Special Registers	DB2 Accounting Data Fields	WLM Classificati
SQL_ATTR_INFO_ACCTSTR	setClientAccountingInformation	CLIENT_ACCTNG	QMDAAINFO (accounting string)	AI
SQL_ATTR_INFO_APPLNAME	setClientApplicationInformation	CLIENT_APPLNAME	QWHCEUTX (end user transaction name)	PC
SQL_ATTR_INFO_USERID	setClientUser	CLIENT_USERID	QWHEUID (end user id)	SPM
SQL_ATTR_INFO_WRKSTNNAME	setClientWorkStation	CLIENT_WRKSTNNAME	QWHCEUW (end user workstation name)	SPM

Note: CLIENT_USERID & CLIENT_WRKSTNAME concatenated in WLM filter ‘SPM’

WLM Classification – Set Client Information

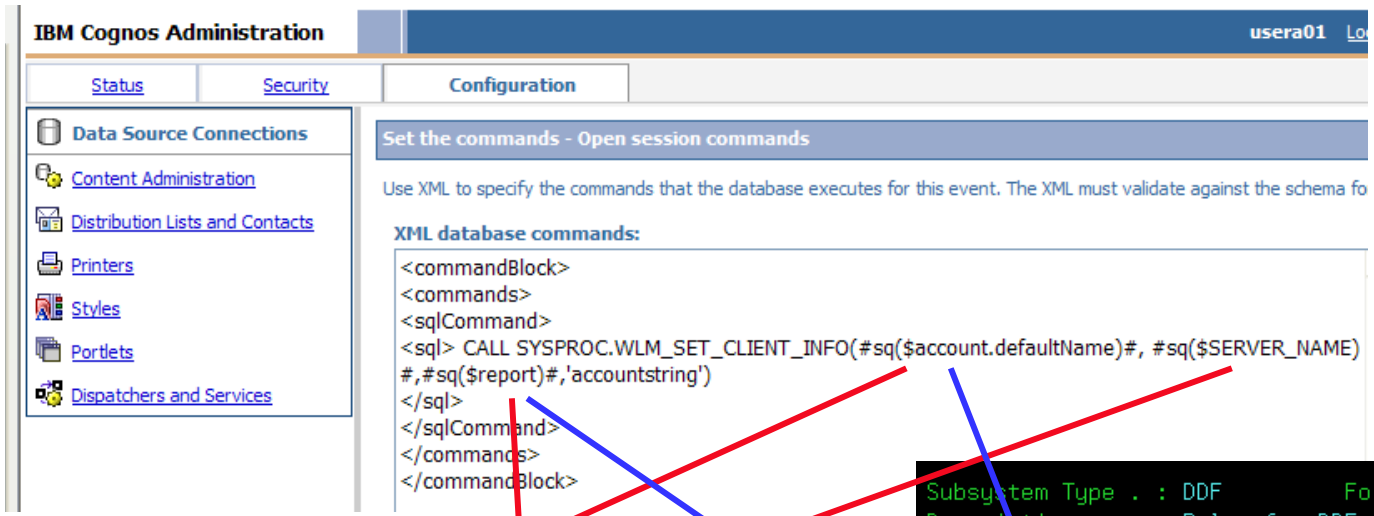

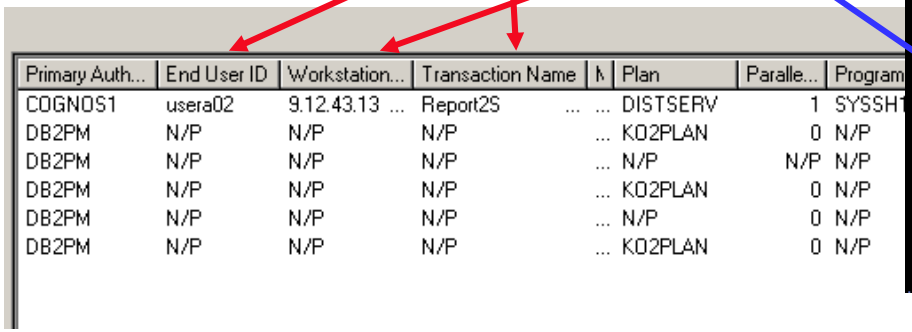

Primary Auth...	End User ID	Workstation...	Transaction Name	N	Plan	Paralle...	Program
COGNOS1	usera02	9.12.43.13 ...	Report2S	...	DISTSERV	1	SYSSH1
DB2PM	N/P	N/P	N/P	...	KO2PLAN	0	N/P
DB2PM	N/P	N/P	N/P	...	N/P	N/P	N/P
DB2PM	N/P	N/P	N/P	...	KO2PLAN	0	N/P
DB2PM	N/P	N/P	N/P	...	N/P	0	N/P
DB2PM	N/P	N/P	N/P	...	KO2PLAN	0	N/P

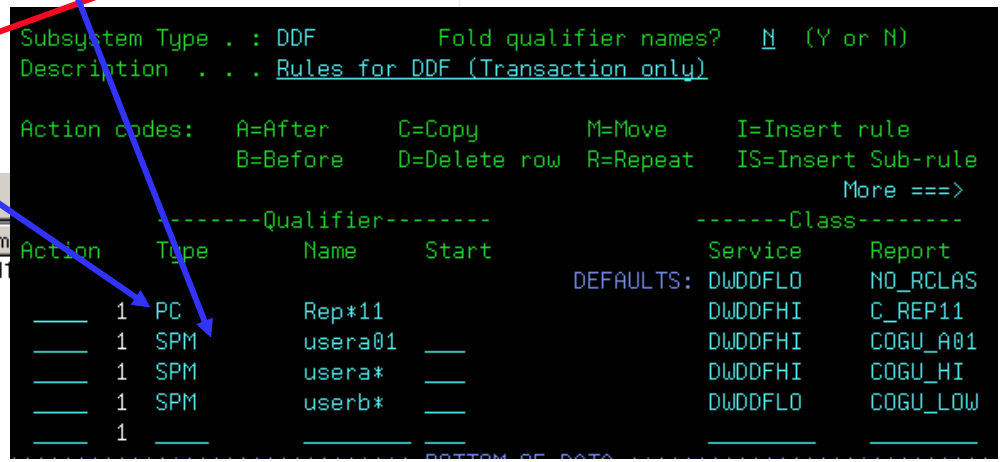

Increased flexibility in resource management & monitoring !!

CRITICAL FOR DATA WAREHOUSE / BI WORKLOADS

Expand WLM Policy to Cover DW Queries

- Collect DB2 accounting data
- Store data in DB2PE Performance Database
- Retrieve data from database and download to spreadsheet
- Use Pivot Table to analyze data

client_transaction	class1_su_cpu	class1_iip_su_cpu	total_class1_su	class1_elapsed
Report 1	5,759	6,715	12,474	0.610031
Report 2	6,885	5,590	12,475	0.591942
Report 3	6,200	6,276	12,476	1.0109722

Workload Analysis

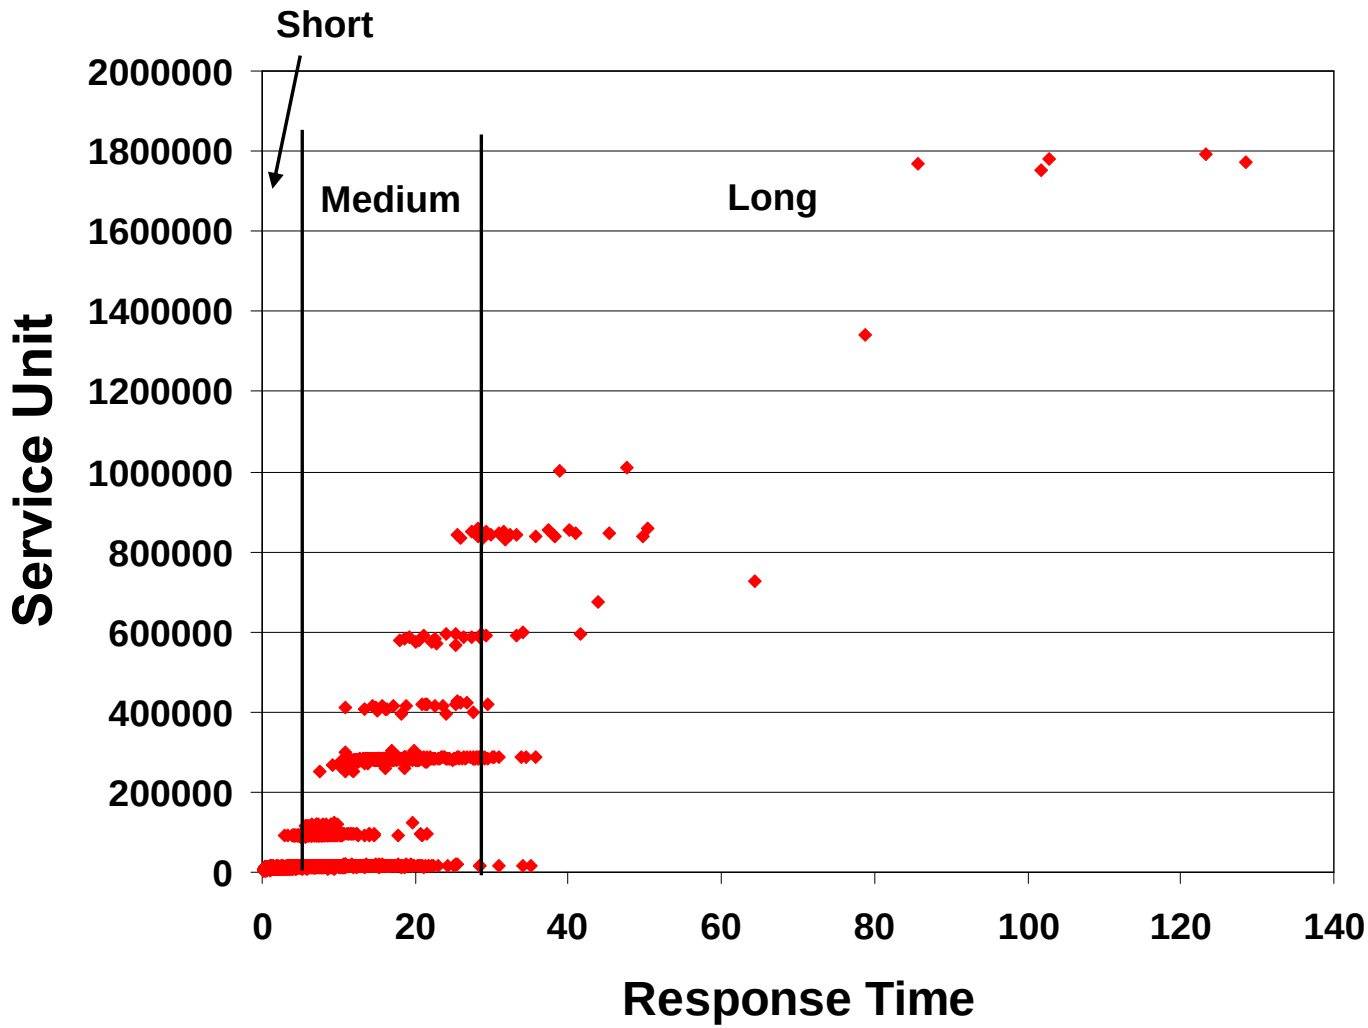

Sample WLM Policy – Dedicated DW LPAR Study

Description	Service Class	Prd	Imp	Duration (SU)	Type	Goal
High Priority BI DDF Queries	DWDDFHI	1	2	25,000	Resp Time	90% < 3 sec
		2	2	100,000	Resp Time	80% < 15 sec
		3	3	1,000,000	Velocity	30
		4	4		Velocity	10
Standard BI DDF Queries	DWDDFST	1	3	25,000	Resp Time	90% < 3 sec
		2	4	500,000	Velocity	10
		3	5		Discretionary	
Scheduled Reports	DWSCHED	1	4	2,000,000	Velocity	30
		2	5		Velocity	10
Refresh	DWREFSH		4		Velocity	10

Query Parallelism

- Goal: high concurrency to lower elapsed times
- High number of table partitions → high degrees of parallelism
- High degree of parallelism → more virtual storage, consume processing power
- Significant amount of thread related storage is still below the bar (V9)
- zParm PARAMDEG is available to cap the maximum degree per query
  - Recommend 2 x number of CPs as starting value
- Hidden zParm, SPRMPTH, defines the threshold of enabling parallelism
  - Noticeable query parallelism overhead for short running queries
  - 120 milliseconds as the default
  - Consider increasing it to a few seconds (e.g., 5 seconds)
- For tables completely pinned in buffer pools in memory:
  - Do not turn off (VPSEQT = 0) sequential prefetch
    - Some customers do this to increase efficiency, don't do this because….
  - Parallelism enablement (VPPSEQT) is based on VPSEQT setting
- Apply APAR PM06953 / PTF UK58871 / RSU1012 for DB2 9 (single enclave support for CP query parallelism)

ETL

Data Sharing Mode

- Direct data movement from OLTP DB to DW DB if transformation can be processed via SQL
- Significantly faster than using intermediate staging area
- Watch out for lock escalation
  - Commit often
  - LOCKMAX 0
  - LOCK TABLE EXCLUSIVE
- Group buffer pool dependency if refresh and queries run on different members
- Triggers
  - Direct updates to DW DB but elongate transaction response times

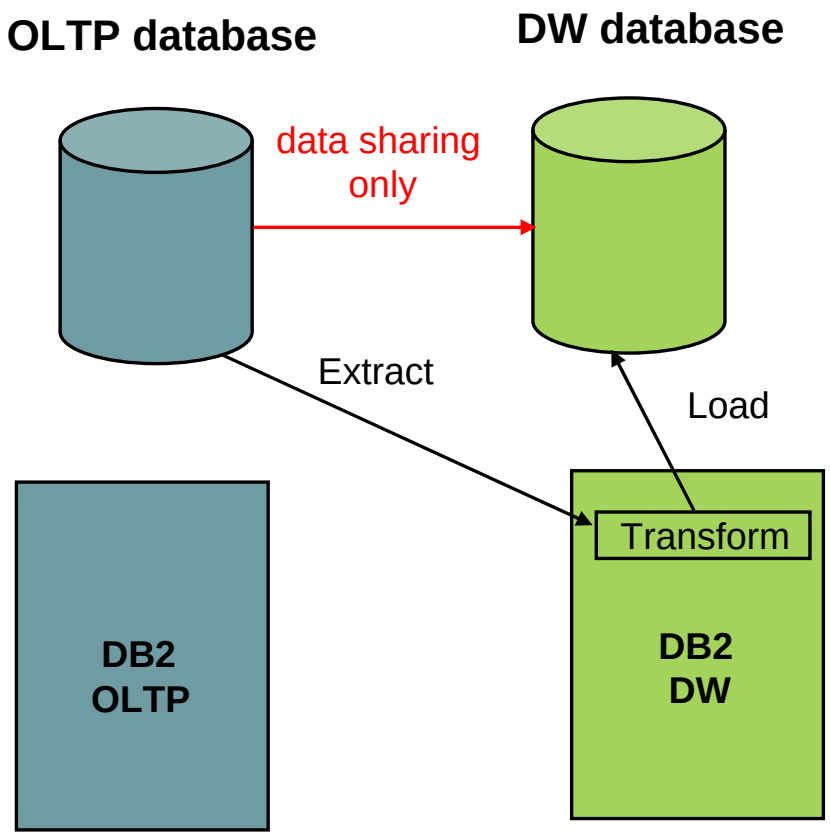

Enhanced Query Capability

Data Sharing Mode

- Problem statement – provide a list of recommended books based on current book browsed
- Single query accessing OLTP and DW databases in the same unit of work:
  - Fetch recommended books from pre-built MQT in DW DB
  - Obtain recent purchases by the same customer from OLTP DB
  - Eliminate duplicates prior to sending result set back to user
- Queries running in DW LPAR does not take away CPU cycles from OLTP transactions
- Read only access minimizes database contention with OLTP transactions

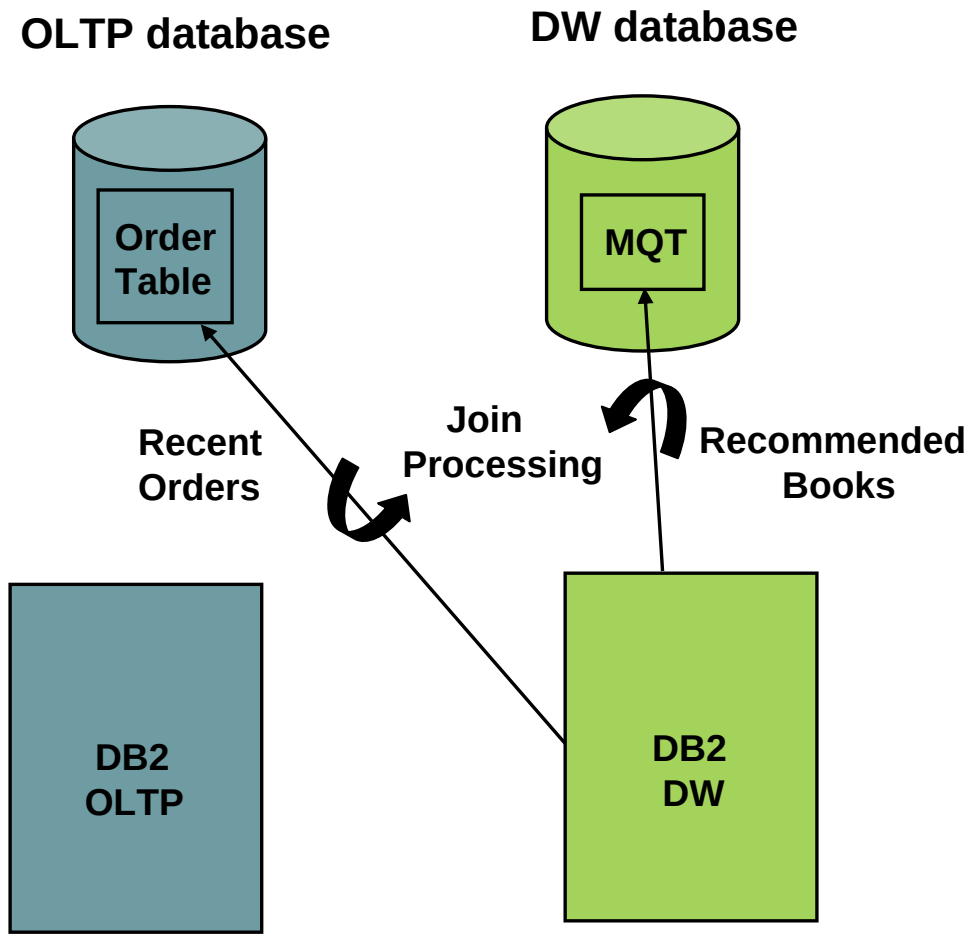

Spike OLTP Workload

Left bars – baseline
Right bars – spike OLTP workload

Throughput

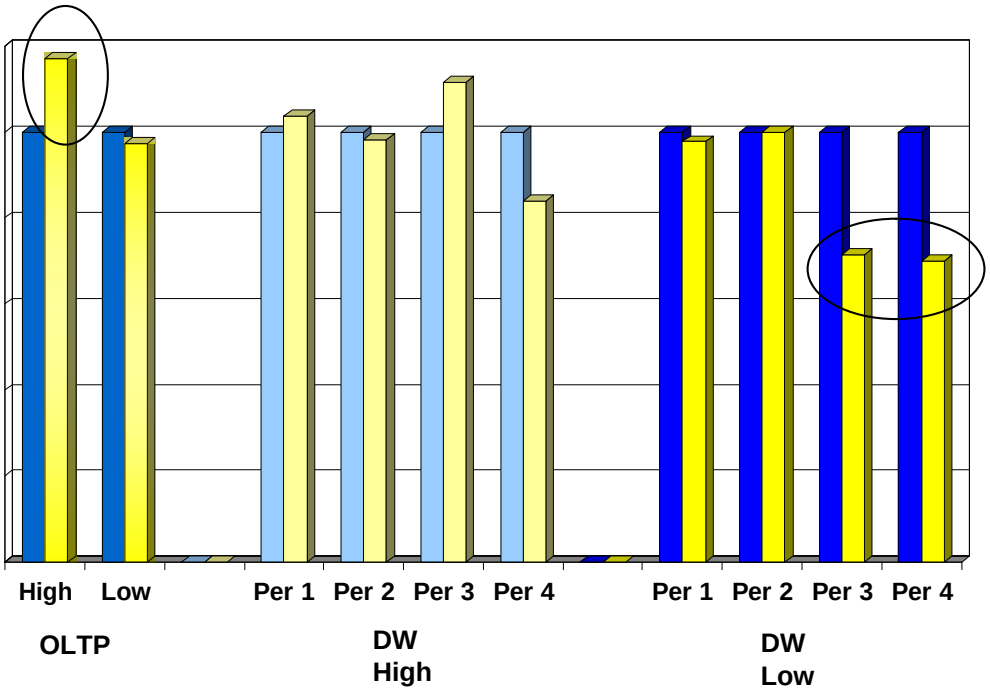

Response Time

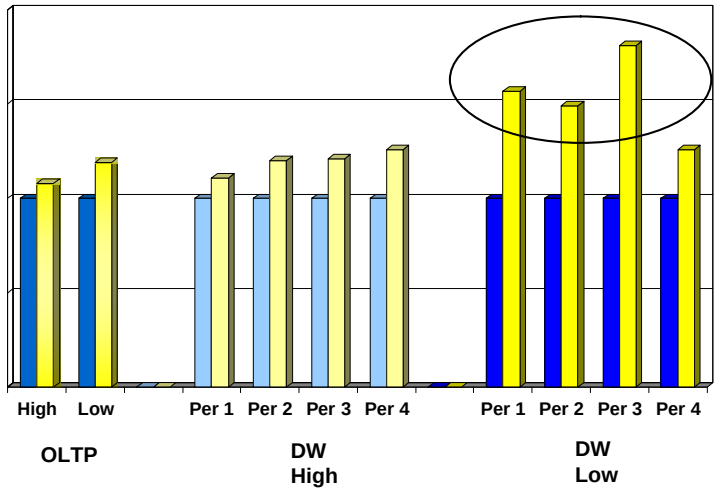

CPU Utilization

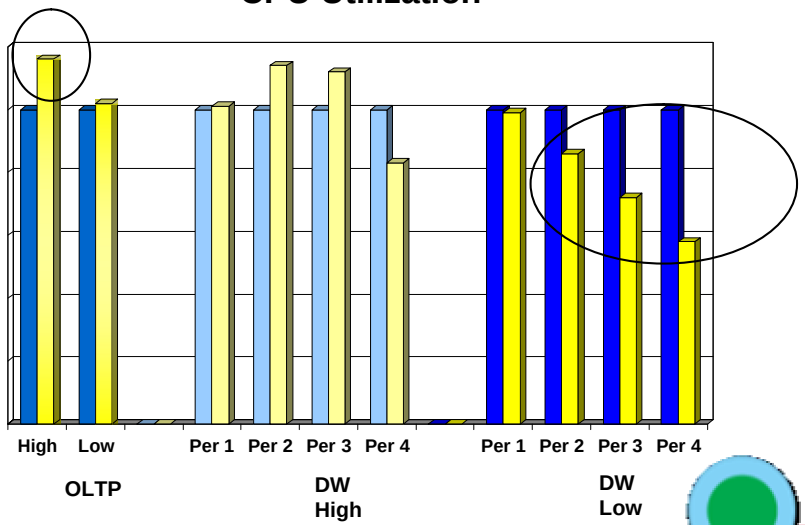

Spike DW Workload

Left bars – baseline
Right bars – spike DW workload

Throughput

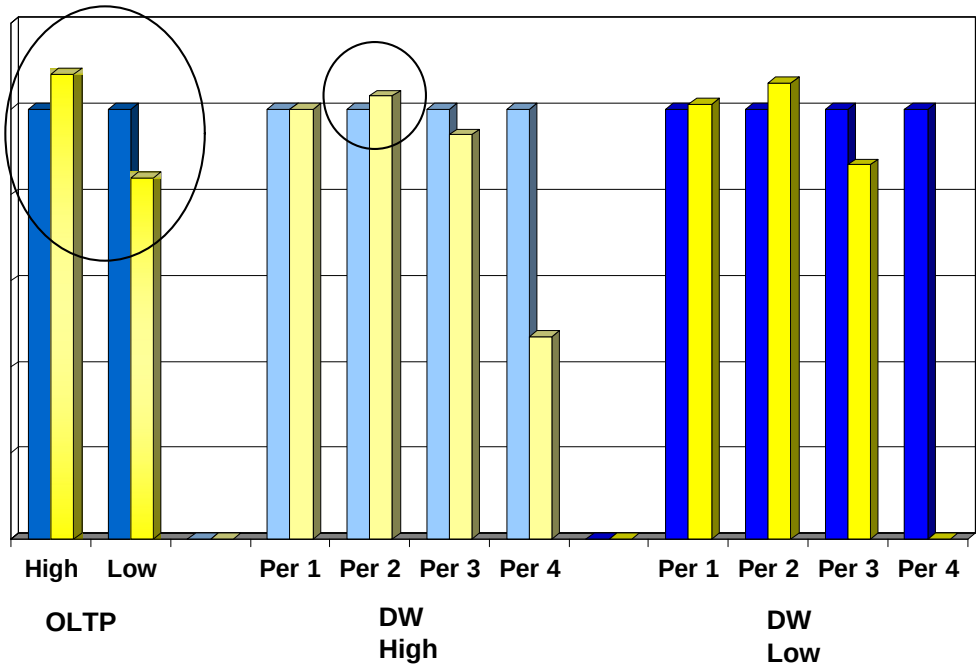

Response Time

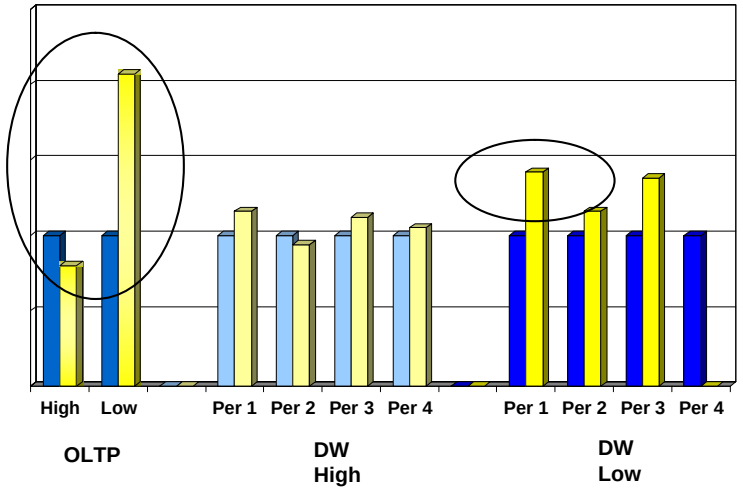

CPU Utilization

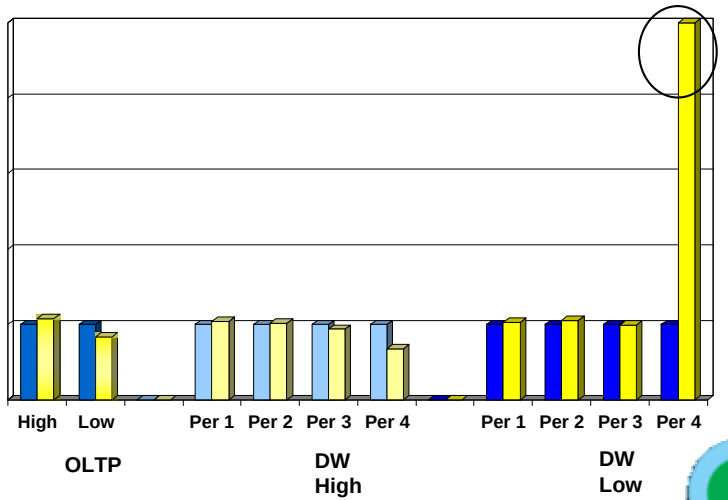

IIPHonorPriority

- Setting it to NO means z/OS will try to contain all zIIP eligible workloads on zIIPs.
- Minimize usage of general processors to reduce cost
- Potential degradation in response time and throughput if zIIPs are highly utilized
- What does it mean in a co-location environment?

IIPHonorPriority = NO

Left bars – baseline
Right bars – IIPHonorPriority = NO

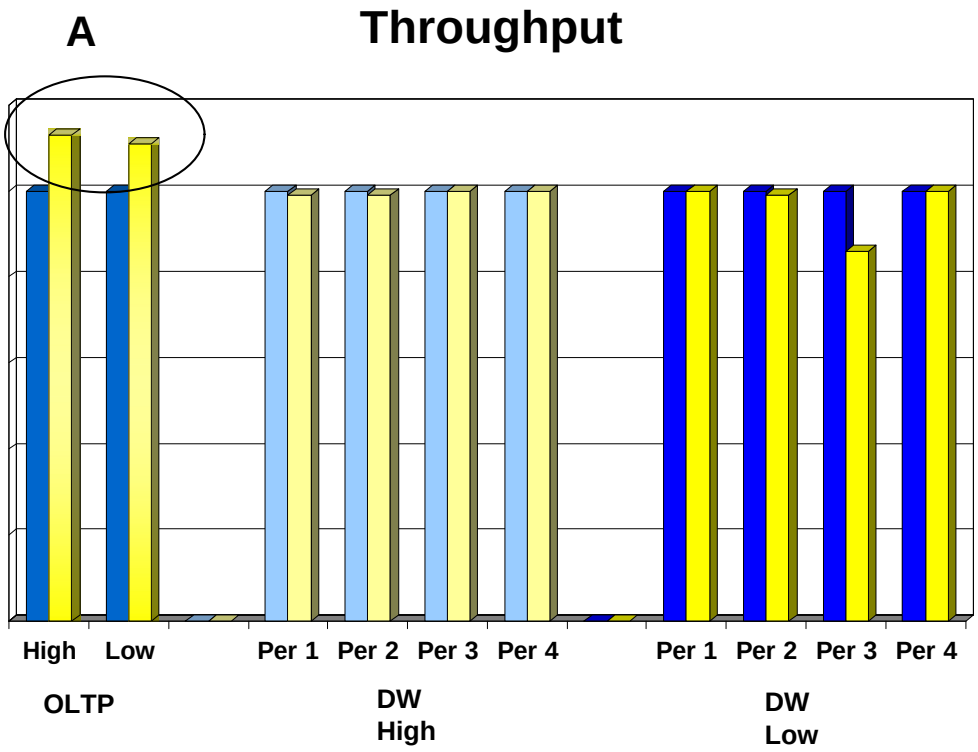

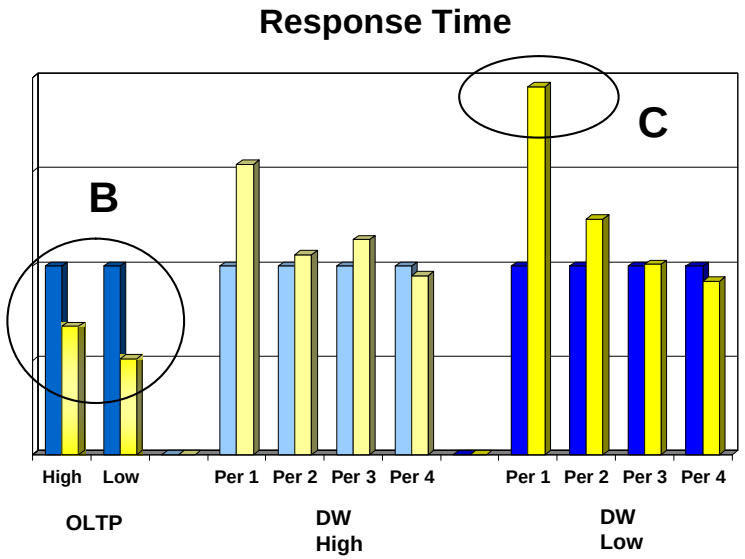

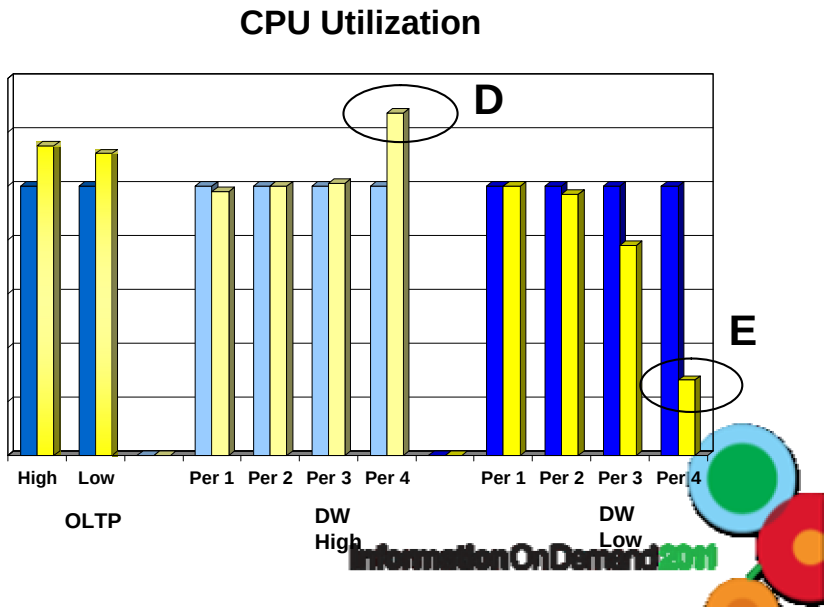

Websites for Additional Information:

IBM Smart Analytics System 9700 Webpage:
http://www.ibm.com/software/data/infosphere/smart-analytics-system/

IBM Smart Analytics Optimizer Webpage:
http://www.ibm.com/software/data/infosphere/smart-analytics-optimizer-z/

Data Warehousing and Analytics
http://www.ibm.com/software/data/infosphere/data-warehousing/

Data Warehousing and Business Intelligence on System z
http://www.ibm.com/software/data/businessintelligence/systemz/

Terabyte Club for System z BI customers
http://www.ibm.com/software/data/businessintelligence/systemz/terabyte-club.html

Data Governance on System z
http://www.ibm.com/software/data/db2imstools/solutions/compliance.html

Data Warehouse Community in the World of DB2 for z/OS.
http://db2forzos.ning.com/group/datawarehousebusinessintelligenceonsystemz

Some Key Redbooks

- Getting Started with the IBM Smart Analytics System 9600
  - http://www.redbooks.ibm.com/abstracts/sg247902.html?Open
- Enterprise Data Warehousing with DB2 9 for z/OS
  - www.redbooks.ibm.com/abstracts/sg247637.html?Open
- 50 TB Data Warehouse Benchmark on IBM System z
  - www.redbooks.ibm.com/abstracts/sg247674.html?Open
- Housing Transactional and Data Warehouse Workloads on System z
  - www.redbooks.ibm.com/redpieces/abstracts/sg247726.html?Open
- InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence
  - www.redbooks.ibm.com/abstracts/sg247813.html?Open
- DB2 for z/OS: Data Sharing in a Nutshell
  - www.redbooks.ibm.com/abstracts/sg247322.html?Open
- System Programmer's Guide To: Workload Manager
  - www.redbooks.ibm.com/abstracts/sg246472.html?Open

Communities

- **On-line communities, User Groups, Technical Forums, Blogs, Social networks, and more**
  - Find the community that interests you …
    - **Information Management** **ibm.com**/software/data/community
    - **Business Analytics** **ibm.com**/software/analytics/community
    - **Enterprise Content Management** **ibm.com**/software/data/content-management/usernet.html

- **IBM Champions**
  - Recognizing individuals who have made the most outstanding contributions to Information Management, Business Analytics, and Enterprise Content Management communities
    - **ibm.com**/champion

Thank You!
Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
  - Your smart phone or web browser at: iodsmartsite.com
  - Any SmartSite kiosk onsite
  - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech

Acknowledgements and Disclaimers:

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.