


IBM DB2 for z/OS Data Warehouse Performance

Presenter: Mike Biere
IBM WW Marketing Mgr.
IBM Silicon Valley Lab
mbiere@us.ibm.com
Session: 006



IBM Information
ON Demand 2010

INFORMATION-LED
TRANSFORMATION



IBM

January 21 - SINGAPORE • January 26 – MALAYSIA • January 28 - THAILAND

Key Objectives

- Understand the value of the System z based data warehouse and business intelligence solution
- Introduce the business intelligence and data warehousing on System z solution
- Describe in some technical depth the key value points that the solution provides
- Provide additional sources of information and contact points to further explore the System z solution

Objective Number 1

Understand the value of the System z based data warehouse and business intelligence solution

Leverage your investment in System z

Help build a smarter planet

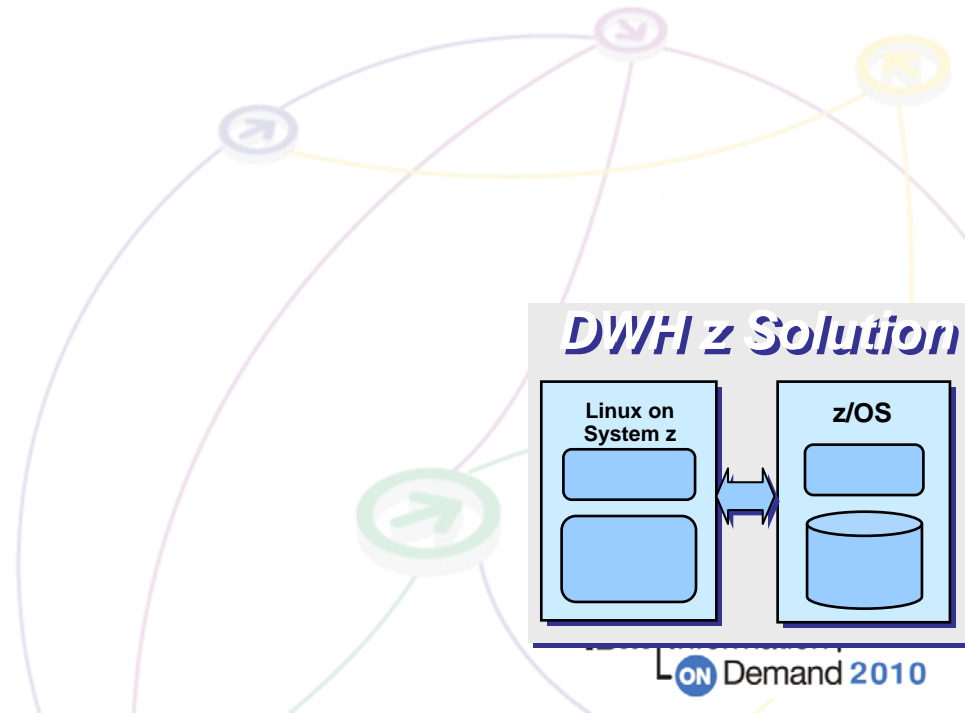
Synergize with System z, z/OS and DB2

Deliver information to your business with Mission Critical Business Intelligence

Leverage your investment in System z

- Infrastructure
 - Hardware and Software
 - Parallel Sysplex
- Processes
 - Development
 - Change Management
 - Security
 - Backup and Recovery
 - Business Continuity
- People
 - Skills
 - Business Knowledge

Take what you have
and make good use of it.




Help build a smarter planet

→ System z is GREEN

- System z utilizes 65% less power than the comparable number of distributed processing power
- Less power means a smaller carbon footprint
- Less hardware means less natural resources utilized and less waste generated when its useful life is over

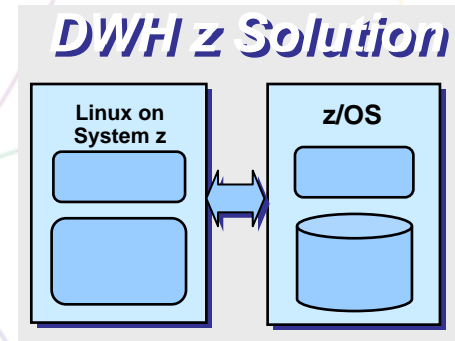
→ System z Saves Green \$\$\$

- Drive down costs (Labor, Storage, Security Breaches, Outages)
- Driving more workload to a single platform leverages economies of scale




Synergize with System z, z/OS and DB2

- System z, z/OS and DB2 designed with integration in mind since their inception
- Deliver unique value with integrated development and features
- These unique features include, but are not limited to:
 - Data Sharing
 - Industry-Leading Workload Management
 - Data and Index Compression
 - Growing incrementally without changes
- This is a level of cooperation that other platforms cannot achieve



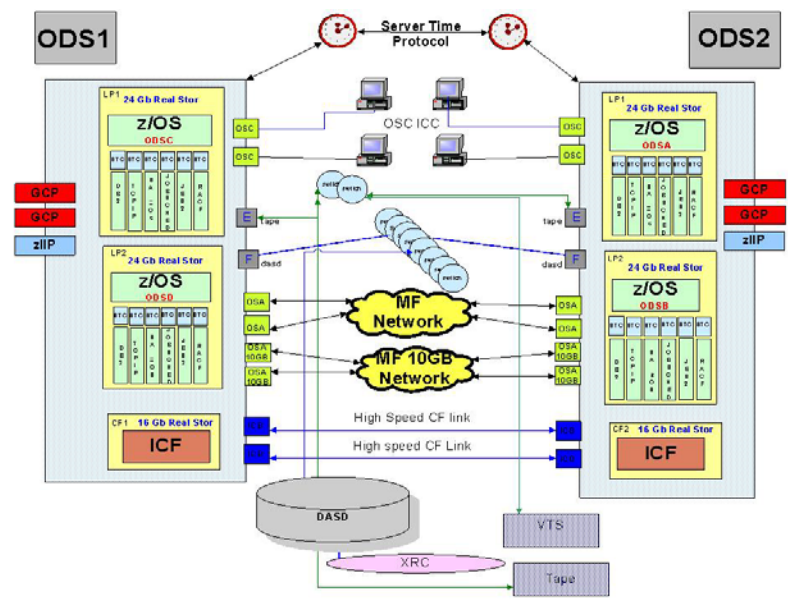
Deliver information to your business with Mission Critical Business Intelligence

- Critical business operations depend on the availability of information stored in the DW
- Embedding BI content in operational applications requires matching Service Level Agreements
- If your business needs the data warehouse to run, then it is Mission Critical and should be on a Mission Critical platform.



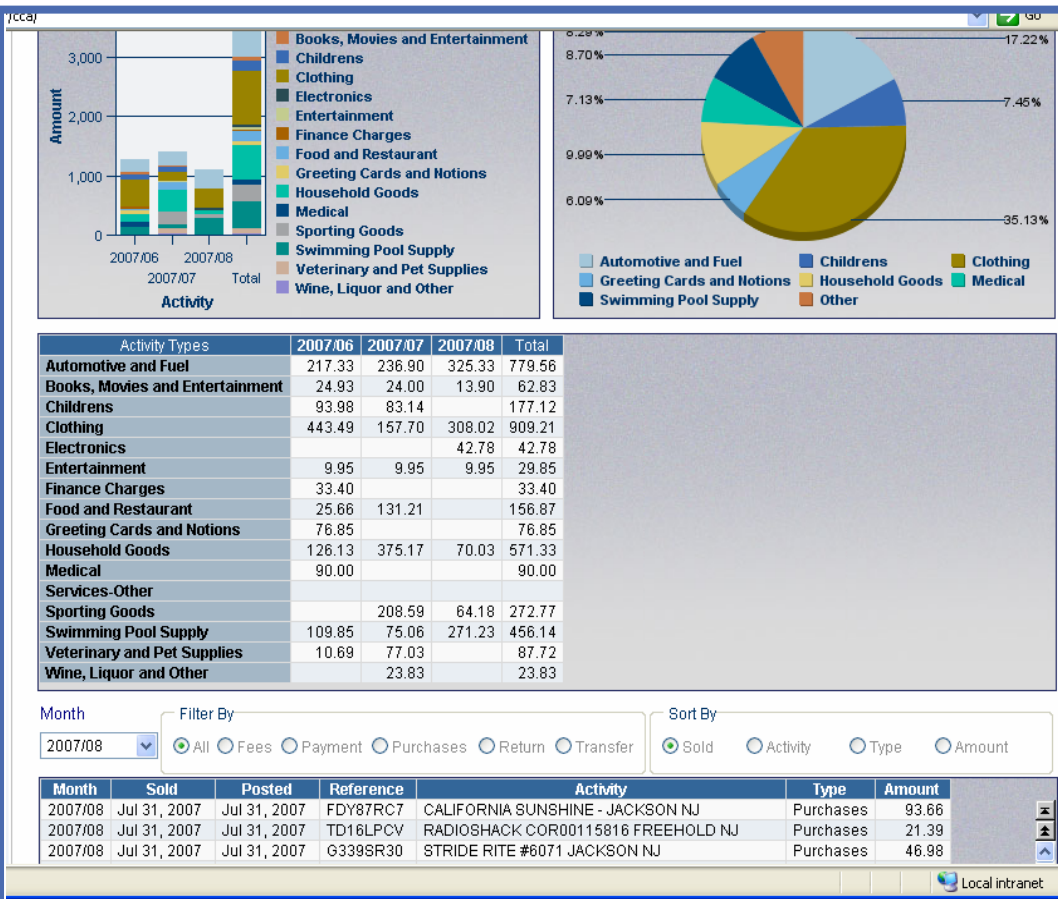
New requirements causing many more to consider System z

- “Real-Time” data access with Dynamic Warehousing
 - Monthly, Weekly, and even Daily feeds are no longer acceptable
 - “Fresh” data is in demand
- Example
 - Customer facing 24/7 Web Access for account information
 - Highly Available Operational Data Store
 - Increase system responsiveness, availability and reduce the load on the back-end systems



Example – Credit Card Self Service


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



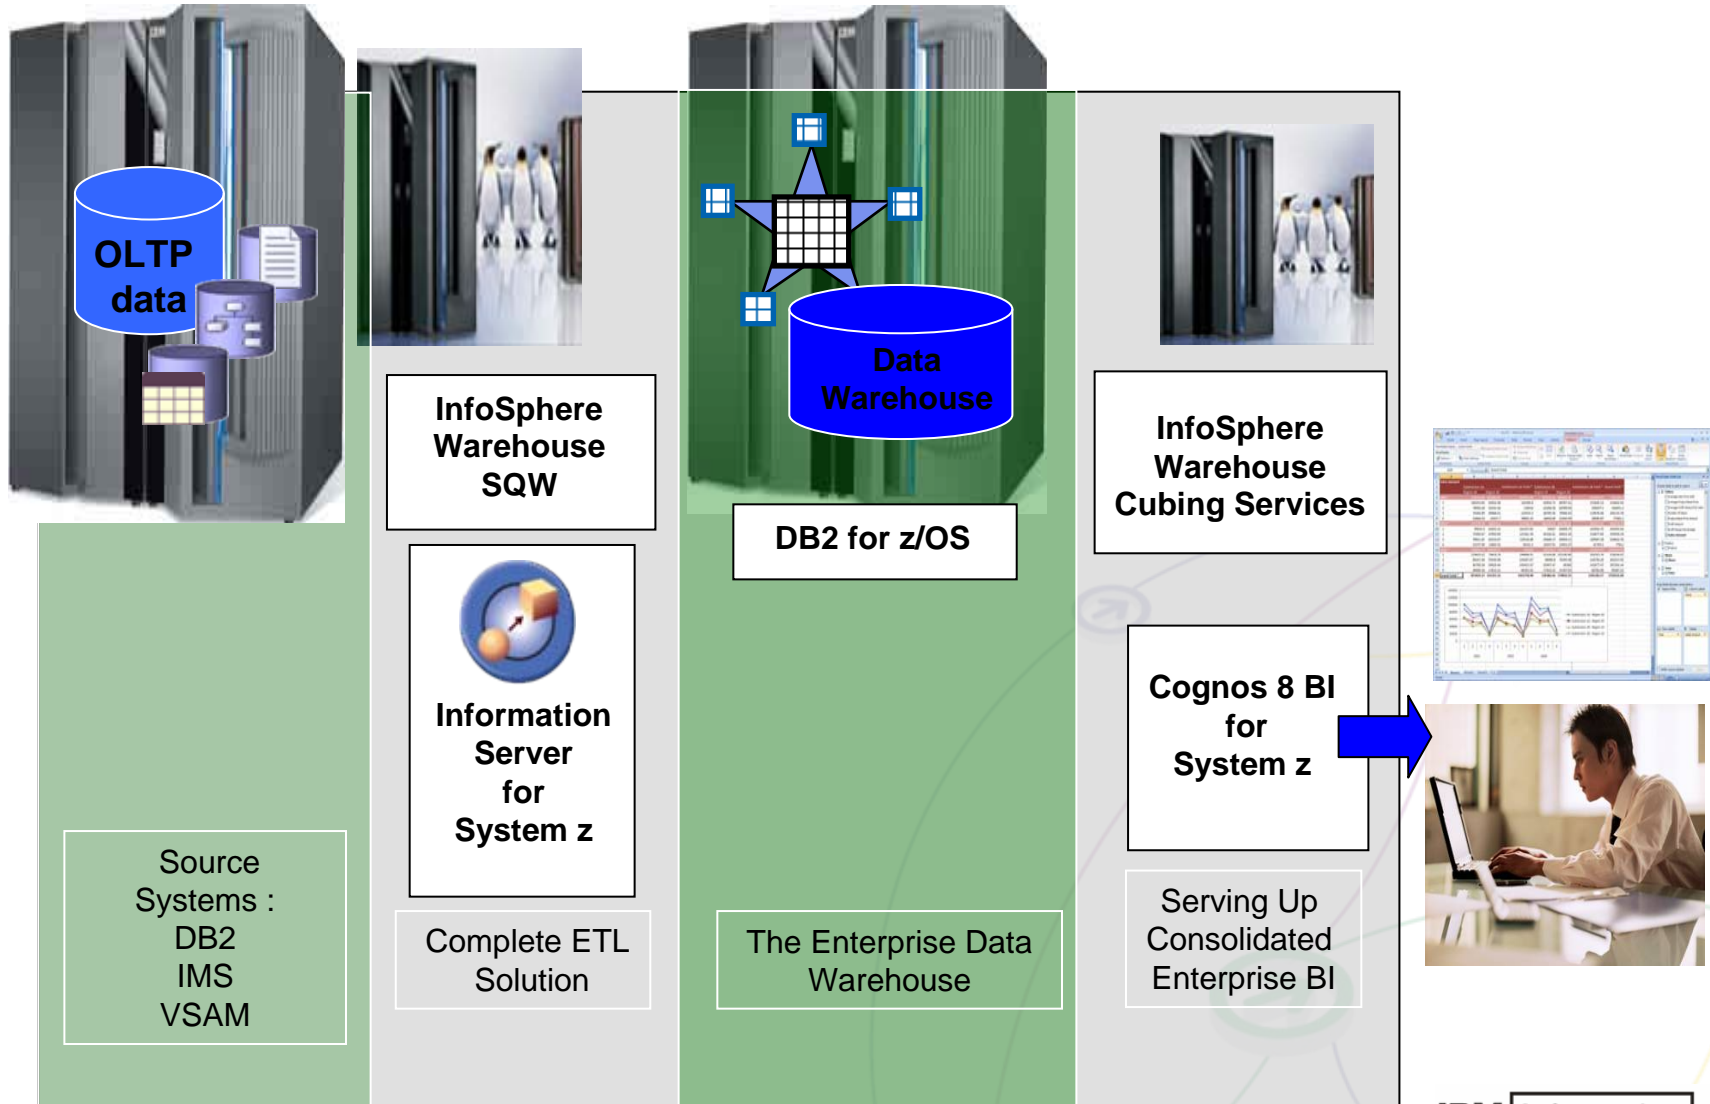
- Go beyond the transaction level and deliver information like “categorical expense reporting” via grids and trend graphs
- Since the information is not “portable” this drives further “stickiness”
- Establish differentiation from competitors

Objective Number 2


An introduction to the System z based business intelligence and data warehouse solution.



Data Warehouse and BI Architecture for System z




Specialty Processors in a DB2 for z/OS Warehouse Solution



ICF – Uniquely allows a Data Warehouse or ODS database to share data 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

IBM zIIP leveraged by DW workloads

1. Business Intelligence applications via DRDA[®] over a TCP/IP connection
2. Complex Parallel Queries
3. DB2 Utilities for Index Maintenance



zIIP Experiences

- Benchmark and internal workloads indicate redirection between 50% and 80% with a typical mix of queries
- RMF and OMEGAMON reports provide projection of possible redirection, amount actually redirected, and the amount of spillage
- Number of zIIP processors can not be more than the number of general processors in a physical server. However, an LPAR can be configured to contain more zIIPs than general processors.

Queries	Elapsed Time	Elapsed Time Min	CPU Total Min	CPU (CP + zIIP) / Elapsed	% Offload to zIIP (Measured)	% Offload to zIIP (Projected)
1	0:53:17	53.28	813.15	15.3	51	80
2	0:01:03	1.05	5.03	4.8	79	79
3	0:49:05	49.08	106.27	2.2	57	65
4	0:40:57	40.95	36.53	0.9	77	77
5	0:41:56	41.93	132.28	3.2	57	73
6	0:23:25	23.42	84.03	3.6	79	79
7	1:03:25	63.42	182.12	2.9	58	67
8	1:08:54	68.90	196.53	2.9	60	72
9	3:50:31	230.52	3436.28	15	51	79
10	0:40:46	40.77	100.98	2.5	56	70
11	0:15:35	15.58	7.03	0.5	64	64
12	0:26:49	26.82	92.40	3.4	79	79
13	0:53:21	53.35	203.22	3.8	46	65
14	0:25:30	25.50	91.72	3.6	72	77
15	0:44:14	44.23	156.03	3.5	72	72
16	0:14:54	14.90	35.18	2	57	57
17	0:28:14	28.23	7.42	0.3	44	44
18	0:30:00	30.00	313.15	10.4	61	79
19	1:01:17	61.28	11.43	0.2	48	48

DB2 VUE: An alternative to a DB2 Monthly License Charge

→ DB2 for z/OS Value Unit Edition

- Offering DB2 for z/OS as One Time Charge (OTC) software for eligible, net new workloads on z/OS
- Net new workloads limited to those that qualify for System z New Application License Charge (zNALC)
- The new workload can only run in a new zNALC partition, machine or sysplex

→ DB2 for z/OS as OTC

- Provides a purchasing option
- Control z/OS stack costs (pay only for what you use)
- Simplified measuring of costs
- May address some release migration timing issues

→ www.ibm.com/software/data/db2/zos/edition-vue.html

IBM Information Server for System z

- ✓ *COMPLETE* All products
- ✓ *CONSISTENT* Same functionality in the Linux for System z products
- ✓ *COST EFFECTIVE* Can leverage lower-cost IFL MIPS with native z/OS data

IBM Information Server for System z

Understand



- InfoSphere** Information Analyzer for Linux on System z
- InfoSphere** Business Glossary for Linux on System z
- InfoSphere** Data Architect

Cleanse



- InfoSphere** QualityStage for Linux on System z

Transform



- InfoSphere** DataStage for Linux on System z
- InfoSphere** DataStage for z/OS
- InfoSphere** DataStage MVS

Deliver



- InfoSphere** Federation Server & Classic Federation
- InfoSphere** Data Event Publisher & Classic Data EP
- InfoSphere** Replication Server & Classic Replication

Platform Services

Parallel Processing Services



Connectivity Services



Metadata Services



- InfoSphere** Metadata Workbench for Linux on System z

Administration Services




Deployment Services



- InfoSphere** Information Services Director for Linux on System z

IBM Cognos 8 BI for Linux on System z



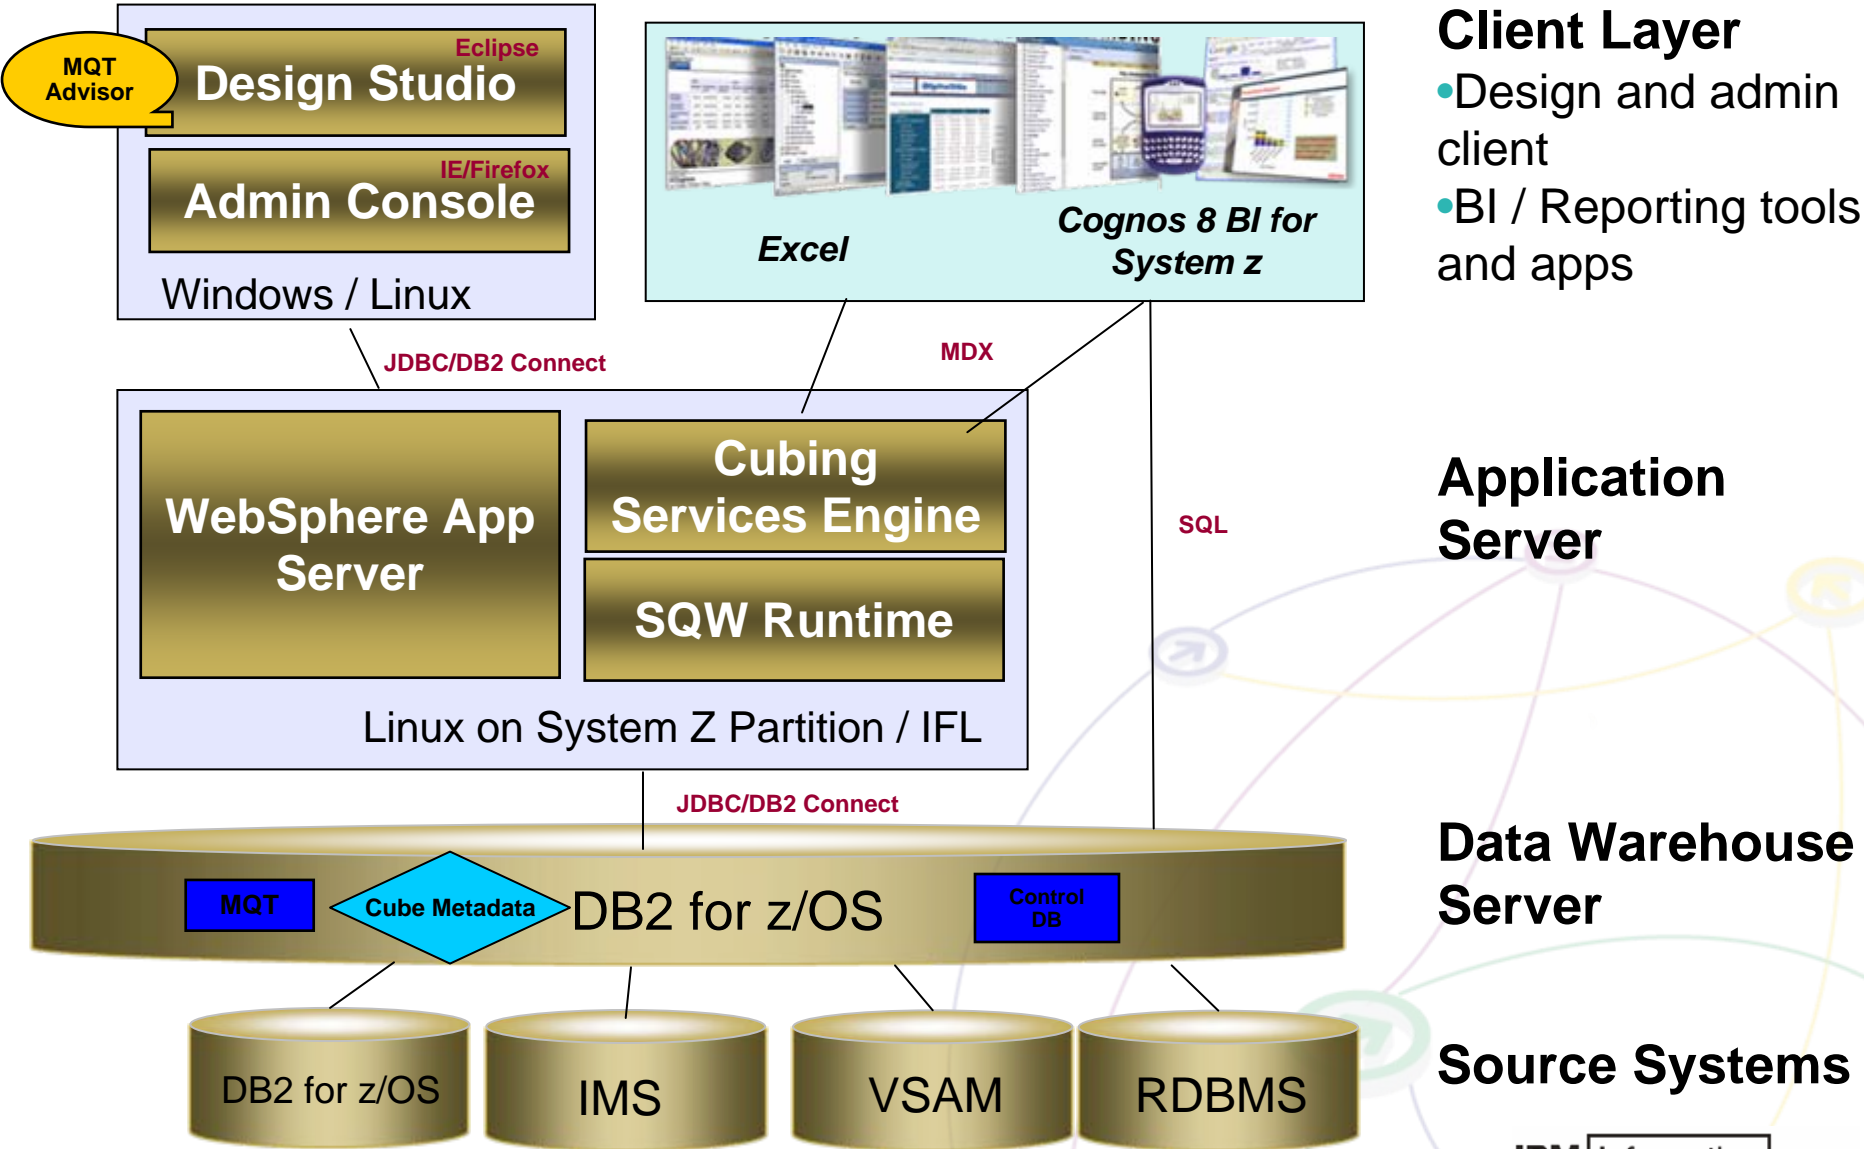
- Adhoc query, reporting and analysis (Query Studio, Report Studio & Analysis Studio)
- Dashboards and charting (Cognos Connection & Report Viewer)
- Event management (Event Studio)
- Integration with Microsoft Office (Go! Office and CAFE)
- Cube building (Transformer)
- Go! Mobile
- Go! Dashboards

Introduction to IBM Cognos 8 BI for Linux on System z

“10 TB Study”

www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101437

InfoSphere Warehouse on System z



Client Layer

- Design and admin client
- BI / Reporting tools and apps

Application Server

Data Warehouse Server

Source Systems

Why Linux on System z for the solution?

- Unify the infrastructure – Get it all “in the box”
 - Manageability and Environmental benefits
- Significant cost savings
 - MIPs charged at IFL rate ... NOT z/OS rate
 - All processing is on Linux for System z, except the z/OS data access
 - Minimizes impact on other z/OS software costs
 - DB2 access qualifies for zIIP specialty engine
- Keep your data access and information integration processes close to your data
 - Eliminate "wire" connectivity – data can flow over hipersockets
 - Simplify, less parts



Objective Number 3


What are some of the key value points that the System z based solution provides?



Platform Options

- IBM is the only company that can provide solutions across multiple platforms
- IBM platform solutions are not intended to be “Me Too”
 - Emphasize the unique functionality on each platform
 - Align advantages of the platform to the business requirements
- System z customers have options
- Four primary DB2 options today
 - Data Warehousing on DB2 for z/OS on System z
 - IBM InfoSphere Balanced Warehouse on System p
 - IBM InfoSphere Balanced Warehouse on System x
 - DB2 for Linux on System z with Data Partitioning

Creating a solid foundation with DB2 for z/OS



- Built on highly reliable hardware and operating system
- Dynamic Configuration Parameters
- Online Utilities
- Online Schema Changes

Latest features that support DW and BI

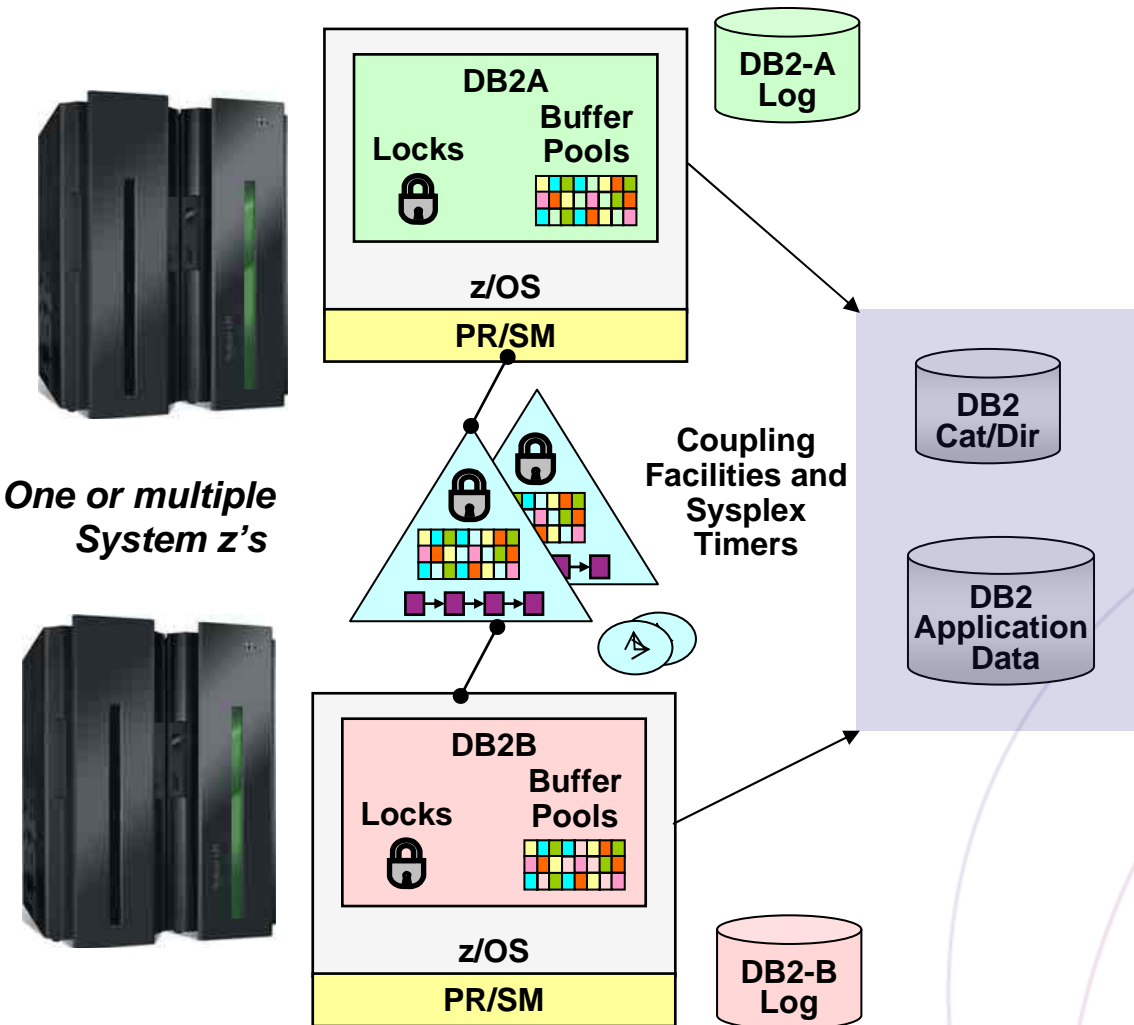
V8

- **Materialized Query Tables**
 - Improved performance for query
- **Multi-row insert and fetch**
 - Improved speed of warehouse ETL and query
- **Online Schema evolution**
 - Improves availability and efficiency
- **Longer Table and Column names**
 - Increased compatibility with ETL and BI tools
- **Increased SQL vocabulary**
 - Increased compatibility with ETL and BI tools
- **64 bit support**
 - Expanded size capability

V9

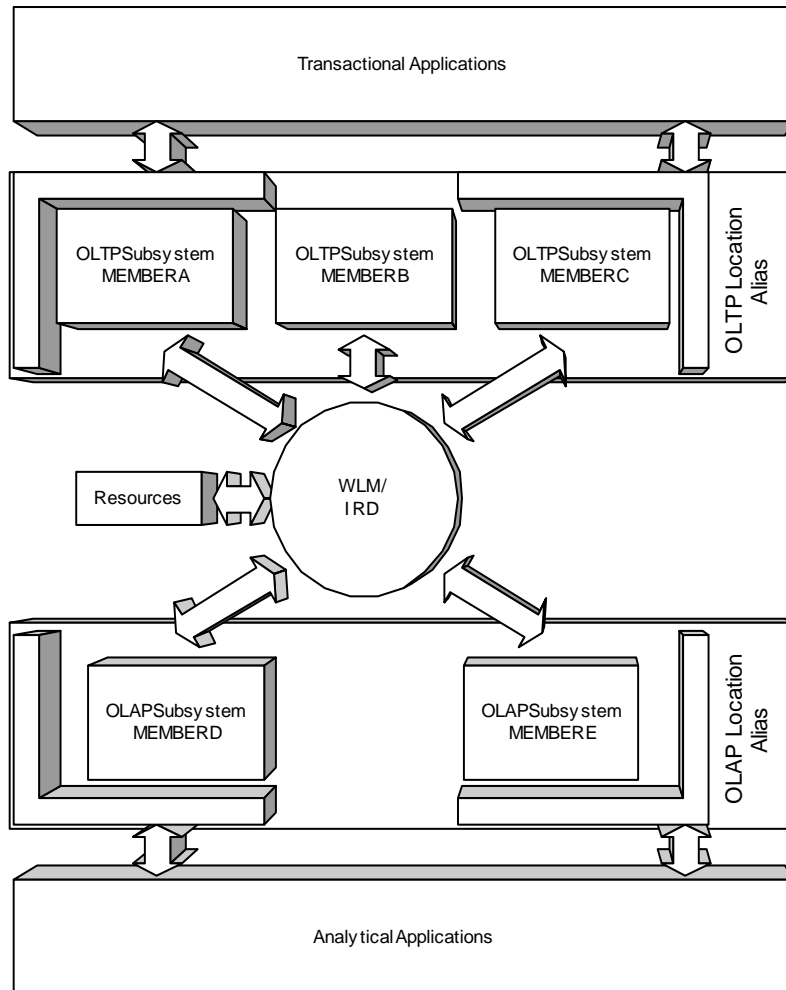
- **Index over Expression**
 - Improved performance on ETL and BI tools
- **Index Compression**
 - Improved use of space (up to 50%) and reduced CPU
- **Truncate Table**
 - Tooling portability and quickly reuse space
- **Not Logged Table Space**
 - More efficient for temporary tables such as Staging tables
- **Universal Table Space – partition by growth**
 - Easier to manage growth
 - Best of segmented tables and partitioning
- **Utility Improvements**
 - More online utility operations, reduced CPU

Building on the foundation with data sharing



- A Coupling Facility (CF) is a specialty processor which coordinates data access
 - Provides high-speed caching and lock processing
- DB2 Data Sharing Group
 - Up to 32 DB2 subsystems
 - All group members share the same DB2 catalog and directory
 - Multiple data sharing groups can be defined within the same Parallel Sysplex
- Highly Available
 - Rolling Operating System and DB2 maintenance and upgrades
 - Add members without repartitioning
 - Maintain application transparency
 - Little to no performance impact

DW Solution Architecture using DB2 for z/OS



- Centralized and Consolidated
- Transactional and Warehouse data in one system
- All members see all data, but each member is optimized for a particular workload
- Location aliases for transaction routing
- Shared Resources managed by Workload Manager (WLM) and the Intelligent Resource Director (IRD)
- Single subsystem option for non-data sharing environments


Workload Management

- Mixed workloads need to be managed
- Traditional workload management approach is problematic
 - Screen queries before they start execution
 - Time consuming for DBAs
 - Not always possible
 - All business queries are important and deserve service
 - Large queries will get into the system
 - Running these queries can hurt system performance
 - Canceling the queries is bad for public relations and wastes resources


Workload Management

Think about this:

Short query submitter behavior:



Large query submitter behavior:




Who's impacted more **real time** ??



The ideal workload manager policy for data warehousing:

Consistent favoring of shorter running work.....




keep em short

through WLM period aging

** no need to pre-identify shorts either

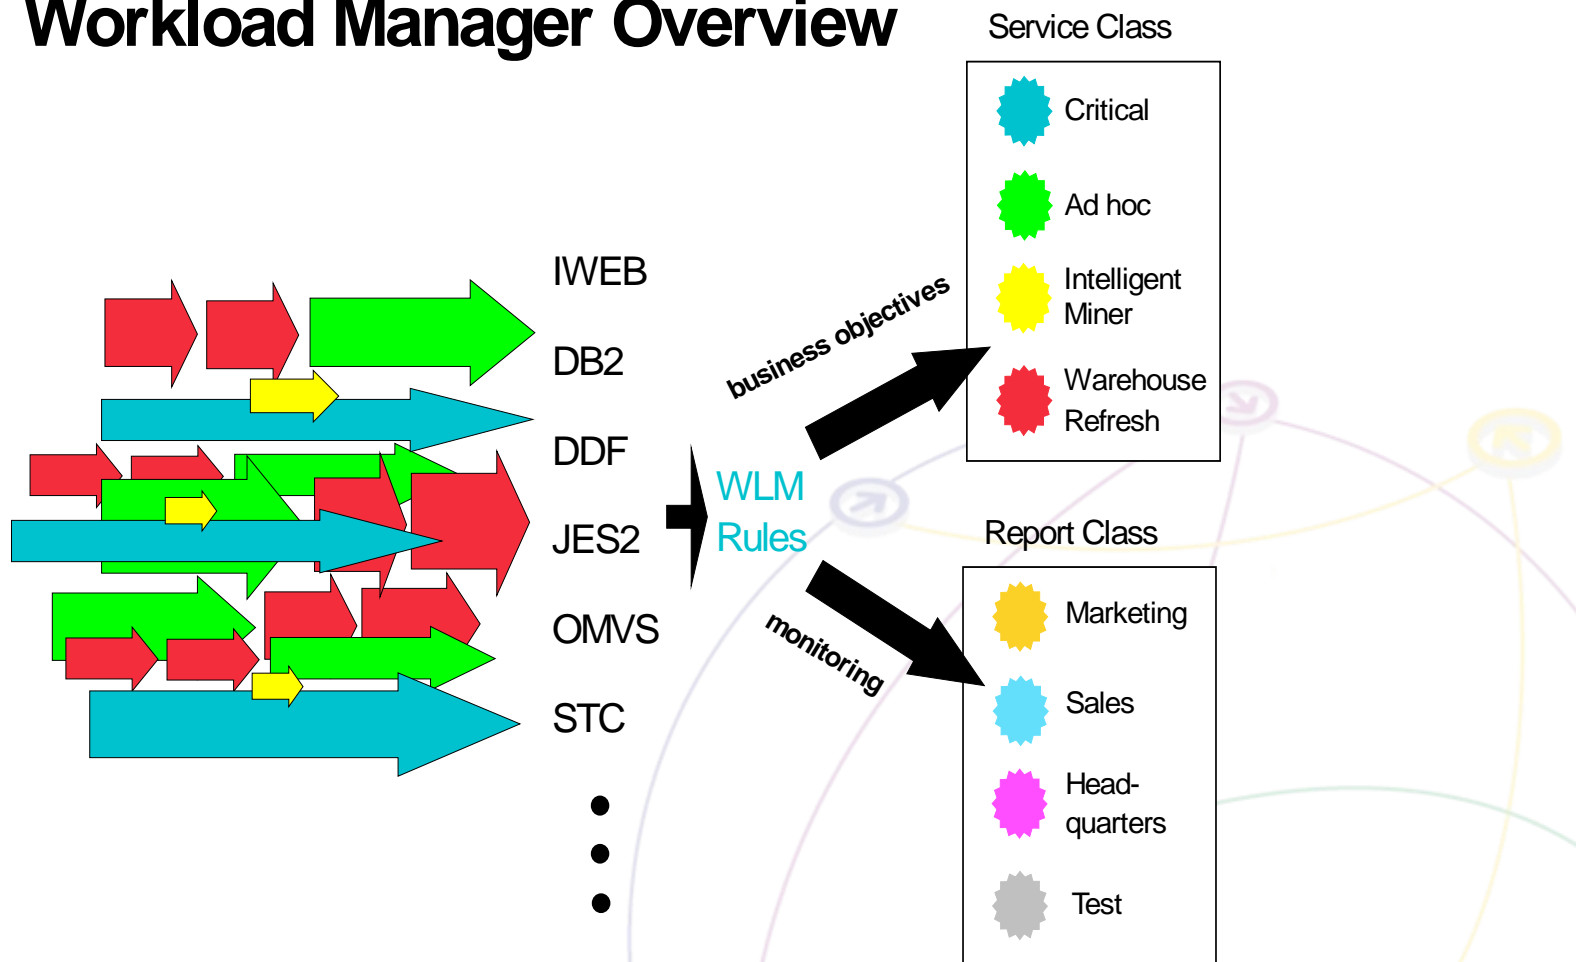
with select favoring of critical business users

through WLM explicit prioritization of critical users

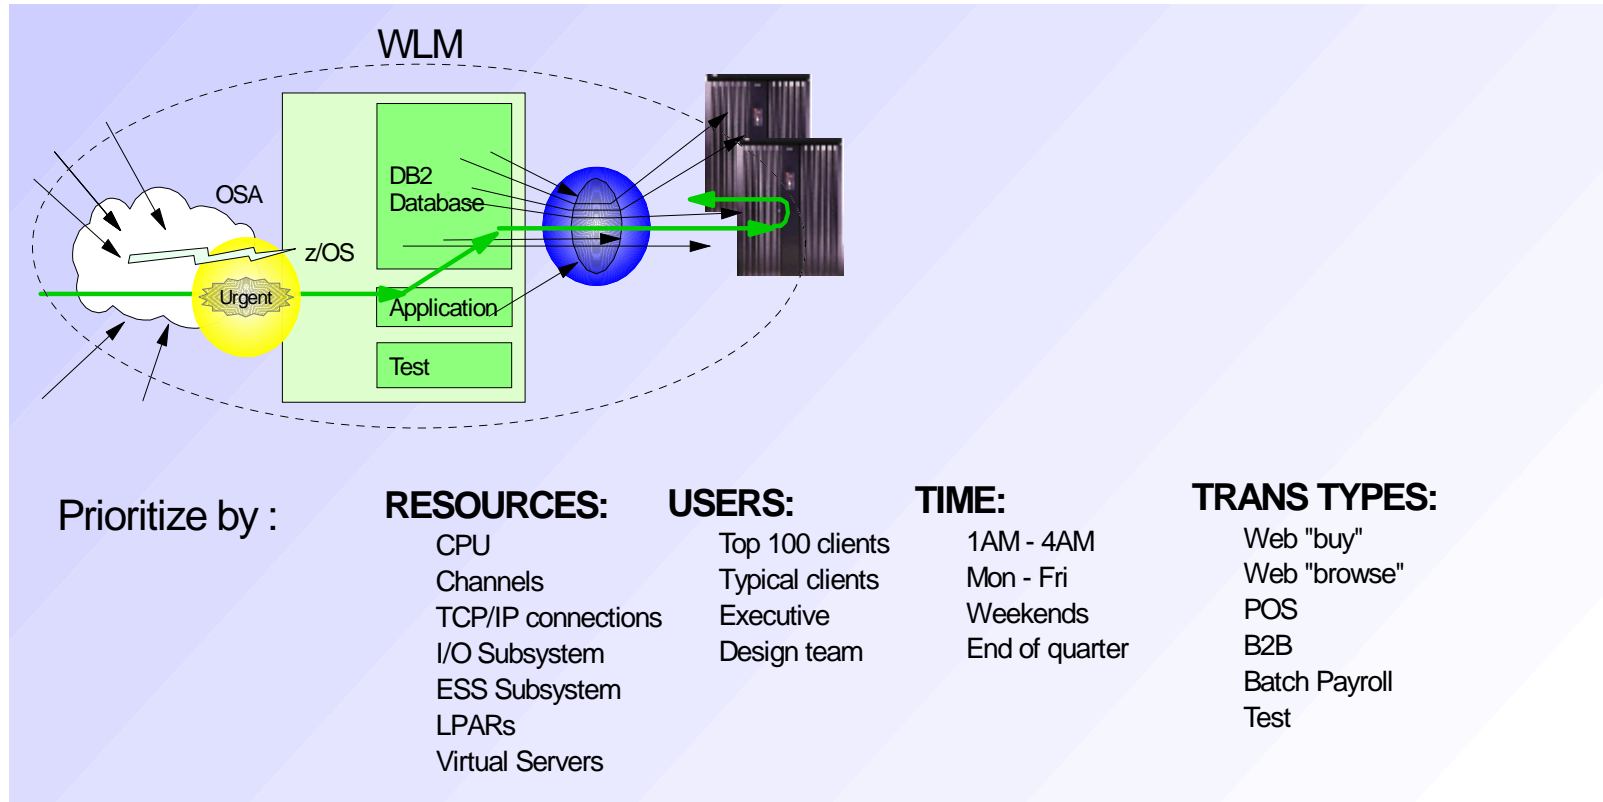


Workload Management

Workload Manager Overview

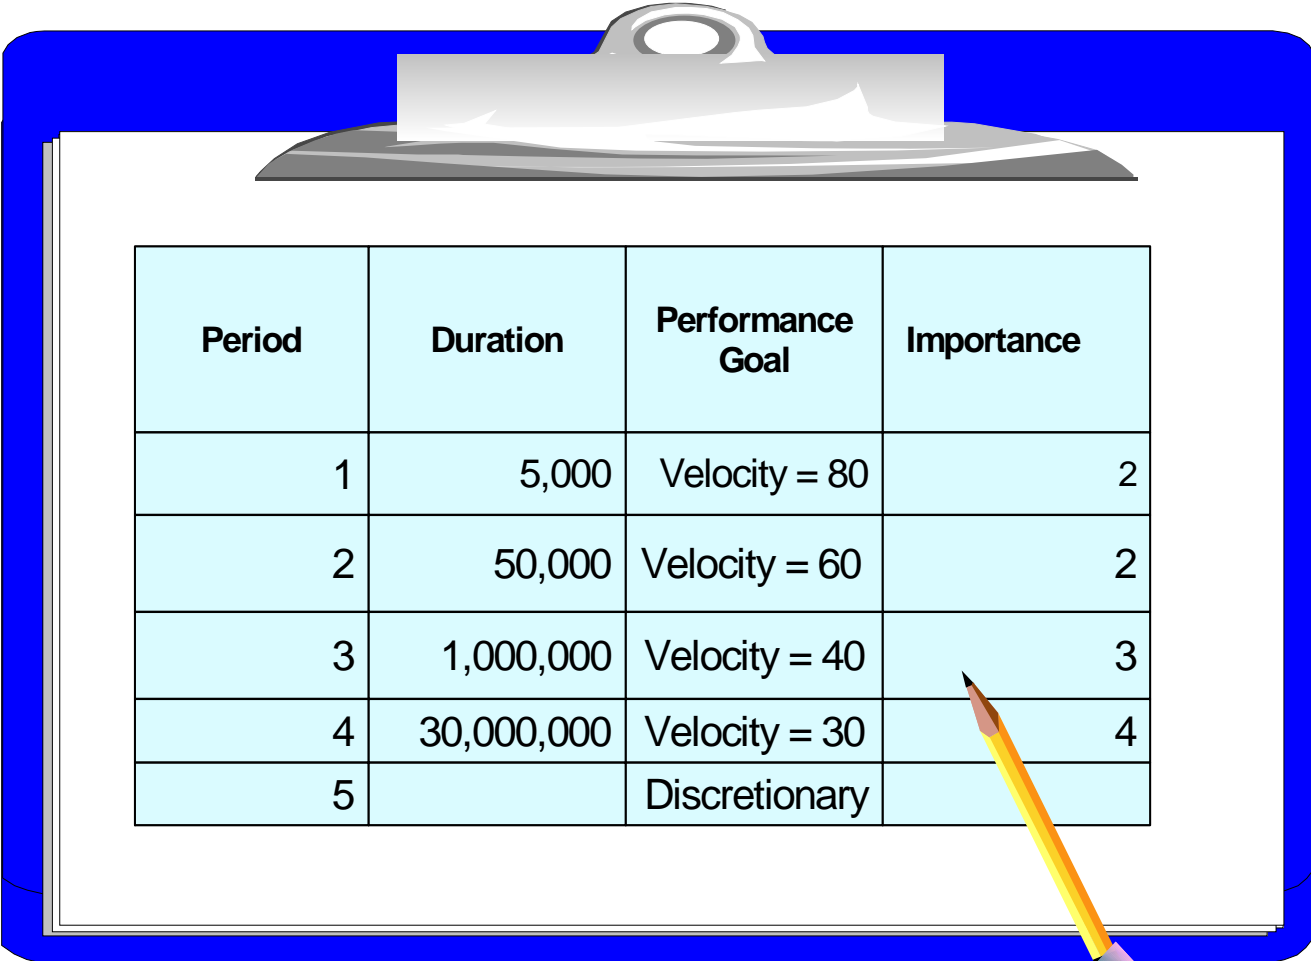


Leveraging Workload Manager to Prioritize



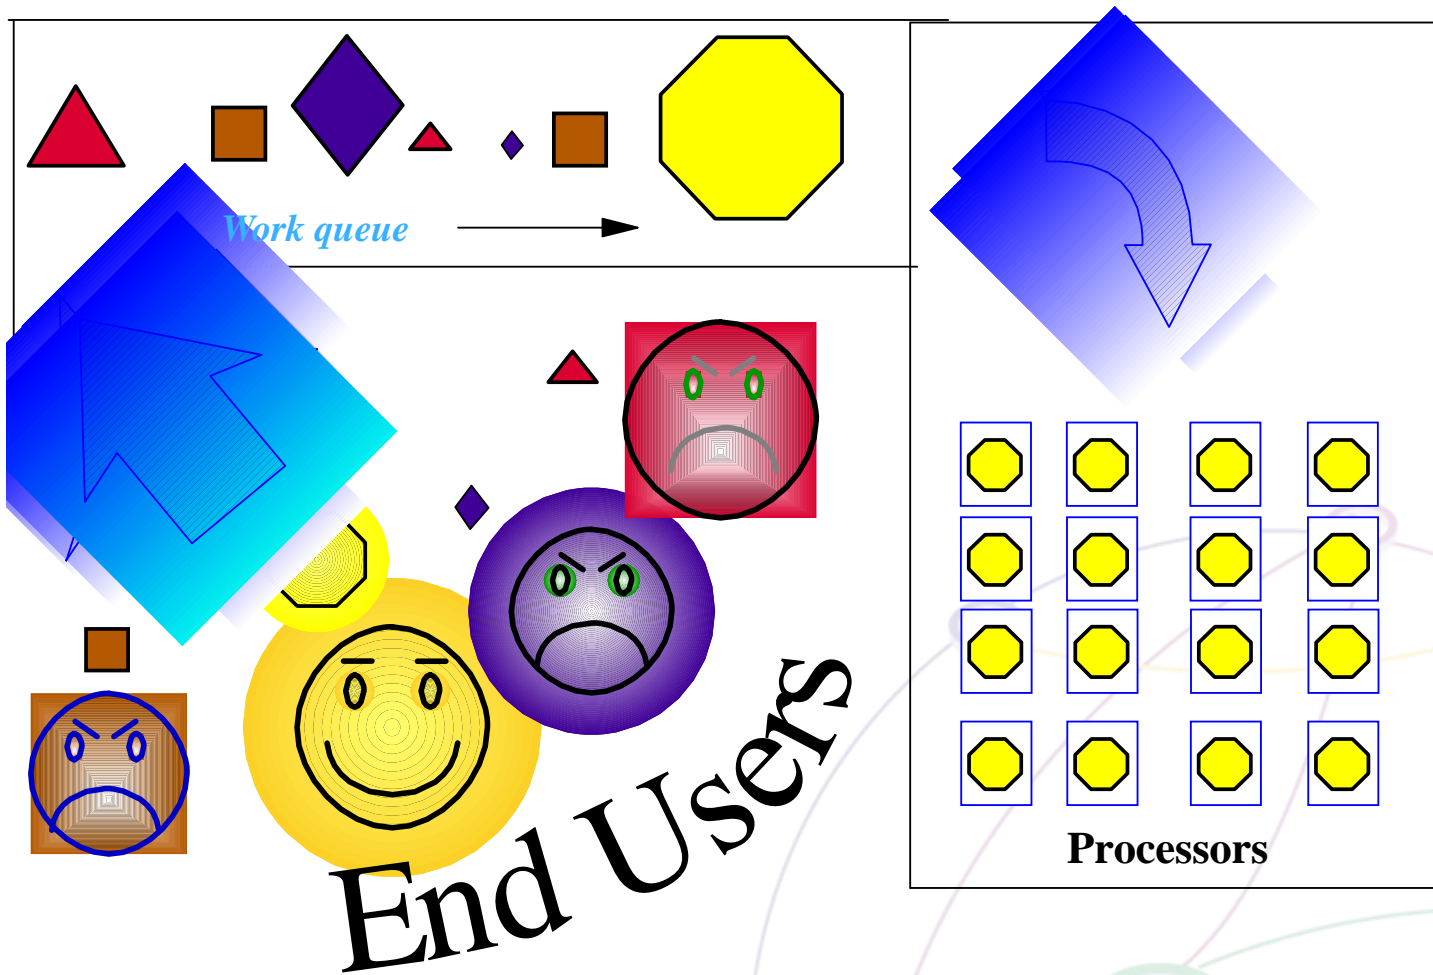
System z average system utilization often exceeds 80%, and System z servers are designed to handle sustained peak workload utilization of 100% without service level degradation to high priority workloads.

Service Classification



Period	Duration	Performance Goal	Importance
1	5,000	Velocity = 80	2
2	50,000	Velocity = 60	2
3	1,000,000	Velocity = 40	3
4	30,000,000	Velocity = 30	4
5		Discretionary	

Query Monopolization

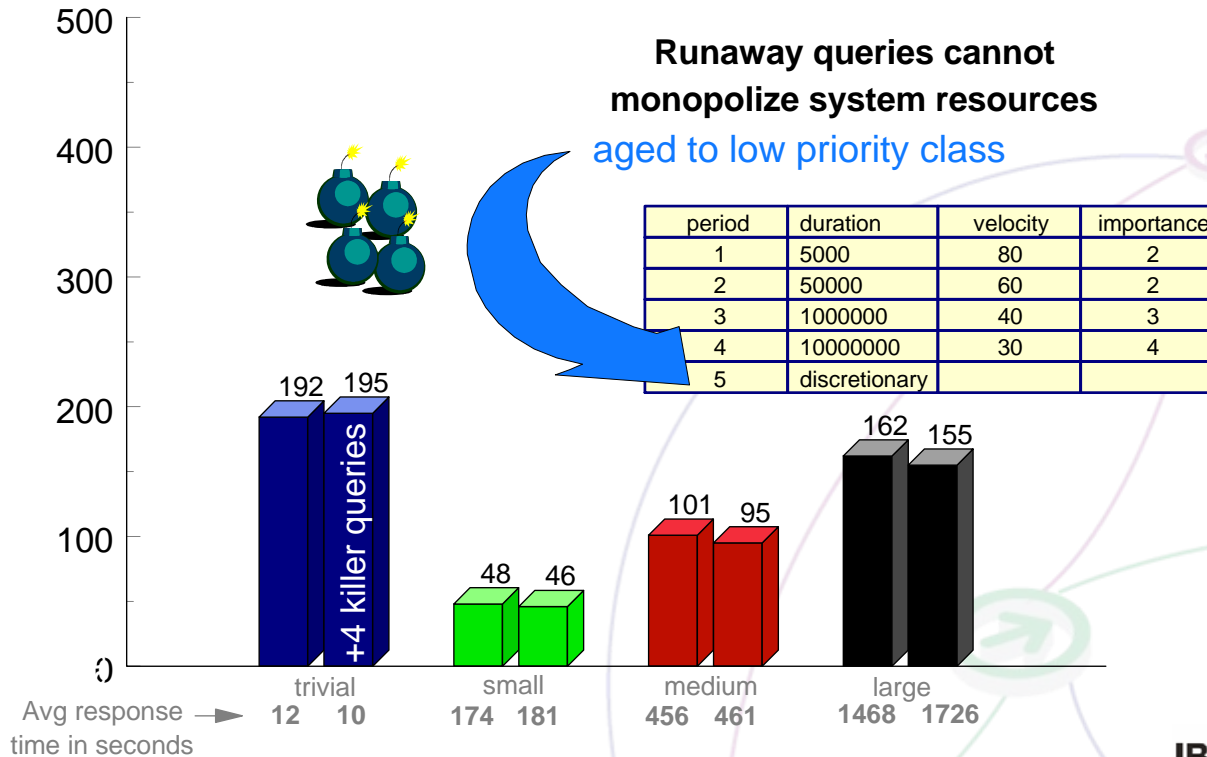


Query Monopolization Solution

50 Concurrent Users

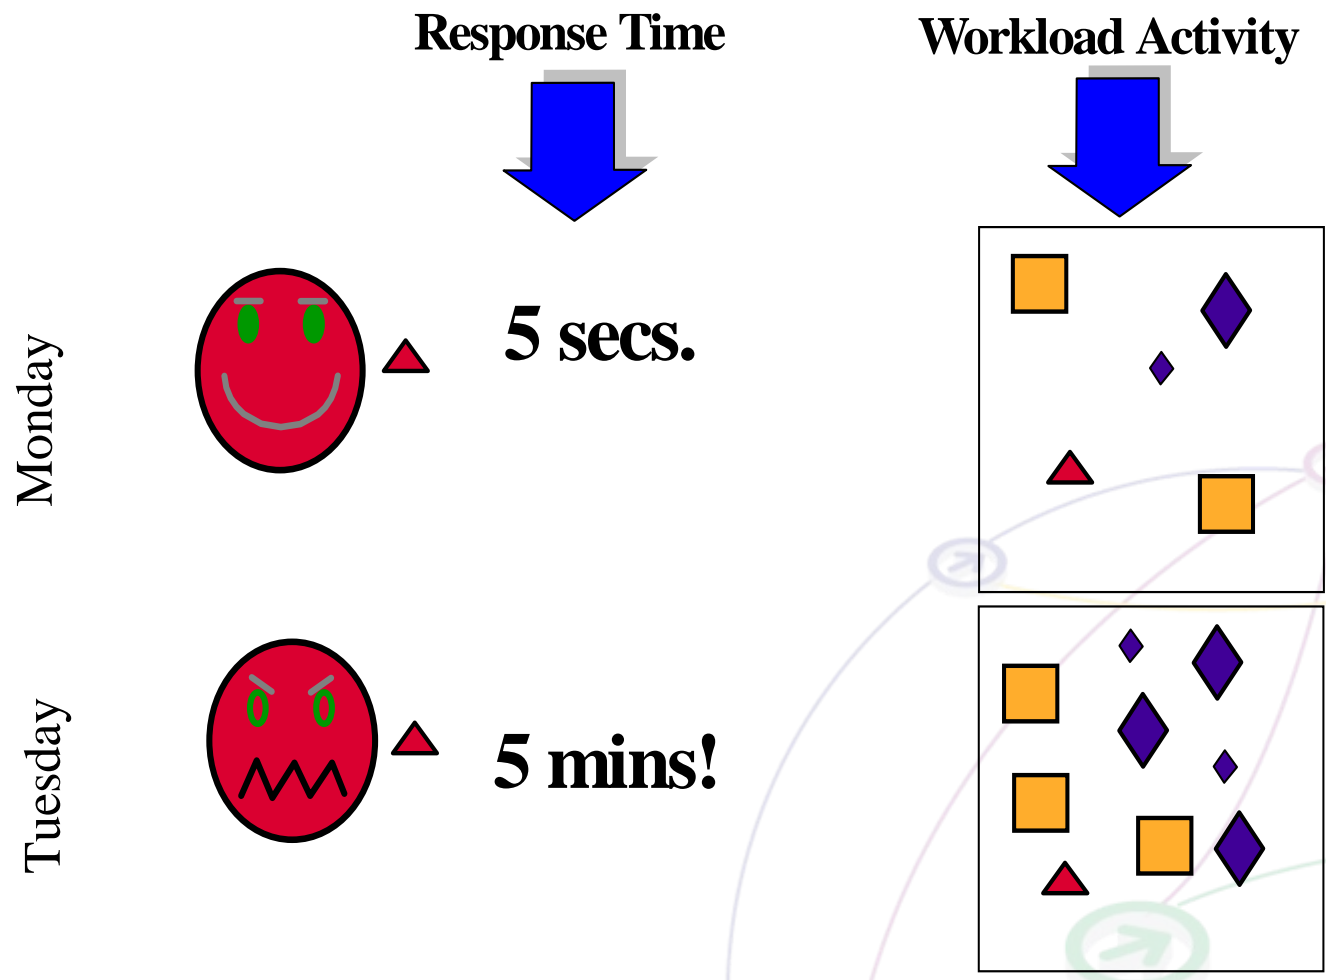


Killer Queries



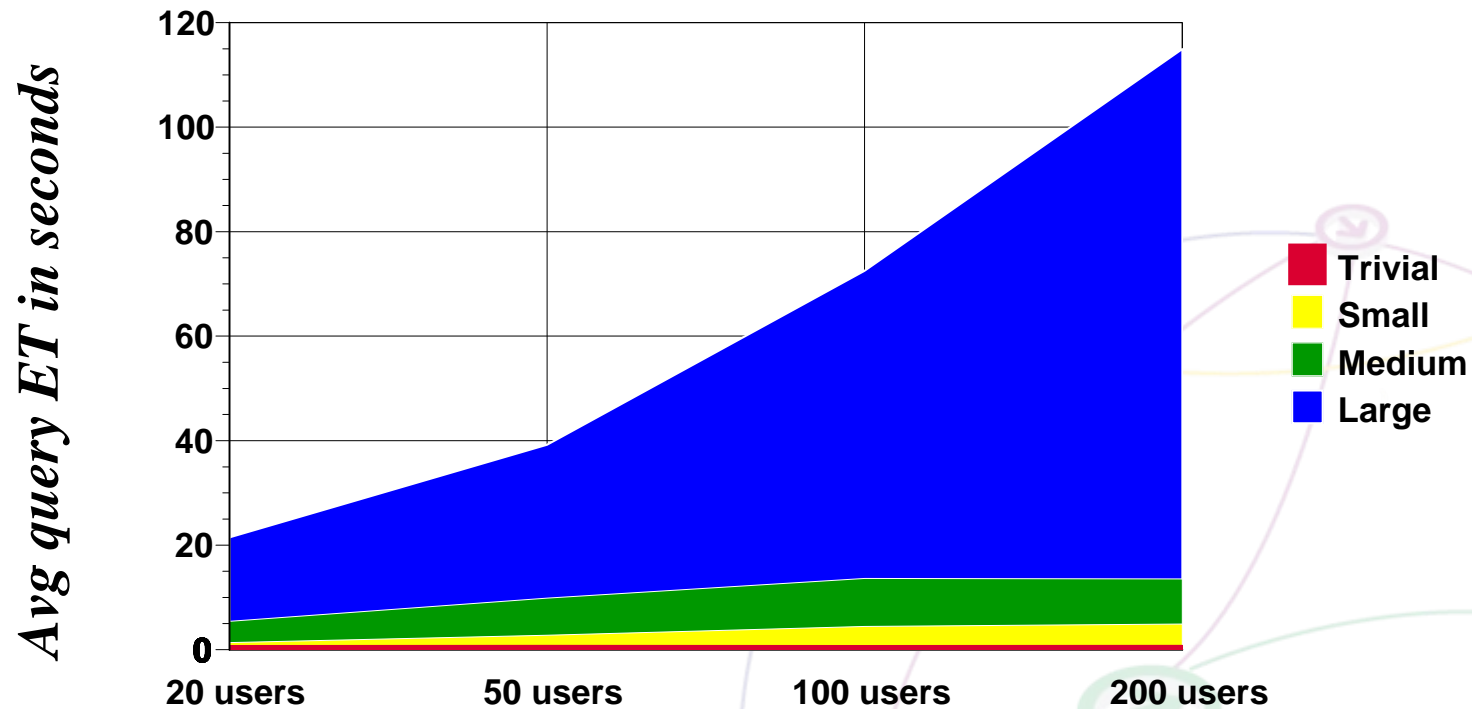


Inconsistent Response Times for Short Running Queries



Consistent Response Time

Consistent Response Time for Short-running work



High Priority Queries / Dynamic Warehouses

Service Classification

Service Class = QUERY	Period	Duration	Performance Goal	Importance
	1	5,000	Velocity = 80	1
	2	20,000	Velocity = 70	2
	3	1,000,000	Velocity = 50	3
	4		Velocity = 30	


Service Class CRITICAL	Period	Duration	Performance Goal	Importance
	1		Velocity = 90%	1

DB2 for z/OS Parallelism

V3

I/O Parallelism
Single Execution Unit
Multiple I/O streams


TCB
(originating task)



V4

CP Parallelism
Multiple Execution Units
Each has a single I/O stream


SRBs
(parallel tasks)



V5

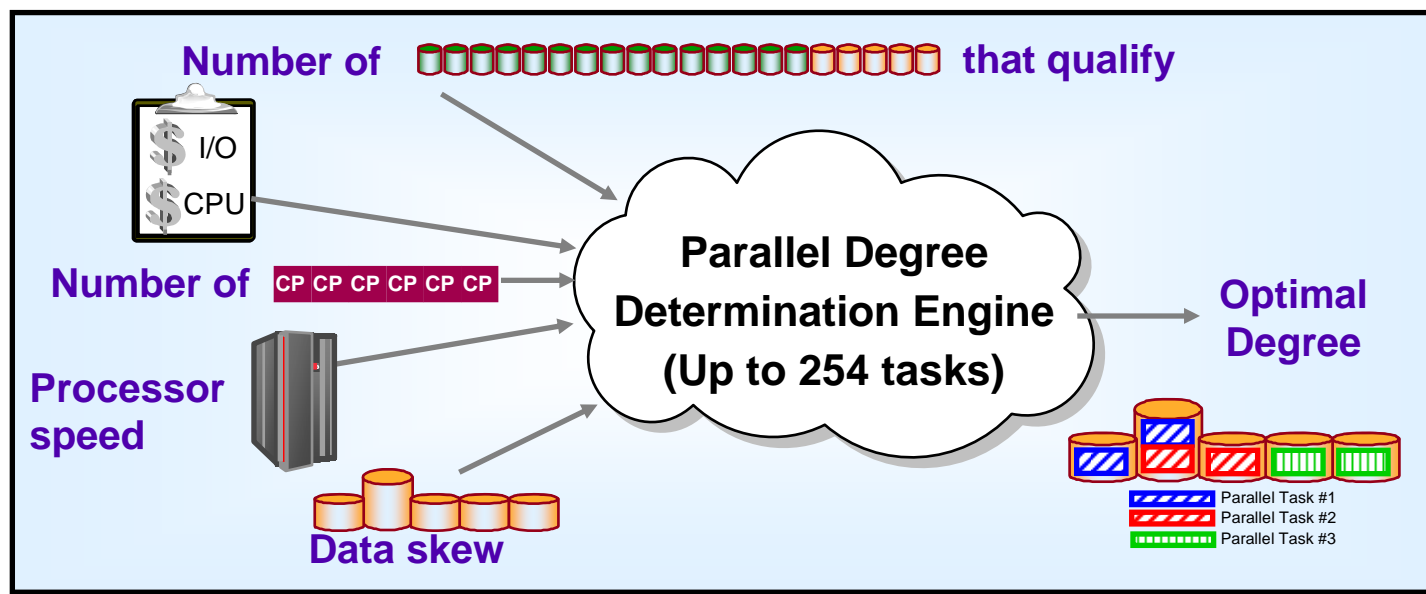
Sysplex Query Parallelism
Sysplex Query Parallelism

Parallel tasks spread across Parallel Sysplex



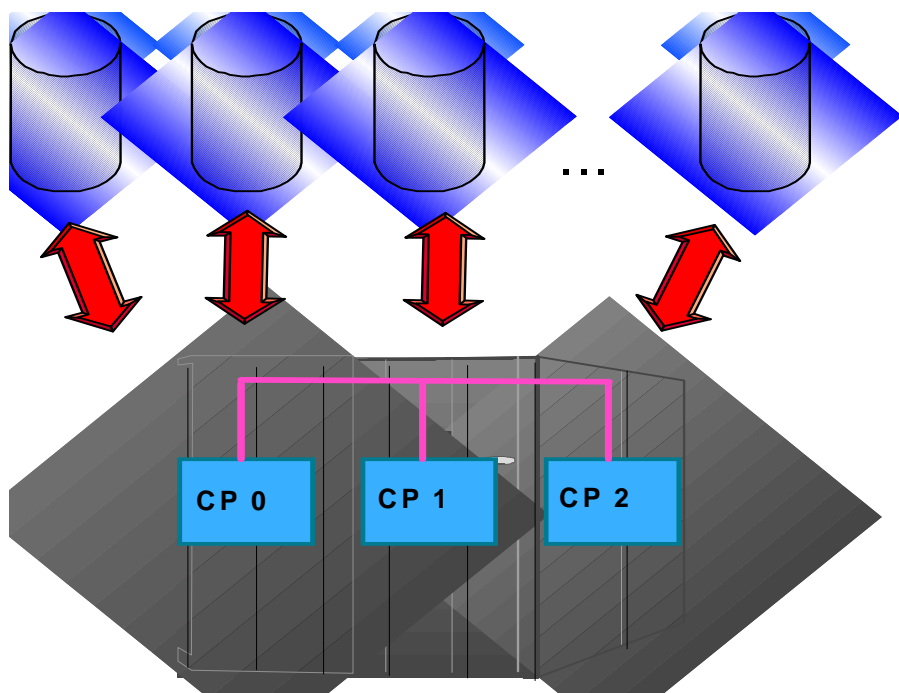
Parallel Degree Determination

- DB2 for z/OS has the flexibility to choose the degree of parallelism



- Degree determination is done by the DBMS -- not the DBA
- Using statistics and costs of the query to provide the optimal access path with low overhead taking data skew into consideration

Trust with Limits – Set the Max. Degree of Parallelism



Set the maximum degree between the # of CPs and the # of partitions

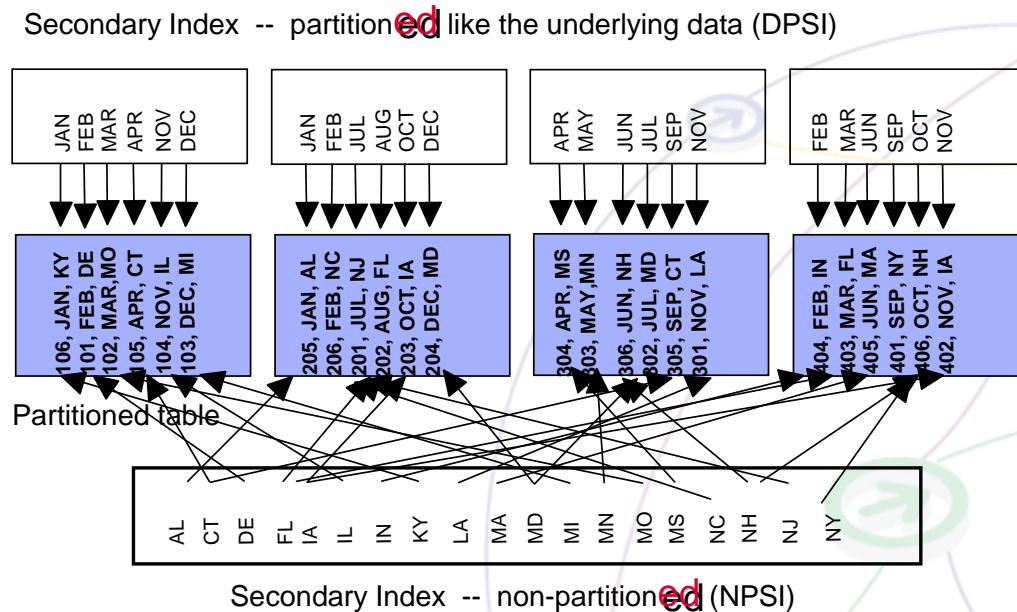
CPU intensive queries - closer to the # of CPs

I/O intensive queries - closer to the # of partitions


Data skew can reduce # of degrees

DB2 for z/OS Partitioning (Range)

- ➔ Partitioning in DB2 for z/OS V8 and beyond is defined at the table level
- ➔ Maximum of 4096 partitions
- ➔ Effectively cluster by two dimensions
- ➔ Partition by Growth – DB2 9 feature that relates (in a way) to hash

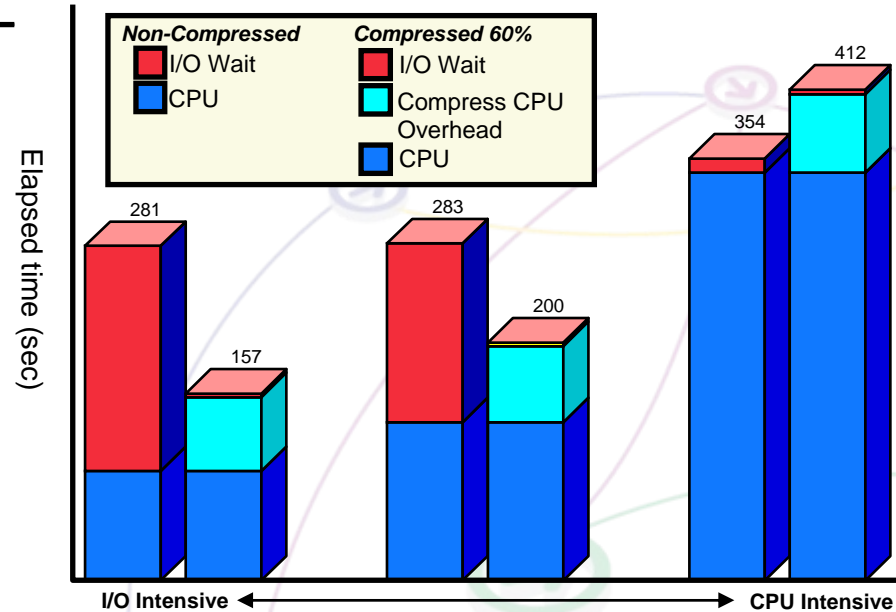


Hardware-Assisted Data Compression




← Compression Ratios Achieved

→ Effects of Compression on Elapsed Time



Index Compression

- Solution provides page-level compression
 - Data is compressed to 4K pages on disk
 - 32K/16K/8K pages results in up to 8x/4x/2x disk savings
 - Use DSN1COMP to estimate
 - No compression dictionaries
 - Compression on the fly
 - No LOAD or REORG required



Index Compression Versus Data Compression

→ There are differences between index and data compression


	Data	Index
Level	Row	Page (1)
Comp in DASD	Yes	Yes
Comp in Buffer Pool	Yes	No
Comp in Log	Yes	No
Comp Dictionary	Yes	No (2)
'Typical' Comp Ratio	10 to 90%	25 to 75% (3)
CPU Overhead (4)	In Accounting	In Accounting and/or DBM1 SRB

1. No compression or decompression at each insert or fetch; instead, it is done at I/O time
2. LOAD or REORG not required for compression; compression on the fly
3. Based on very limited survey and usage
4. CPU time impact under study – sensitive to index BP Hit Ratio, larger index BP recommended, higher impact on relatively unique indexes with long keys




Noteworthy zPARMs in a Data Warehouse Environment

Recommendation for DWH	Comments
CDSSRDEF=ANY,	ANY Allow parallelism for DW. any: parallelism, 1: no parallelism
CONTSTOR=NO,	NO For best performance, specify NO for this parameter. To resolve storage constraints in DBM1 address space, specify YES. See also: MINSTOR
DSVCI=YES,	YES The DB2-managed data set has a VSAM control interval that corresponds to the buffer pool that is used for the table space.
MGEXTSZ=YES,	YES Secondary extent allocations for DB2-managed data sets are to be sized according to a sliding scale
MINSTOR=NO,	NO Recommendation: For best performance, specify NO for this parameter. To resolve storage constraints in DBM1 address space, specify YES. See also: CONTSTOR
OPTCCOS4=ON,	enables fix PK26760 (inefficient access plan)
OPTIXIO=ON	OPTIXIO=ON: Provides stable I/O costing with significantly less sensitivity to buffer pool sizes. (This is the new default and recommended setting).
OPTIORC=ON	OPTIORC=ON – explanation???
PARAMDEG=X,	#Processors <= X <= 2*#Processors If concurrency level is low, the ratio can be higher.
SRTPOOL=8000,	8000



Objective Number 4

Where can additional information be found and who can I contact to get started?





Recent papers

Data Warehousing and Business Intelligence for IBM System z
June 2008

Information Management software

Data Warehousing and Business
Intelligence for IBM System z

ftp://software.ibm.com/software/data/businessintelligence/systemz/DW_BI_IBM_SysZ.pdf

Delivering information you can trust
December 2008

Information Management software

Operational business intelligence:
The renaissance of the mainframe

www.ibm.com/software/data/businessintelligence/systemz/

Information Management software

Data Warehousing with the DB2 Family for System z
Customers
Leveraging Platform Strengths

ftp://software.ibm.com/software/data/businessintelligence/systemz/Data_Warehousing_with_DB2_for_System_z_Jan_2009.pdf

Some key Redbooks



- 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
 - Currently in production SG24-7674
 - www.redbooks.ibm.com/redpieces/abstracts/sg247674.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
- Workload Management for DB2 Data Warehouse, REDP-3927
 - www.redbooks.ibm.com/abstracts/redp3927.html?Open



Other recent articles in the press

IBM Systems
MAGAZINE

- **Enhanced Query Parallelism with zIIP processors**
 - February 2008
www.ibmssystemsmag.com/mainframe/enewsletterexclusive/18822p1.aspx
- **Operational BI and System z**
 - March 2008
www.ibmssystemsmag.com/mainframe/enewsletterexclusive/19475p1.aspx
- **Business Intelligence's New Look: IBM extends its BI portfolio with Cognos 8 BI for Linux on System z**
 - July / August 08
www.ibmssystemsmag.com/mainframe/julyaugust08/features/20870p1.aspx
- **Take the Reins - An Information On Demand Strategy helps deliver a competitive edge for today's businesses**
 - July / August 08
www.ibmssystemsmag.com/mainframe/julyaugust08/coverstory/20860p1.aspx

DMReview.com

- **Three part series: Myths of Doing BI on the mainframe**
 - www.dmreview.com/issues/2007_53/10002140-1.html
 - www.dmreview.com/issues/2007_54/10002171-1.html
 - January 2009 features the final part in the series.

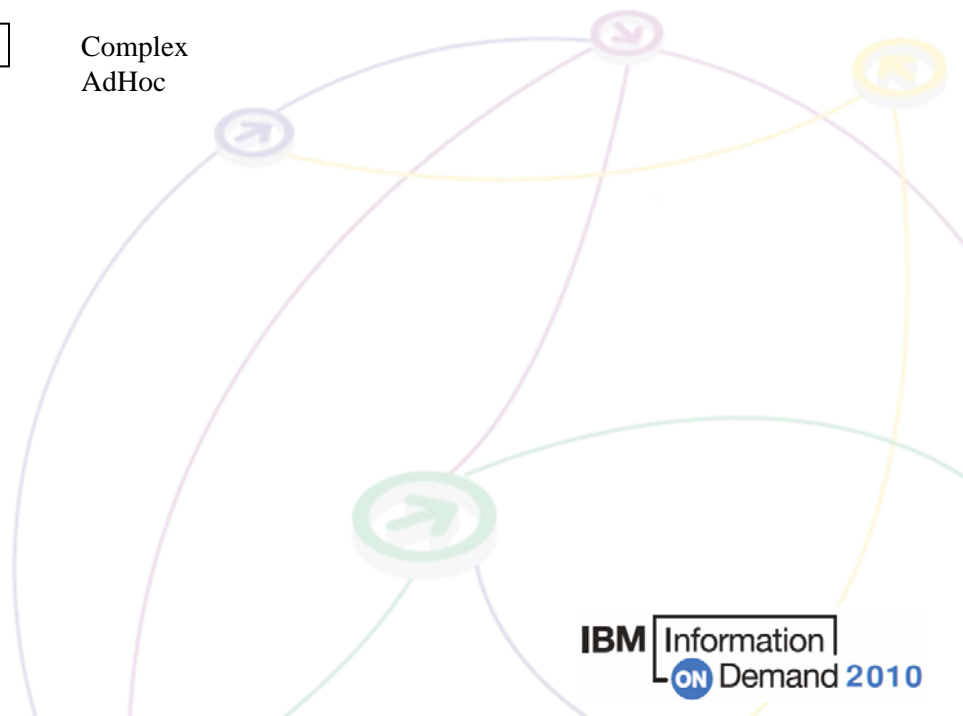
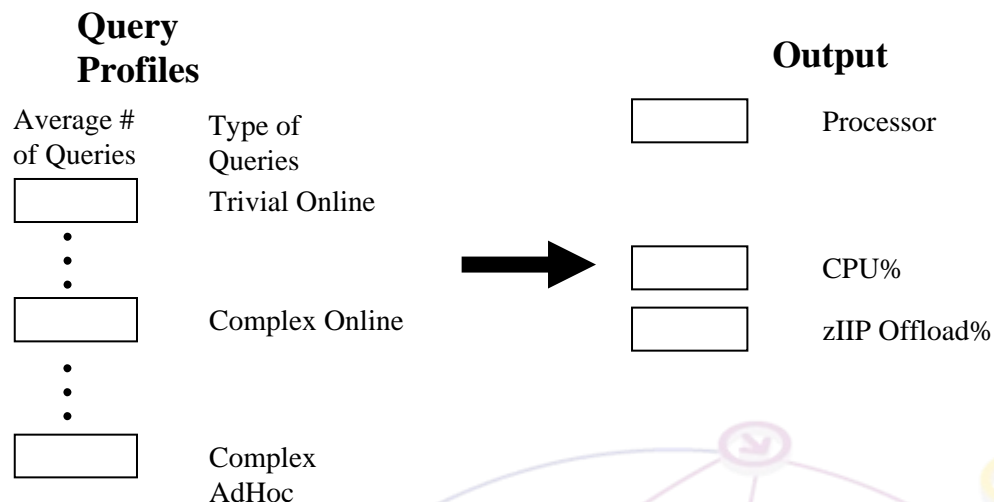
DB2 Sizing Tool

→ Model based on workloads that were run in IBM Lab environments

- Continually refined
- Classify queries into 5 categories
- Provide query category definitions
- Specify % of data "touched" in each query category
- Request size of data warehouse
- Compute required capacity including zIIP offload percentage

→ Alternate method - build a prototype and profile your own workload

→ Consider starting small and growing incrementally (benefit of System z DWH environment)



DB2 for z/OS Data Warehouse Terabyte Club

- Customers who have a Data Warehouse, Operational Data Store, Data Mart, Reporting System or any Decision Support System running on DB2 for z/OS
 - 1 TB or more of uncompressed, raw data
 - No indexes, MQTs etc. included
 - SQL / Script provided to measure
 - Run SQL and submit an application
- Silent Partner
 - Email access to zSwat team
 - T-Shirts
- Full Partner
 - All of the above and...
 - Recognition at IOD
 - zSwat team environment evaluation
 - IBM Executive visit
- www.ibm.com/software/data/businessintelligence/systemz/ter



Thank you!

