



IBM DB2 Analytics Accelerator

Boost your DB2 z/OS Queries

Patrick Hempeler

Client Technical Professional

Pan Europe DWH/BI on System z

Patrick.Hempeler@de.ibm.com



Disclaimer

© Copyright IBM Corporation 2012. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. 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. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. 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 ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, DB2, InfoSphere, Cognos, and InfoSphere Warehouse on System z are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Would You Still Use Google
If It Took 3 Days And 7 People
To Get A Search Result?



days for a single query

constant tuning

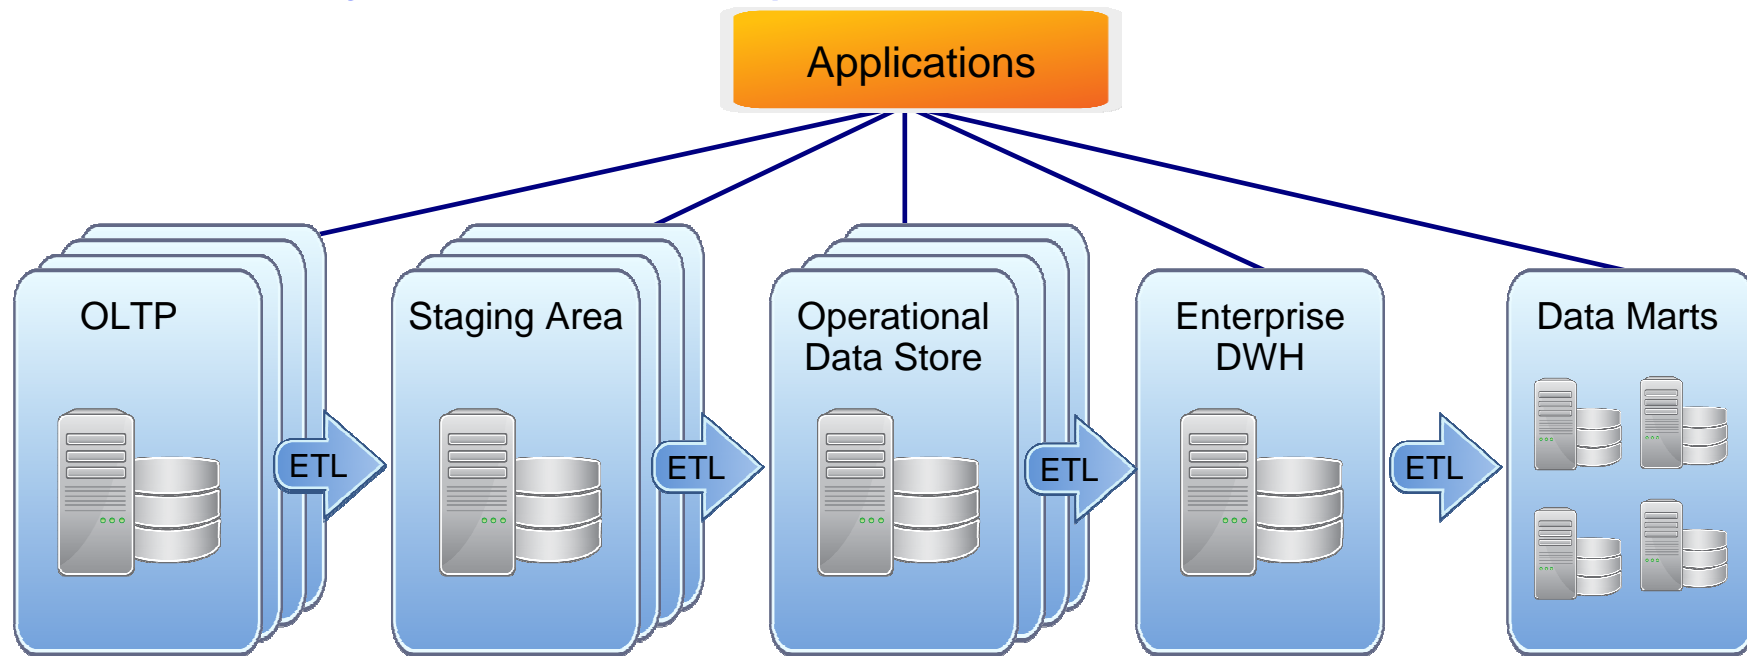
“ **Nearly 70%** of data warehouses experience performance-constrained issues of various types. ”

- Gartner 2010 Magic Quadrant

specialized resources required

months to deploy

Traditional Systems Landscape



Historical reasons:

- Different access patterns
 - impact on performance
- EDW as the data integration hub
 - again, impact on performance
- Different life-cycle characteristics
 - and again, impact on performance
- Different Service Level Agreements (SLA)
 - Lack of broadly available workload management capabilities
 - Choice of lower cost-of-acquisition offerings

Negative ramifications:

- Complexity
 - both in systems management and in applications
- Difficulties in supporting real time analytics
- Inability to match ever more demanding SLA requirements
- High total cost of ownership

Traditional data warehouses are just too complex

They are based on databases optimized for transaction processing—**NOT** to meet the demands of advanced analytics on big data.

- Too complex infrastructure
- Too complicated to deploy
- Too much tuning required
- Too inefficient at analytics
- Too many people needed to maintain
- Too costly to operate

Too long to get answers



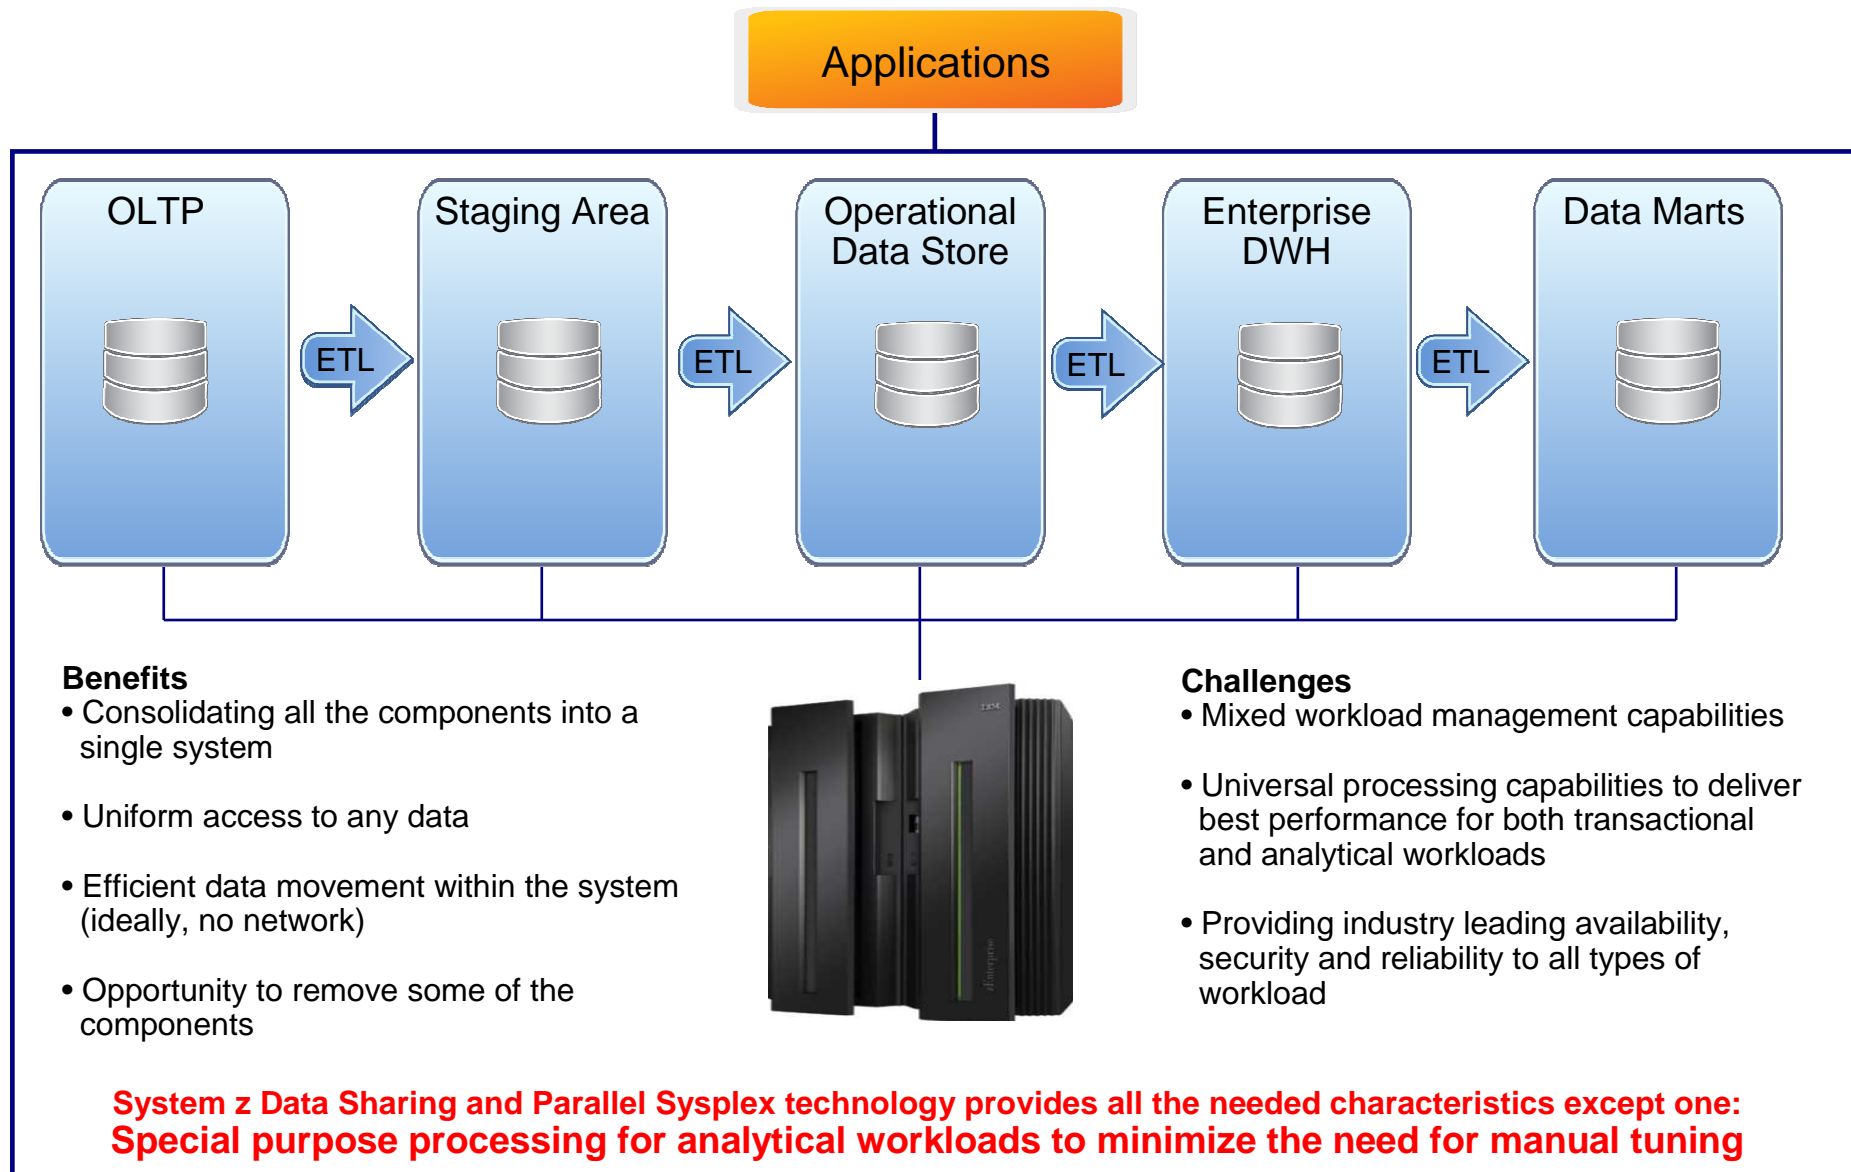
The right data warehouse
is now mission critical.

Data continues to
expand exponentially.

Analytics are becoming more complex as
business demands faster answers.



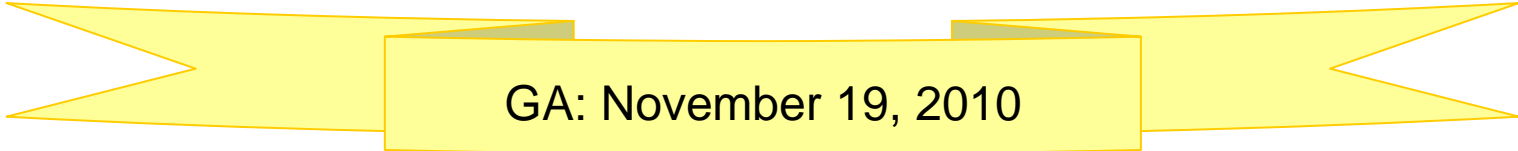
Visionary Systems Landscape



IBM United States Software Announcement
210-266, dated July 22, 2010



**IBM Smart Analytics Optimizer for DB2 for z/OS
enables a new class of business intelligence and data
warehousing workloads for the IBM System z platform**



GA: November 19, 2010

ISAOpt V1 Needed Following Enhancements

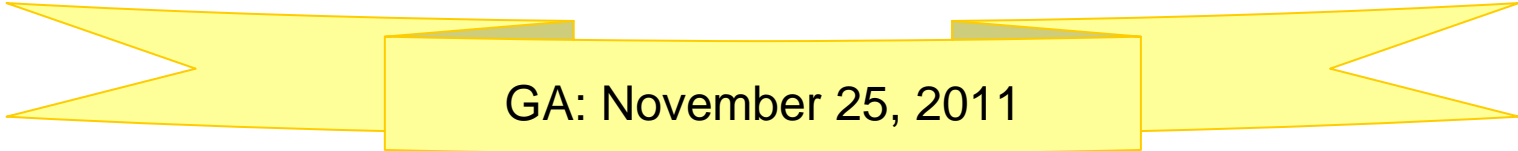
1. Increase applicability by relaxing current off-load restrictions
2. Increase applicability by supporting larger amount of data
3. Support concurrent query execution
4. Improve data currency
5. Support disaster recovery
6. DB2 10 support



IBM United States Software Announcement
211-454, dated October 12, 2011



**IBM DB2 Analytics Accelerator for z/OS , V2.1
enables a new class of business intelligence and data
warehousing workloads for the IBM System z platform**

A yellow ribbon graphic with a central rectangular box containing the text "GA: November 25, 2011".

GA: November 25, 2011

IBM DB2 Analytics Accelerator

Capitalizing on the best of both worlds – System z and Netezza

What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.

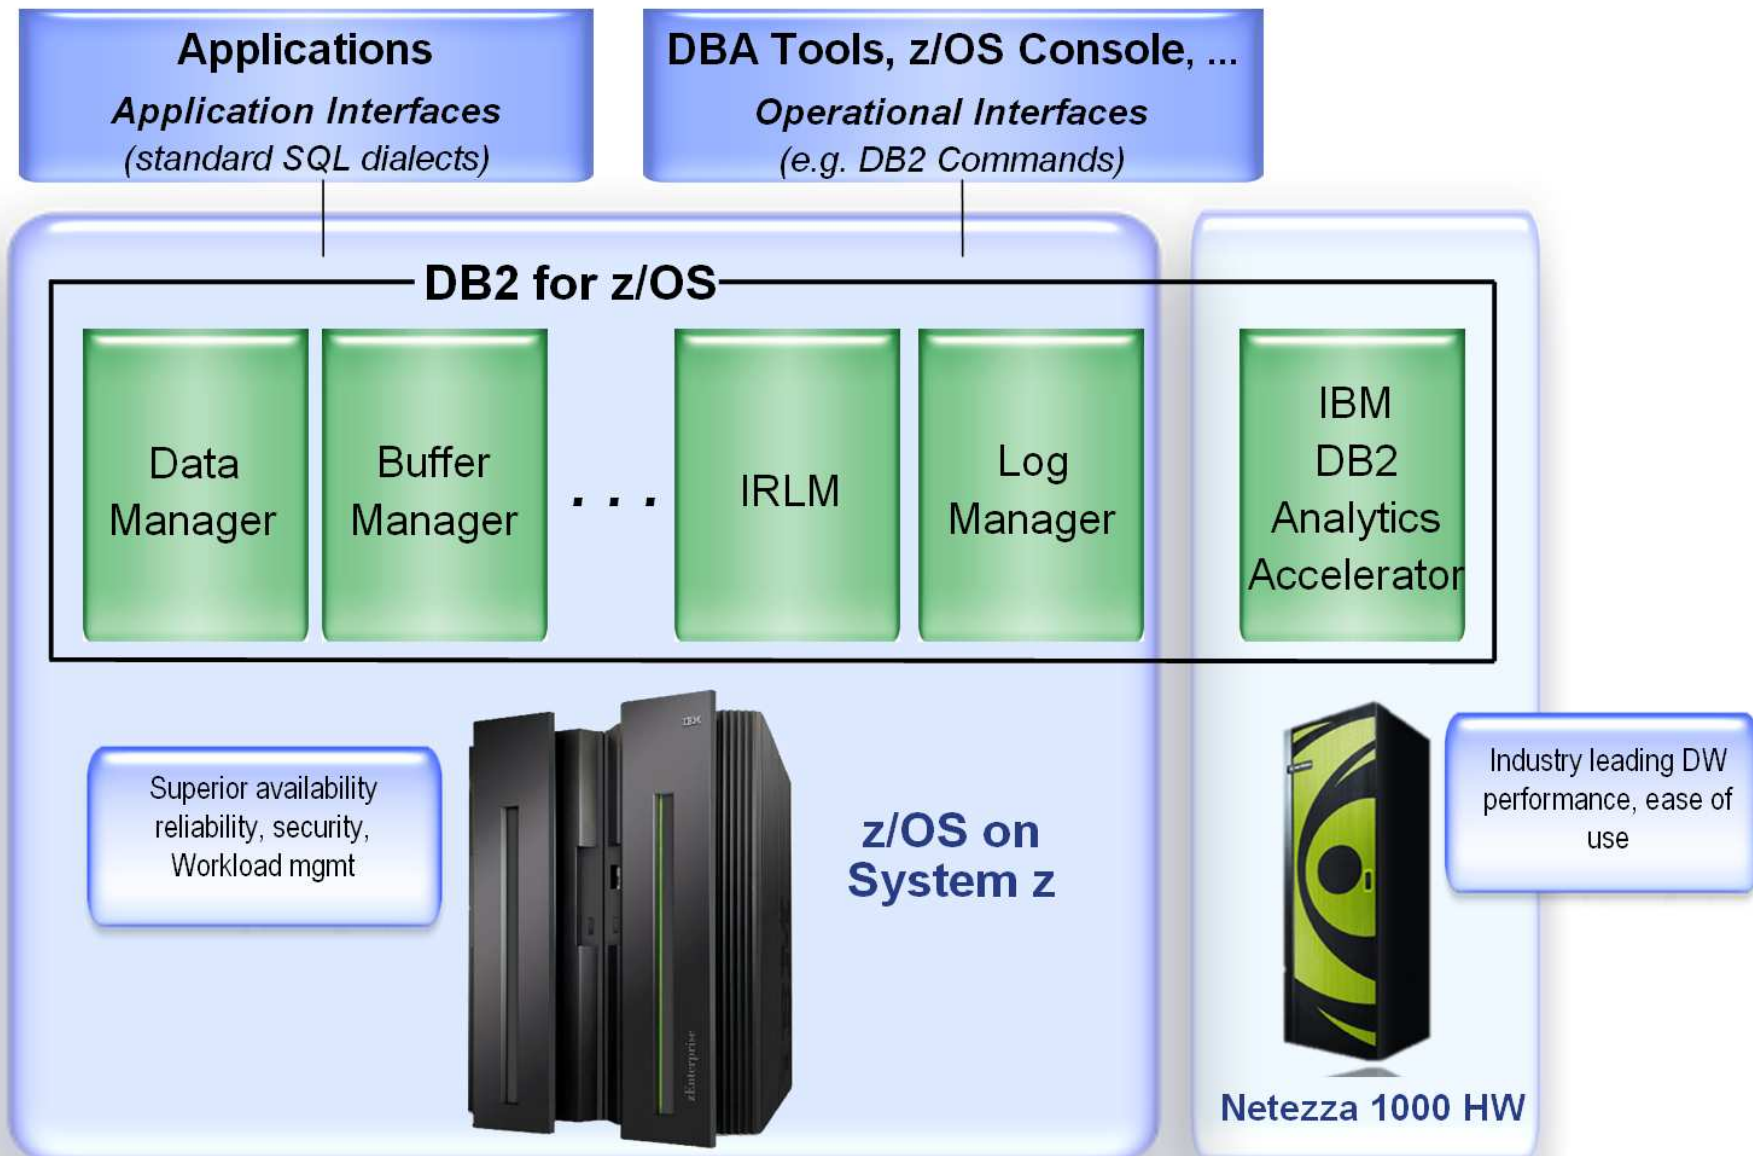


How is it different

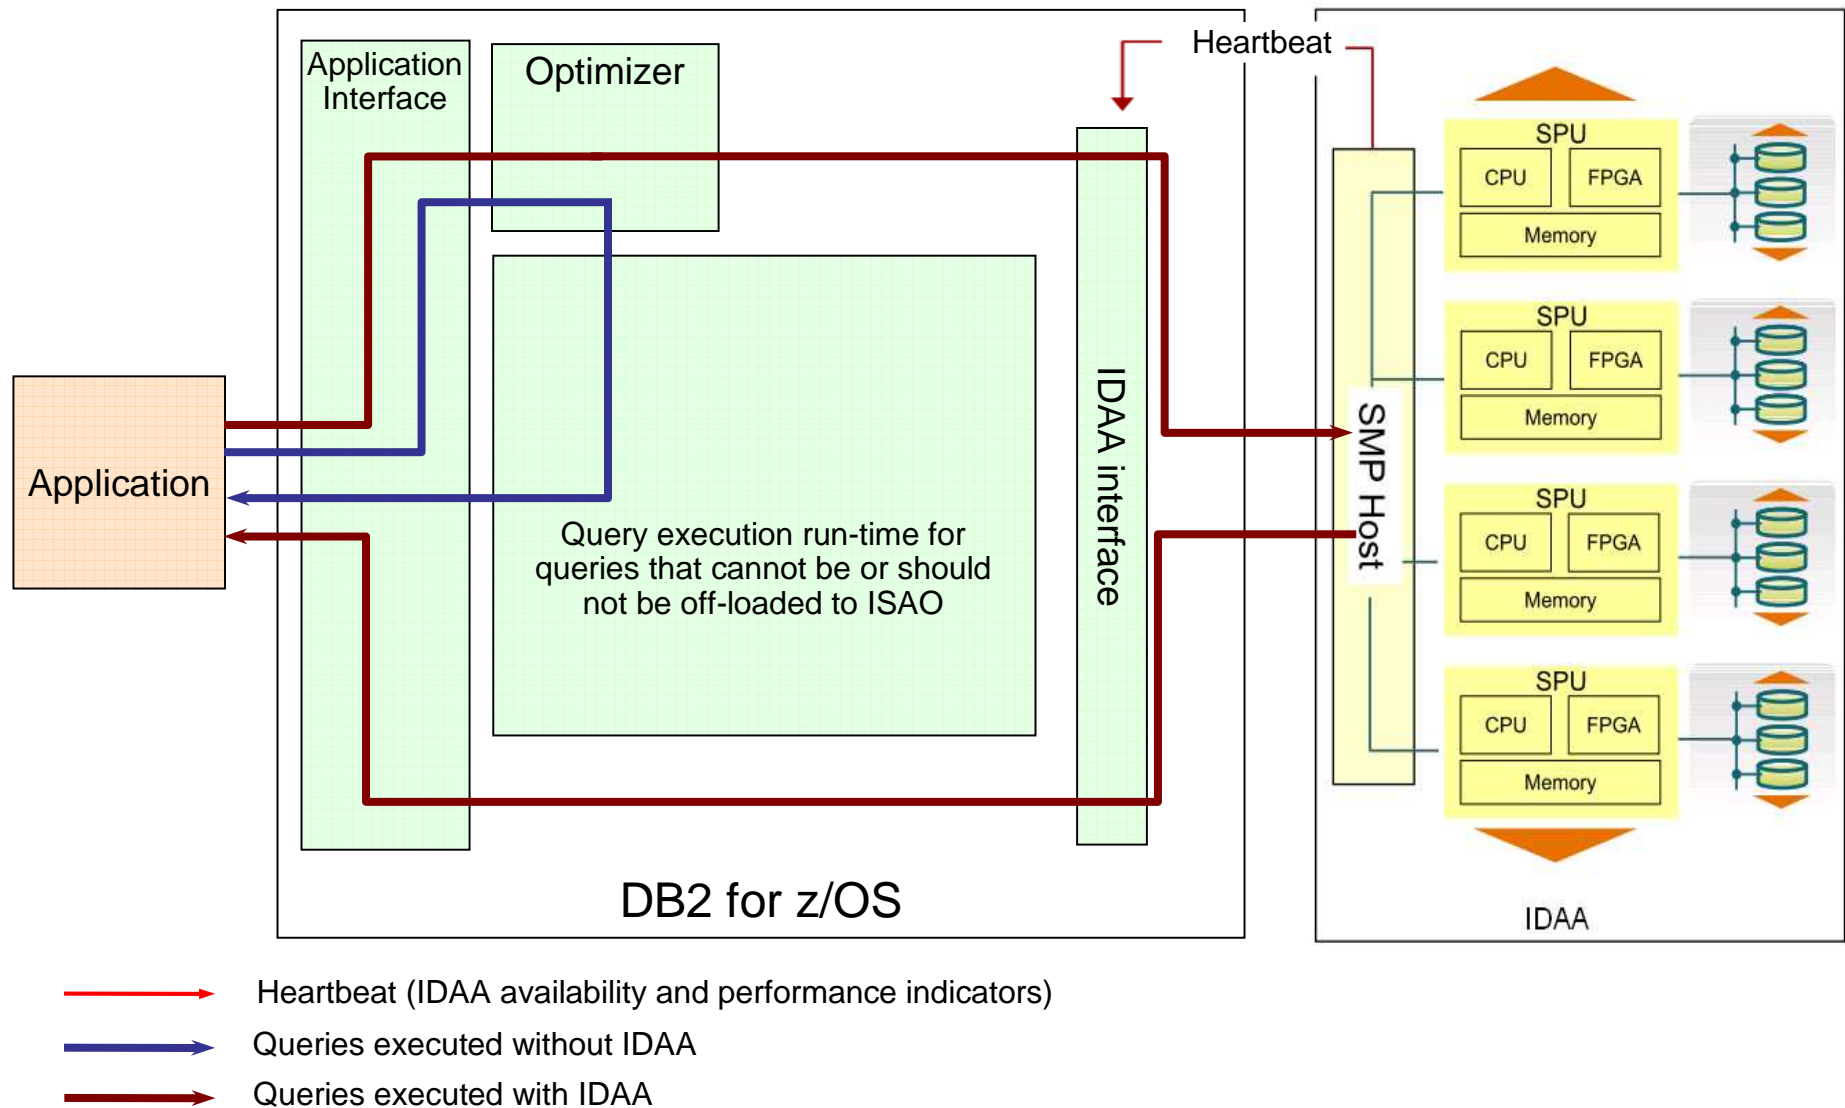
- **Performance:** Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- **Integration:** Connects to DB2 through deep integration providing transparency to all applications.
- **Self-managed workloads:** queries are executed in the most efficient way
- **Transparency:** applications connected to DB2 are entirely unaware of the Optimizer
- **Simplified administration:** appliance hands-free operations, eliminating many database tuning tasks

Breakthrough Technology Enabling New Opportunities

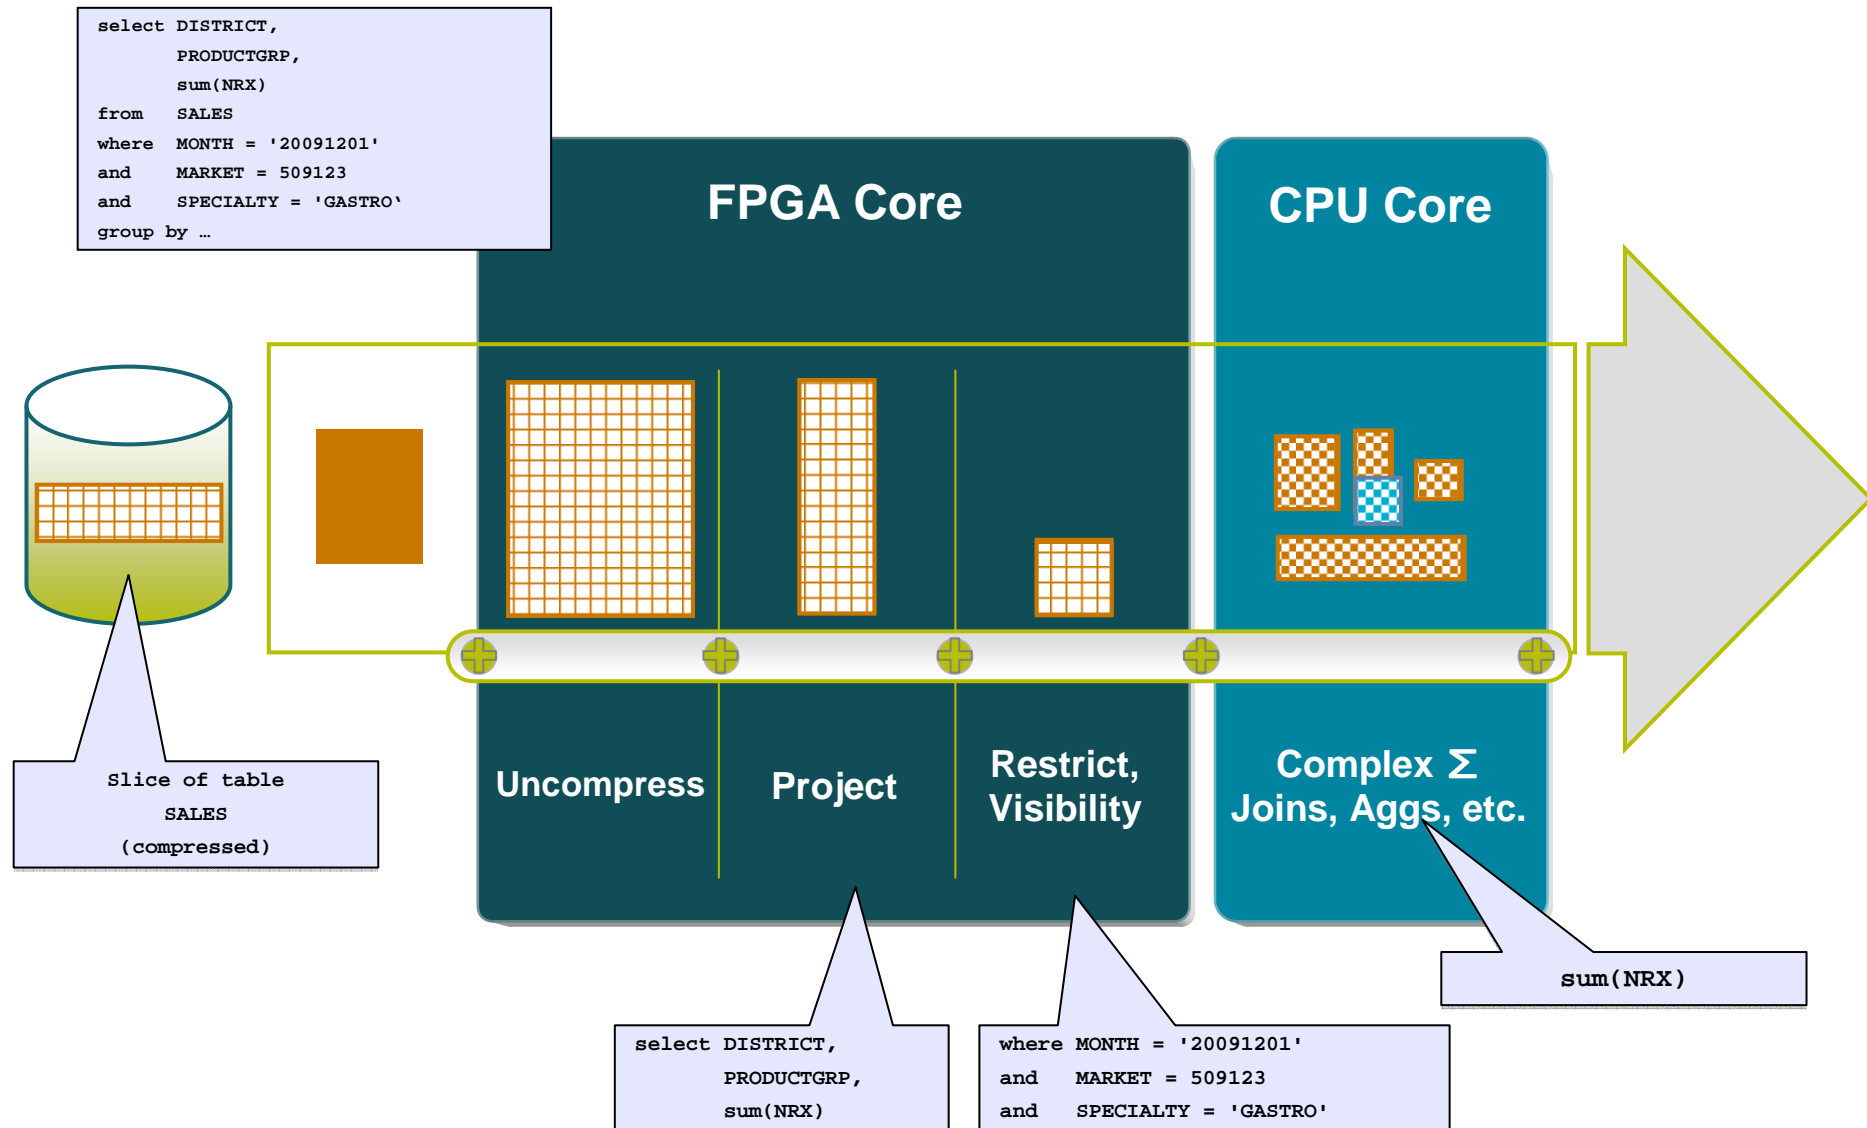
Deep DB2 Integration within zEnterprise



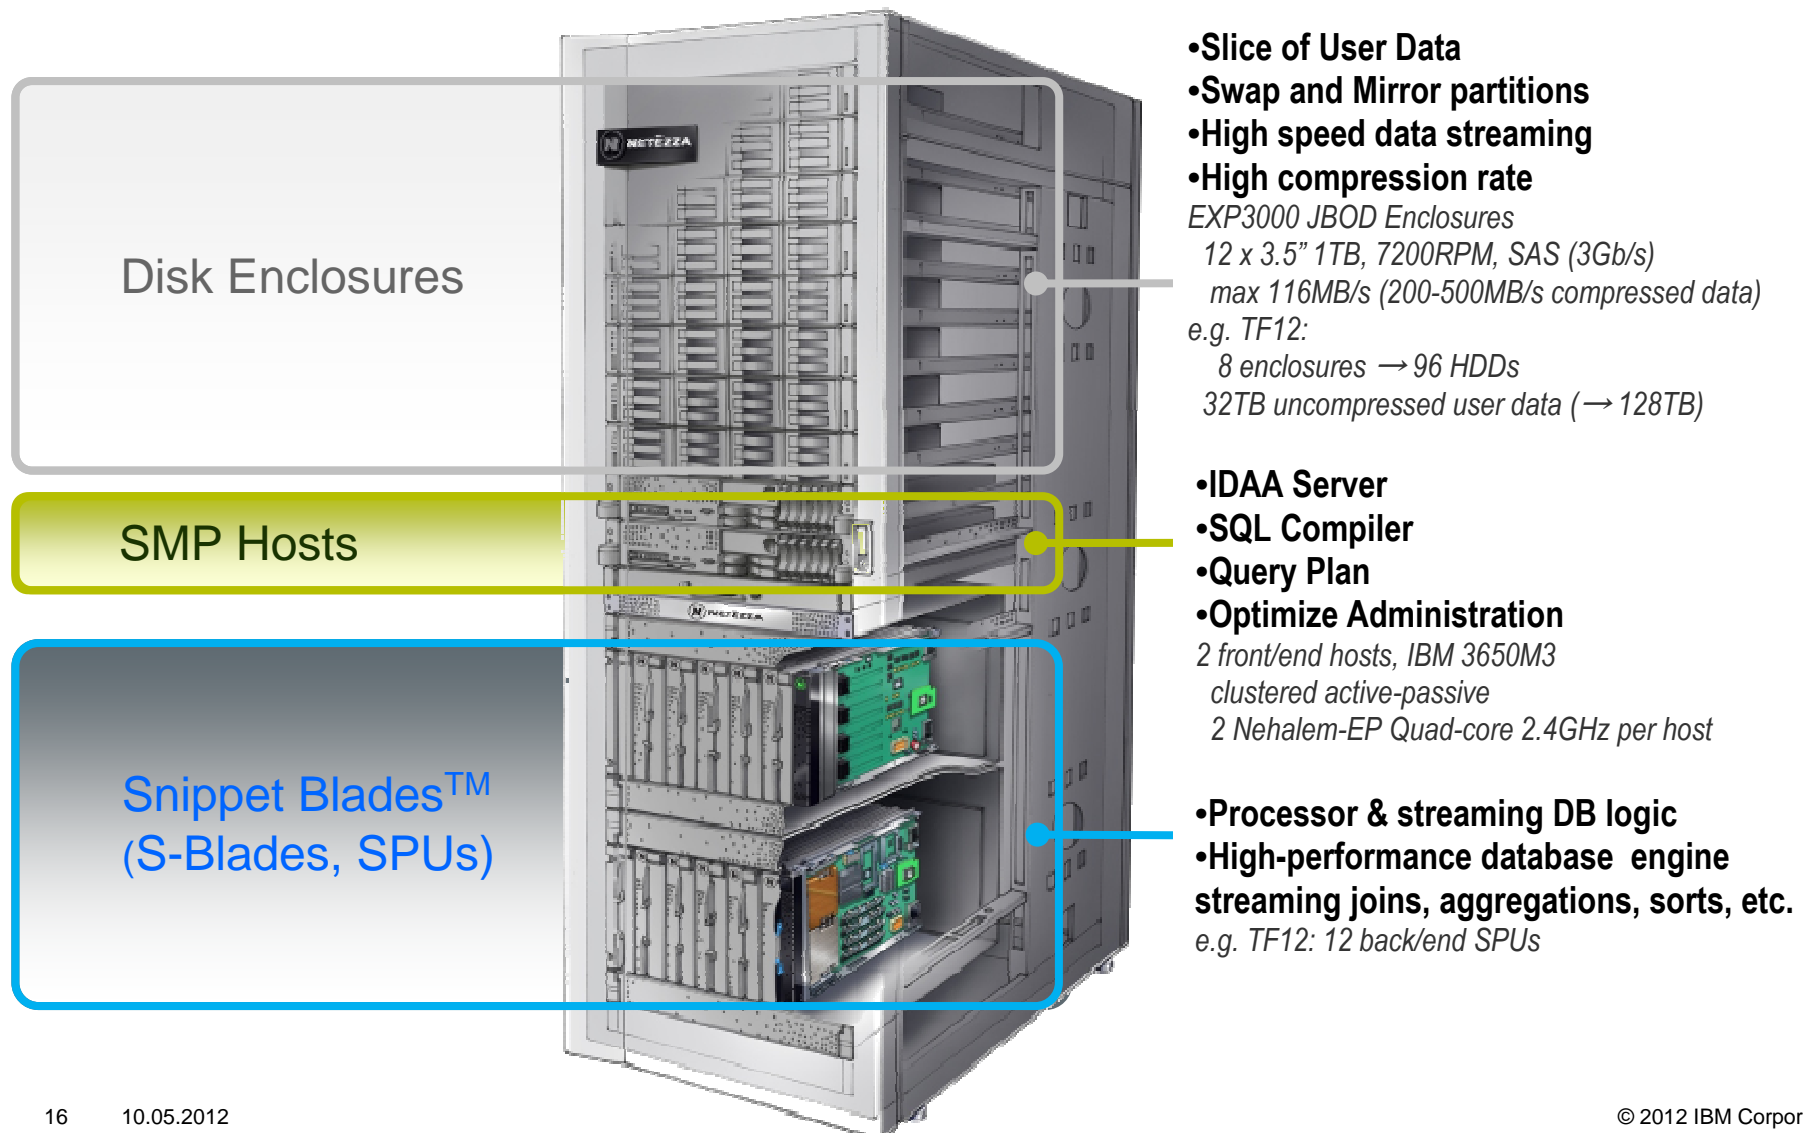
Query Execution Process Flow




The Key to the Speed




Netezza 1000 basic design aspects



The Netezza S-Blade™



Disk Mirroring and Failover



- All user data and temp space mirrored
- Disk failures transparent to queries and transactions
- Failed drives automatically regenerated
- Bad sectors automatically rewritten or relocated

Connectivity Options

Multiple DB2 systems can connect to a single IDAA



A single DB2 system can connect to multiple IDAAs



Multiple DB2 systems can connect to multiple IDAAs



Better utilization of IDAA resources
Scalability
High availability

Full flexibility for DB2 systems:


- residing in the same LPAR
- residing in different LPARs
- residing in different CECs
- being independent (non-data sharing)
- belonging to the same data sharing group
- belonging to different data sharing groups

Network Configuration Options

Option 1 – Simple Direct attachment

- Virtual IP definition both on System z and Netezza
- Only one network link active at a time – if Netezza fails over to standby host, connection might get lost


OSA ports are configured such that only one link is active at a time. If the connection on this link breaks, the other is activated



Option 2 – Additional redundancy or additional CEC requires Switch(es)

- Can address cable failures and Netezza fail over to standby host
- For higher availability requirements, a second switch is required


OSA ports are configured such that packages are sent alternating, using both ports: “multipathperconnection” configuration



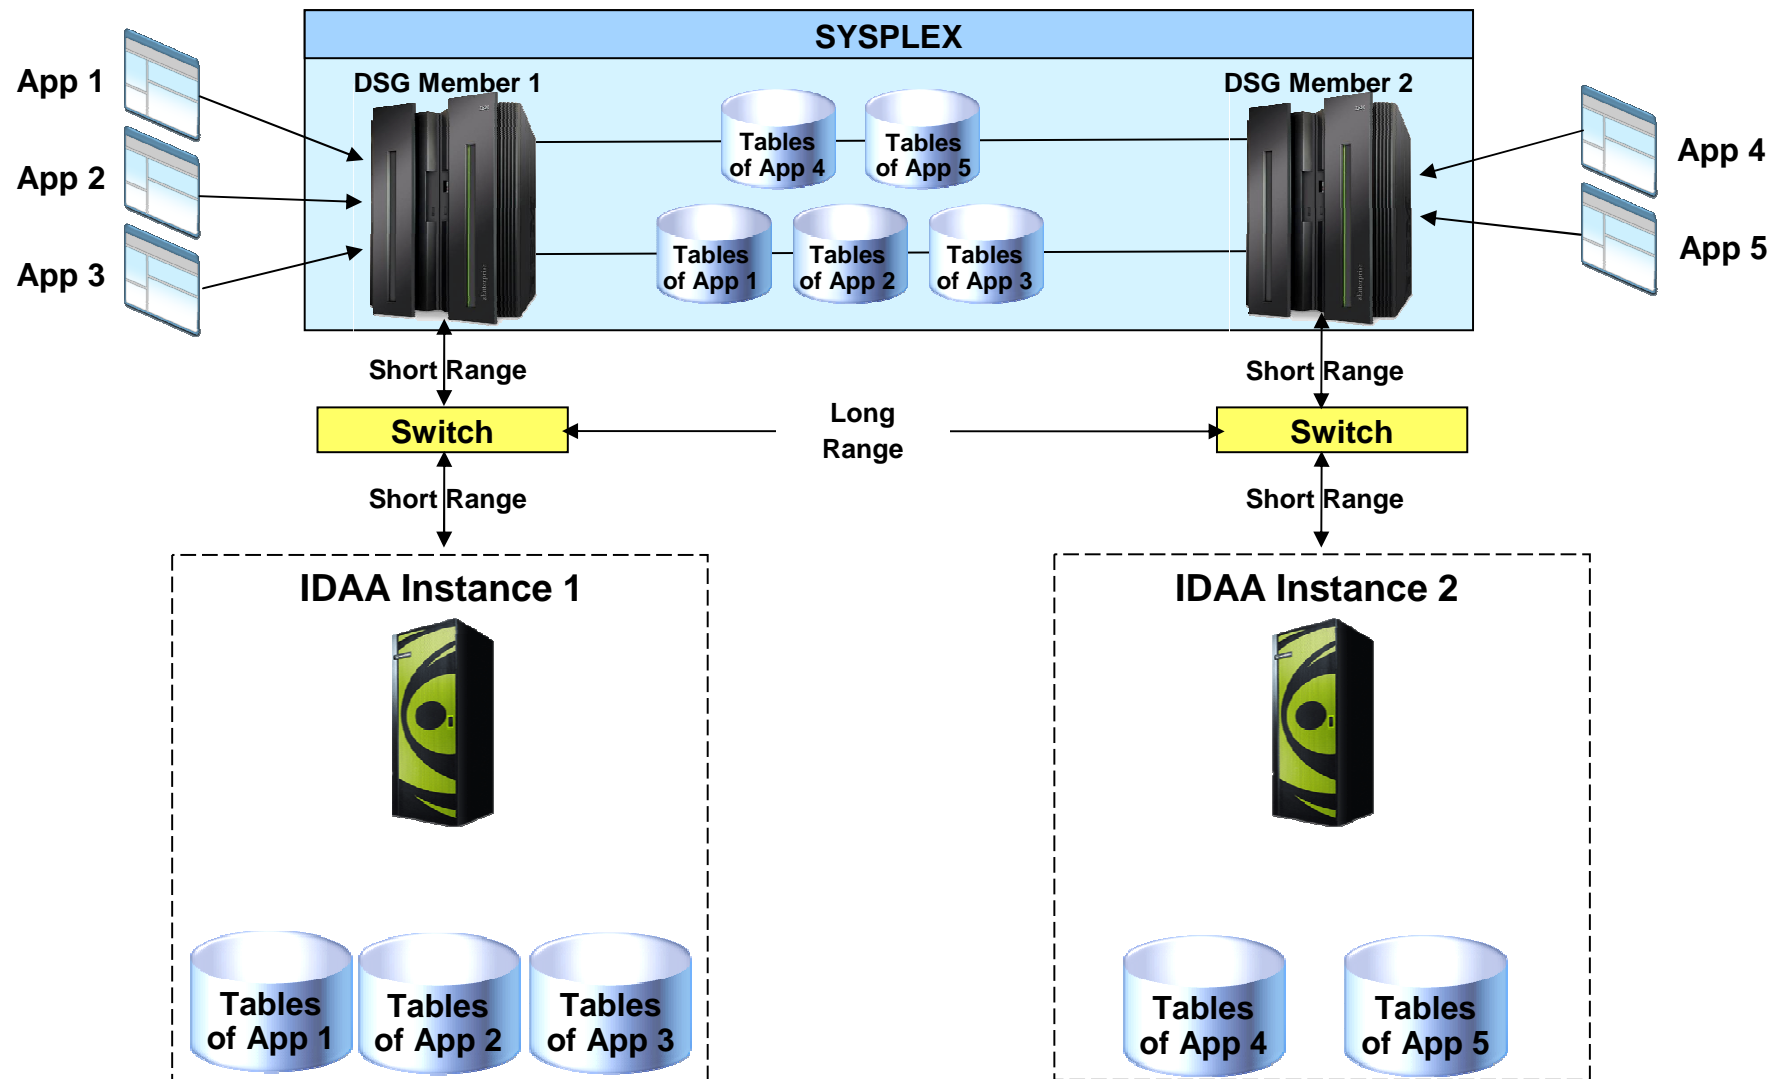
Network Configuration Options *(continued)*

Option 3 – zBX TOR Switch

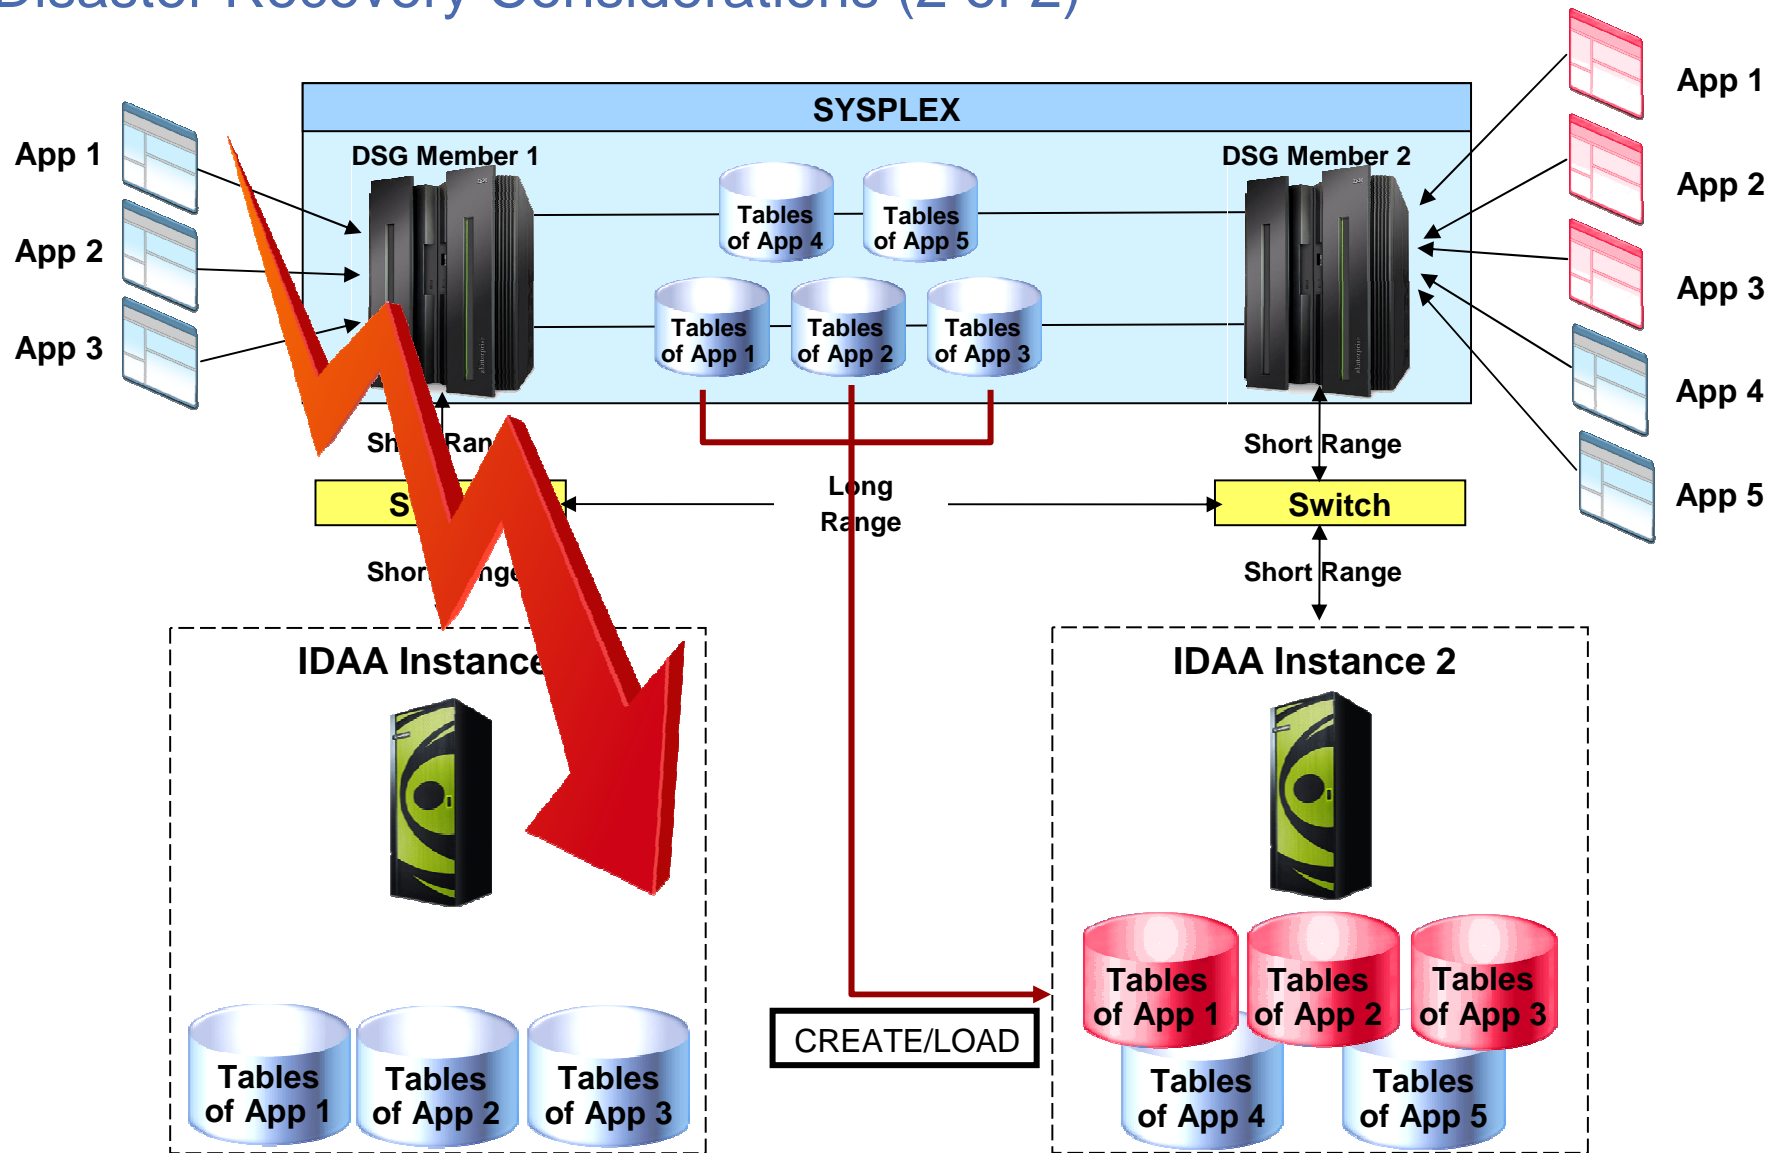
- For clients with an installed zBX connected to zEnterprise, the top-of-rack switch may be leveraged to connect to the IDAA
- Connection between zBX and IDAA can be direct or switched



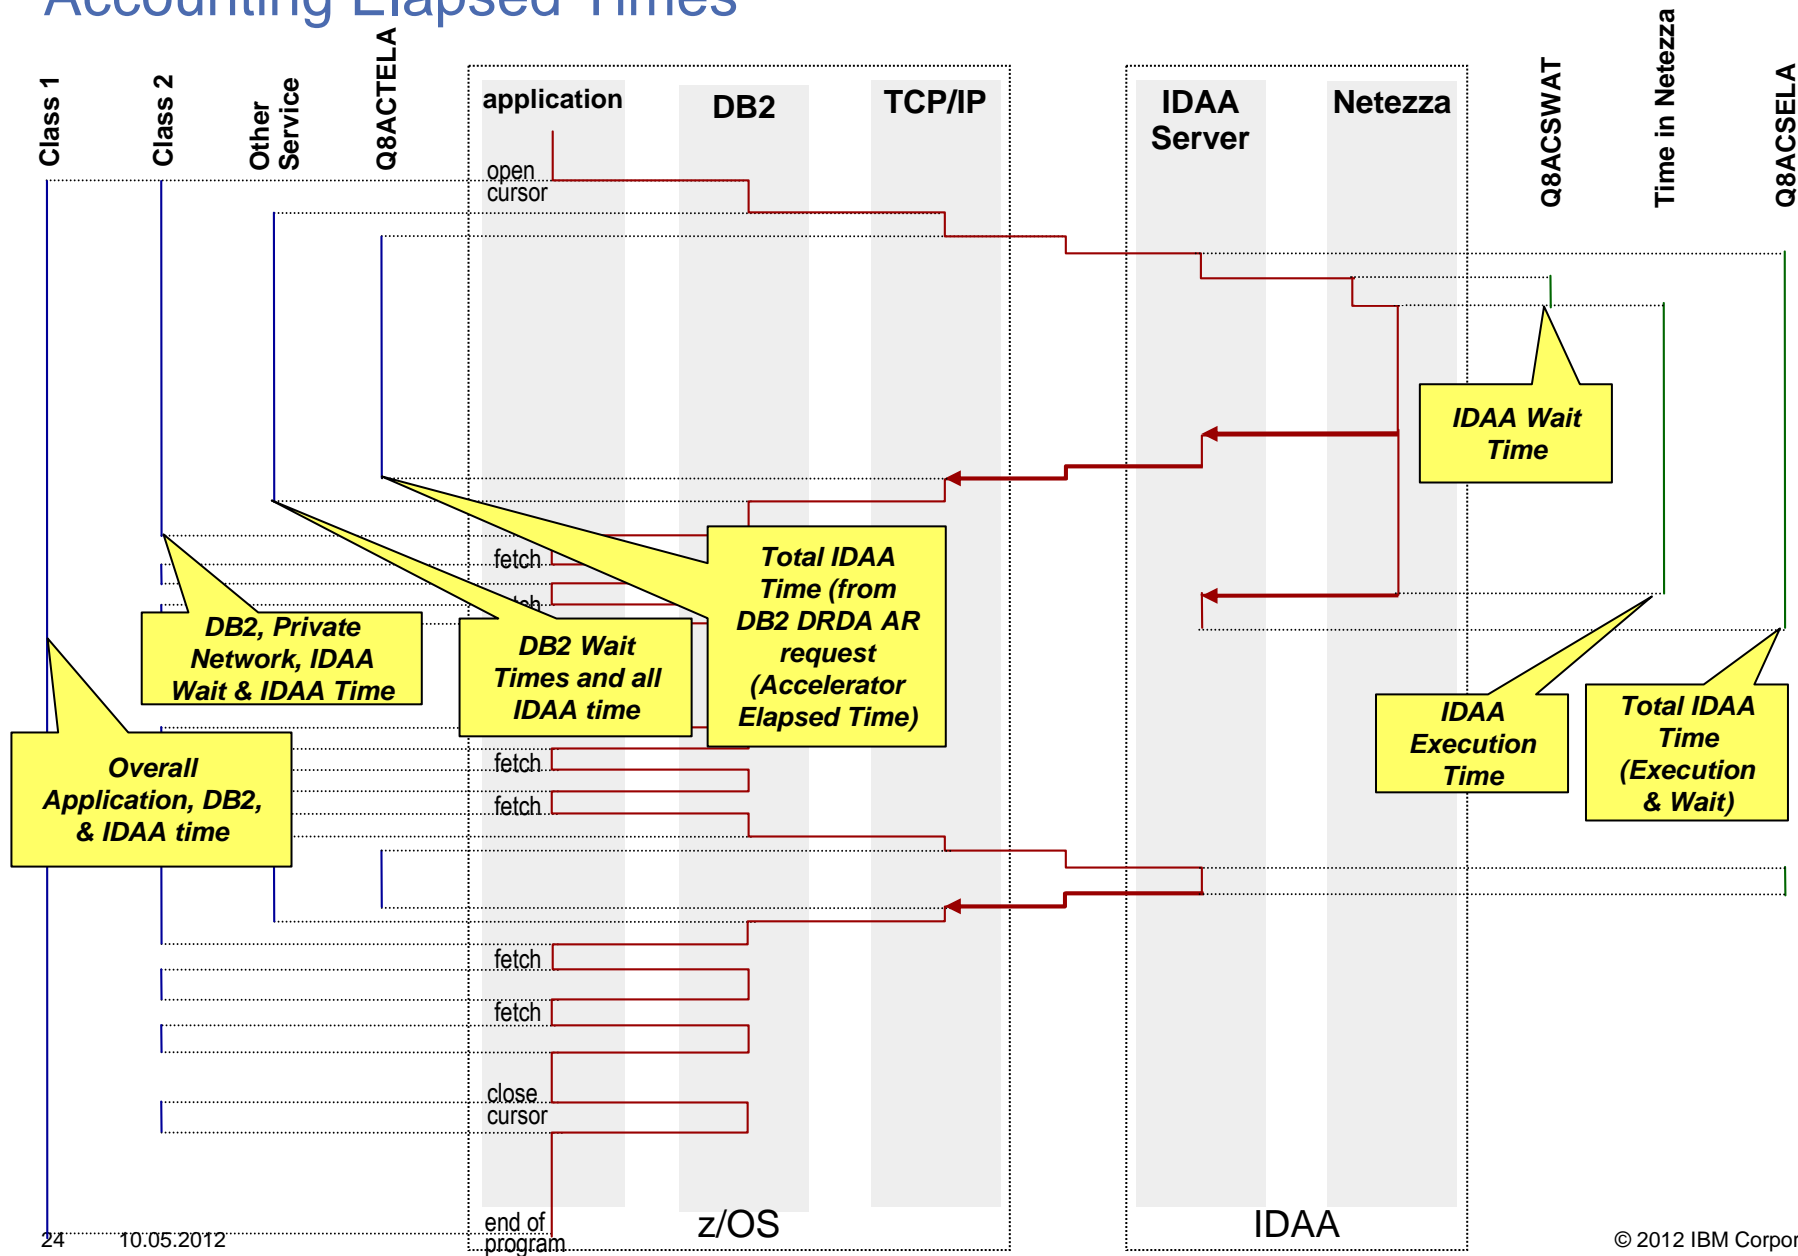
Disaster Recovery Considerations (1 of 2)



Disaster Recovery Considerations (2 of 2)



Accounting Elapsed Times



System Scope Instrumentation

```

1  LOCATION: PMOV91A                OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1)                PAGE: 1-23
    GROUP: N/P                      STATISTICS REPORT - LONG                                REQUESTED FROM: NOT SPECIFIED
    MEMBER: N/P                                                              TO: NOT SPECIFIED
    SUBSYSTEM: V91A                                                         INTERVAL FROM: 09/06/11 21:49:41.35
    DB2 VERSION: V9                                                         TO: 09/06/11 23:41:50.70
                                   SCOPE: MEMBER
  
```

```

----- HIGHLIGHTS -----
INTERVAL START : 09/06/11 21:49:41.35    SAMPLING START: 09/06/11 21:49:41.35    TOTAL THREADS      :    76.00
INTERVAL END   : 09/06/11 23:41:50.70    SAMPLING END   : 09/06/11 23:41:50.70    TOTAL COMMITS      :   109.00
INTERVAL ELAPSED:    1:50:52.248034    OUTAGE ELAPSED:    1:17.097273    DATA SHARING MEMBER:    N/A
  
```

ZGRYPHON	ACCELERATOR	QUANTITY	ZGRYPHON	CONTINUED	QUANTITY
CONNECTS		2.00	AVG QRY QUEUE LEN (3 HRS)		0.00
REQUESTS		9.00	AVG QRY QUEUE LEN (24 HRS)		0.00
REQUESTS TIMED OUT		0.00	HWM QRY QUEUE LENGTH		0.00
REQUESTS FAILED		0.00	DATA SKEW		0.00
BYTES SENT	4630.00		AVG QUEUE WAIT ELAPSED TIME		0.000000
BYTES RECEIVED	224887.00		MAX QUEUE WAIT ELAPSED TIME		0.000000
MESSAGES SENT	27.00		PROCESSING CAPACITY		0.00
MESSAGES RECEIVED	27.00		PROCESSORS		1.62
BLOCKS SENT	0.00		QUERY REQUESTS SUCCESSFUL		1.00
BLOCKS RECEIVED	5.00		QUERY REQUESTS FAILED		1.00
ROWS SENT	0.00		QUERY REQUESTS INVALID		0.00
ROWS RECEIVED	0.00		SHR MEM WORKER NODES (MB)		0.00
SVCS TCP/IP ELAPSED TIME	7.036035		AVG IN USE (MB)		0.00
ACCELERATOR CPU TIME	0.000000		MAX IN USE (MB)		0.00
ACCELERATOR ELAPSED TIME	0.000001				
ACCELERATOR WAIT TIME	0.000000				
CUR ACTIVE REQUESTS	0.00		DISK STORAGE AVAILABLE (MB)		98842.63
MAX ACTIVE REQUESTS	0.00		IN USE (MB)		0.81
			IN USE FOR DB (MB)		0.81
			DATA SLICES		3.25
			MEM COORD AVG IN USE (MB)		0.00
			MEM WORKER AVG IN USE (MB)		0.00

Thread Scope Instrumentation

LOCATION: PMOV91A
GROUP: N/P
MEMBER: N/P
SUBSYSTEM: V91A
DB2 VERSION: V9

OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1)
ACCOUNTING REPORT - LONG
ORDER: PRIMAUTH-PLANNAME
SCOPE: MEMBER

PAGE: 1-7
REQUESTED FROM: NOT SPECIFIED
TO: NOT SPECIFIED
INTERVAL FROM: 09/06/11 21:52:08.31
TO: 09/06/11 23:43:26.65

PRIMAUTH: ADMF001 PLANNAME: db2jcc_a

--- DISTRIBUTED ACTIVITY ---

SERVER	: ZGRYPHON	CONVERSATIONS INITIATED:	1.00	#COMMT(1) SENT:	0	MESSAGES SENT	:	14.33
PRODUCT ID	: AQT	#CONVERSATIONS QUEUED	: 0	#ROLLB(1) SENT:	0	MESSAGES RECEIVED:		14.33
METHOD	: DRDA PROTOCOL	CONVERSATION TERMINATED:	N/A	SQL SENT	: 6.33	BYTES SENT	:	2403.33
REQUESTER ELAP.TIME:	2.402285	#RLUP THREADS	: N/A	ROWS RECEIVED:	0.00	BYTES RECEIVED	:	138728.33
SERVER ELAPSED TIME:	N/A					BLOCKS RECEIVED	:	3.33
SERVER CPU TIME	: N/A							
DBAT WAITING TIME	: N/A							
#DDF ACCESSES	: 3							
#COMMIT(2) SENT	: 0	#BACKOUT(2) SENT	: 0	#BKOUT(2) R.R:	0	#LASTAGN.SENT	:	0
		SUCCESSFULLY ALLOC.CONV:	N/A	TRANSACT.SENT:	1.00	STMT BOUND AT SER:		N/A
		MAX OPEN CONVERSATIONS	: N/A	MSG.IN BUFFER:	0.00	#FORGET RECEIVED	:	0
		#CONT->LIM.BL.FTCH SWCH:	N/A	#PREPARE SENT:	0			
		#COMMIT(2) RESP.RECV.	: 0					

ACCELERATOR	IDENTIFIER	ACCELERATOR	TOTAL VALUE	TOTAL TIME	AVERAGE VALUE	AVERAGE TIME
PRODUCT	AQT02010	OCCURRENCES	3		1.00	
SERVER	ZGRYPHON	CONNECTS	3		1.00	
		REQUESTS	16		5.33	
		TIMED OUT	0		0.00	
		FAILED	0		0.00	
		SENT				
		BYTES	7210		2403.33	
		MESSAGES	43		14.33	
		BLOCKS	0		0.00	
		ROWS	0		0.00	
		RECEIVED				
		BYTES	416185		138728.33	
		MESSAGES	43		14.33	
		BLOCKS	10		3.33	
		ROWS	0		0.00	

ELAPSED TIME		
SVCS TCP/IP	7.206857	2.402286
ACCUM ACCEL	0.000061	0.000020
CPU TIME		
SVCS TCP/IP	0.015023	0.005008
ACCUM ACCEL	0.000000	0.000000
WAIT TIME		
ACCUM ACCEL	0.000000	0.000000

Online Monitoring with Performance Expert Client

Accelerator

Accelerator name	Product ID	Connects to accelerator	Requests sent to accelerator	Successfully executed
TF12	AQT02012			

Main Accelerator: TF12

Accelerator status

Accelerator status

Currently active queries	5	Disk storage available (MB)	33,557,184
Maximum active queries	0	In use (%)	10.14
Successfully executed	0	In use for database (MB)	350,695
Failed	0	Data slices	92
Invalid state	0	Data skew (%)	531.68

Main Accelerator: TF12

Accelerator status

Accelerator

Accelerator name	TF12	Sent	Received
Product ID	AQT02012	Bytes	0
Connects to accelerator	0	Messages	0
Requests sent to accelerator	N/P CTRL	Blocks	0
Successfully executed	0	Rows	0
Timed out	0		
Failed	0		

TCP/IP services elapsed time	0.000000
Elapsed time in accelerator	0.000000
CPU time spent in accelerator	0.000000
Wait time in accelerator	0.000000
Processors in accelerator	96

Accelerator status

ONLINE

Synchronization Options with DB2 Analytics Accelerator

Synchronization options	Use cases, characteristics and requirements
Full table refresh The entire content of a database table is refreshed for accelerator processing	<ul style="list-style-type: none"> ▪ Existing ETL process replaces entire table ▪ Multiple sources or complex transformations ▪ Smaller, un-partitioned tables ▪ Reporting based on consistent snapshot (“check point”)
Table partition refresh For a partitioned database table, selected partitions can be refreshed for accelerator processing	<ul style="list-style-type: none"> ▪ Optimization for (time-) partitioned warehouse tables, appending changes “at the end” ▪ More efficient than full table refresh for larger tables ▪ Reporting based on consistent snapshot (“check point”)
Incremental update (controlled availability) Log-based capturing of changes and propagation to IDAA with low latency (typically few minutes) BETA Phase ongoing	<ul style="list-style-type: none"> ▪ Scattered updates after “bulk” load ▪ Reporting on continuously updated data (e.g., an ODS), considering most recent changes ▪ More efficient for smaller updates than full table refresh

Why Both?

Marrying the best of each

IBM Netezza



Focused Appliance

IBM System z



Mixed Workload System

Capitalizing on the strengths of both platforms while driving to the most cost effective, centralized solution - destroying the myth that transaction and decision systems had to be on separate platforms

Very focused workload

Very diverse workload

The value of a hybrid solution

The best of both worlds

IBM Netezza

Focused Appliance

- High performance analytic queries without tuning by a Database Administrator
- No storage administration
- Simple install requires no additional professional services
- Fastest possible deployment
- High speed batch ingest and 100s of queries per second
- FPGA acceleration
- In-database Mining
- Native MDX OLAP







IBM System z Hybrid Solution with IBM DB2 Analytics Accelerator

Mixed Workload System

- System z attributes for ALL your data (OLTP and DW)
 - Security,
 - Availability,
 - Reliability
 - Recoverability
- Application transparency with simple integration into existing environment
- Fastest performance fit for purpose query engines:
 - DB2 for transactional
 - Accelerator delivers high performance analytic queries without tuning by a Database Administrator
- Cost savings
 - Consolidation
 - Reduced software licenses
 - Single security policies
 - Uniform database management
 - Common skills
 - Common tools
 - Common procedures.

ISAO V1 Needed Following Enhancements

IDAA addresses all of them!

-  1. Increase applicability by relaxing current off-load restrictions
See further documentation for details
-  2. Increase applicability by supporting larger amount of data
Up to 32TB of uncompressed data,
e.g. with 1:4 compression ratio, up to 128TB of user data
-  3. Support concurrent query execution
Exploiting Netezza workload management capabilities
-  4. Improve data currency
Partition-scope update
-  5. Support disaster recovery
Building blocks provided
-  6. DB2 10 support
IDAA supports both DB2 9 and DB2 10*
Requires zEnterprise (z196 or z114)

Feedback from Beta Customer: Fast Time to Value

- IBM DB2 Analytics Accelerator (Netezza 1000-12)
 - ➔ Production ready - 1 person, 2 days
- Table Acceleration Setup ... **2 Hours**
 - DB2 “Add Accelerator”
 - Choose a Table for “Acceleration”
 - Load the Table (DB2 copy to Netezza)
 - Knowledge Transfer
 - Query Comparisons
- Initial Load Performance ...
 - ➔ 5.1 GB in 1 Min 25 Seconds (24M rows)
 - 400 GB in 29 Min (570M rows)
- Actual Query Acceleration ... **1908x faster**
 - ➔ 2 Hours 39 Minutes to 5 Seconds
- CPU Utilization Reduction ... **99% less CPU**
 - ➔ 24M rows: 56.5 CPU seconds to 0.4 CPU seconds



Large Insurance Company

Adding value by Accelerating the Delivery of Business Reporting

				DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Qualifying Rows	Total Rows Returned	Hours	Sec(s)	Hours	Sec(s)	
Query 1	591,941,065	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	591,941,065	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	813,343,052	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	283,105,125	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	591,941,089	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	813,343,052	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	591,941,065	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	813,343,052	3,425,292	724	0:44	2,640	0.0	2	1,320
Query 9	813,343,052	4,130,107	137	0:42	2,520	0.1	193	13

With Accelerated Time to Value

- **IBM DB2 Analytics Accelerator (Netezza 1000-12)**
Production ready - 1 person, 2 days
- **Table Acceleration Setup in 2 Hours**
 - DB2 “Add Accelerator”
 - Choose a Table for “Acceleration”
 - Load the Table (DB2 Loads Data to the Accelerator)
 - Knowledge Transfer
 - Query Comparisons
- **Initial Load Performance**
400 GB Loaded in 29 Minutes
570 Million Rows (Actual: Loaded 800 GB to 1.3 TB per hour)
- **Extreme Query Acceleration - 1908x faster**
2 Hours 39 minutes to 5 Seconds
- **CPU Utilization Reduction**
Up to 35%

Customer Quote: “we had this up and running in days with queries that ran over 1000 times faster”

Swiss Re - IBM System z Client Success Video



http://www.youtube.com/watch?v=xkcp_pJxT5E

You can also benefit!

- Are there **long running queries** that could provide business value if they could be run in seconds vs **30 minutes or more**?
- What about **performance challenges** with **complex and ad hoc queries**?
- Is **modernization of the data warehouse** or **Operational Business Analytics** a topic of interest?
- Are there thoughts about **extending the use of operational platform data** to perform **business analysis** and **daily reporting**?
- Is there **a data warehouse running on System z** or an intention to do so?
- Is an OLAP application **running out of steam** due to growth in data. These are usually a single subject area such as Accounting, Sales or Inventory.
- **The forgotten query**: Have queries been elected to set aside due to performance challenges?
- Is there a **System z/196 or z114** or a plan for one in the near future?

If you answer one or more of the following questions with 'yes' please give me another 3 minutes to show you how we can help you analyzing your acceleration potencial



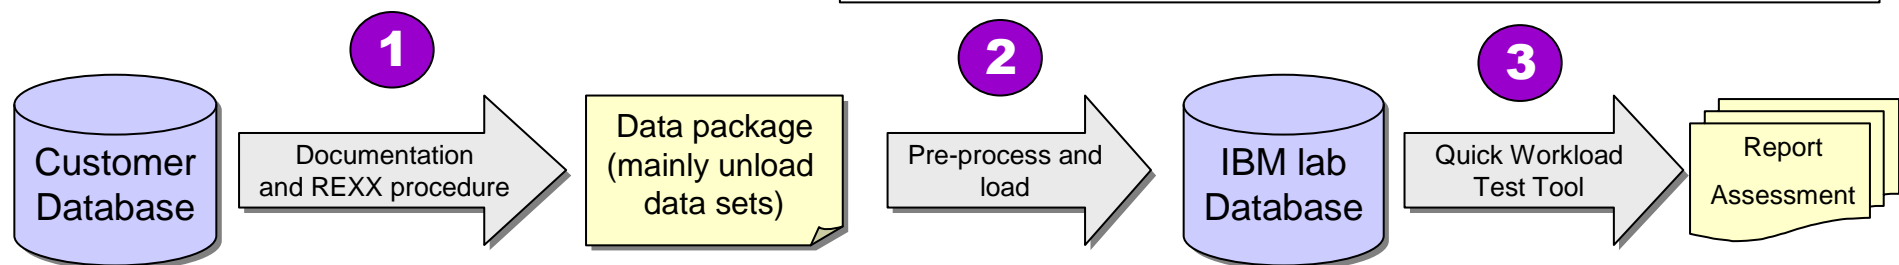
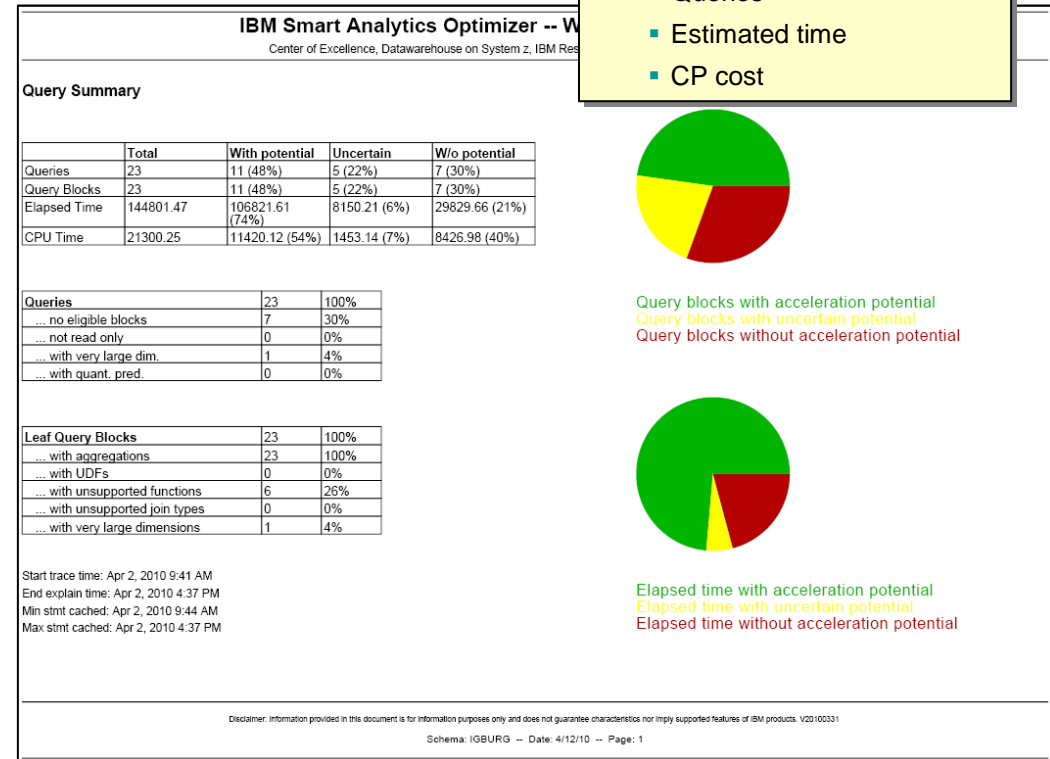
Quick Workload Test

Customer

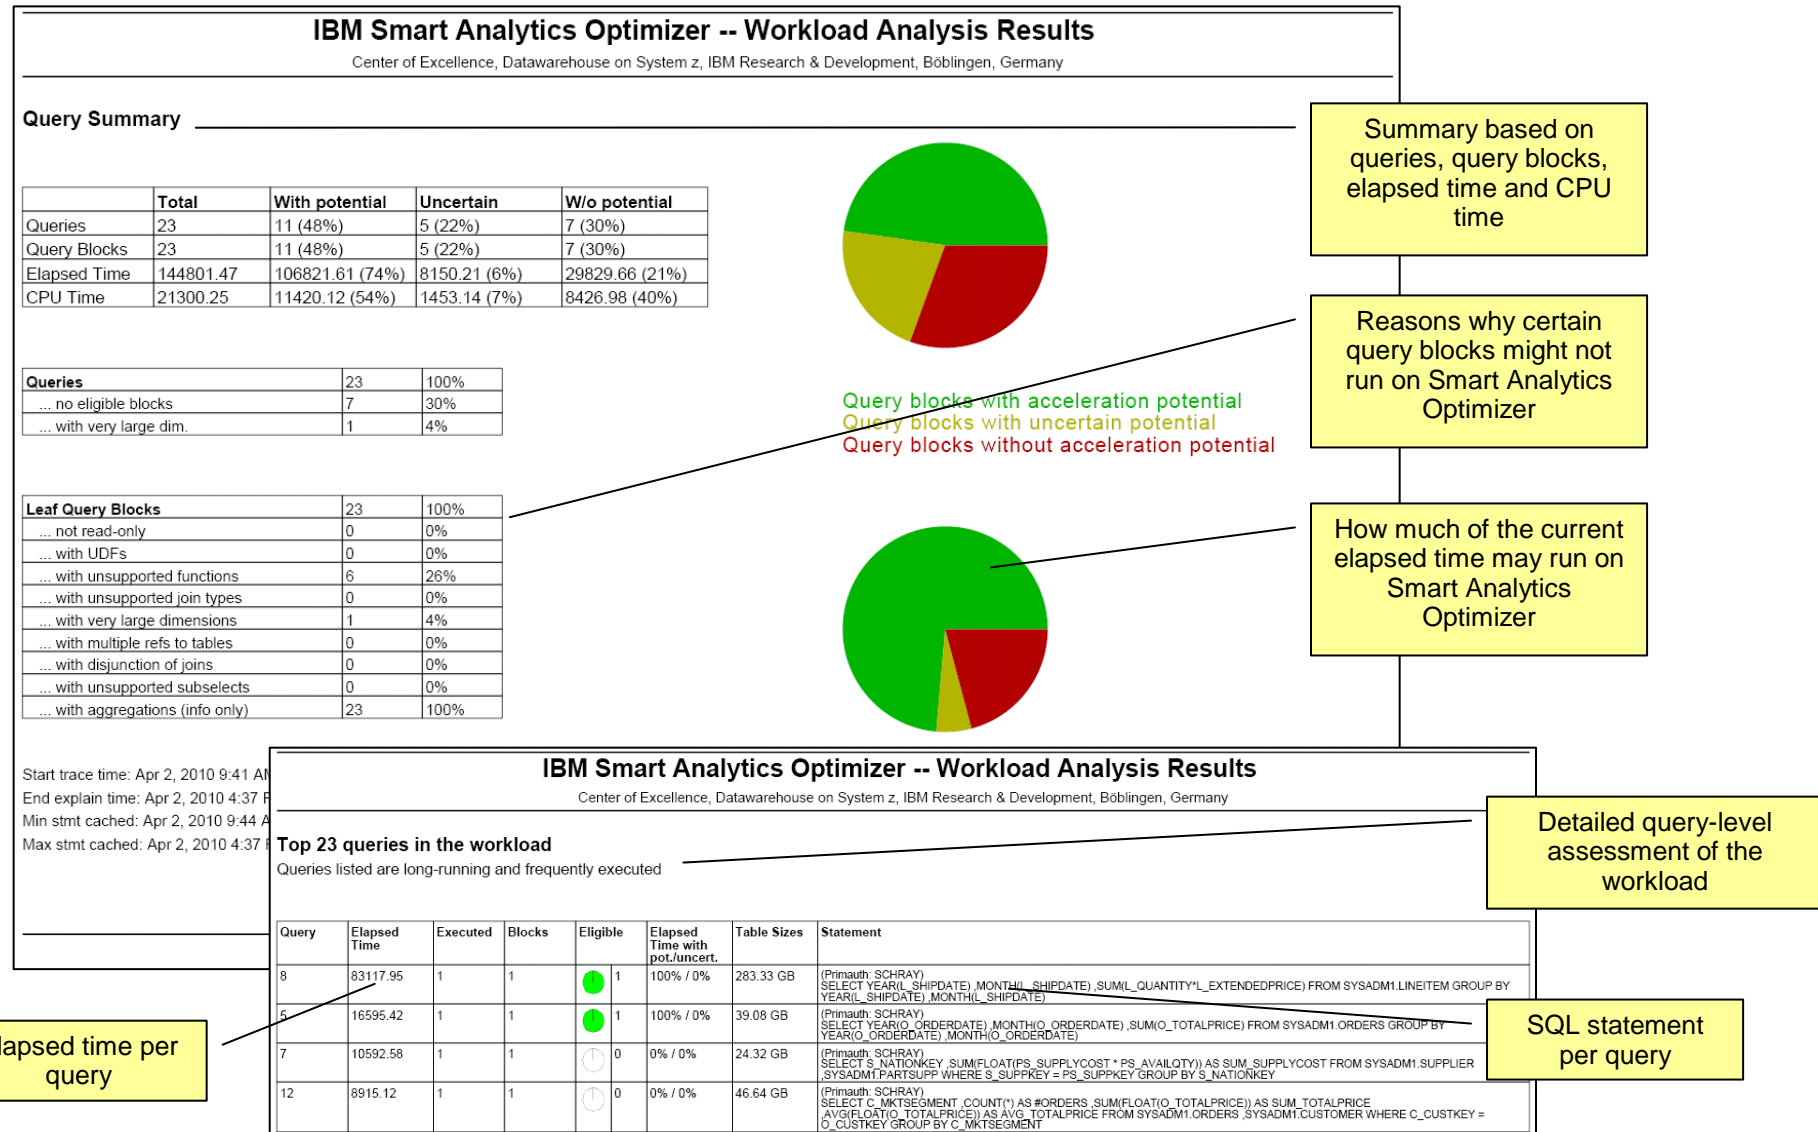
- Collecting information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
- Uploading compressed file (up to some MB) to IBM FTP server

IBM / Center of Excellence

- Importing data into local database
- Quick analysis based on known DB2 Analytics Accelerator capabilities



Workload Assessment Output: Sample PDF Report



Demo



BI mixed Workload with IBM DB2 Analytics Accelerator for z/OS



Demo time in minutes:
Start Queries Stop Queries Reset Demo
Count-down seconds: 0

5 ▾ Concurrent call center users - operational BI

Run

1

2

3


IDAA status

Concurrent users

Queries started

Queries completed

Avg. resp. time (s)



0 ▾ Concurrent power users - complex ad-hoc reports

Run

1

2

3


IDAA status

Concurrent users

Reports started

Reports completed

Avg. resp. time (s)



Setup

SYS1*,PROCESSOR -- % CPU utilization (CP) [8D0460]

Time Range: 04/13/2012 13:55:45 - 04/13/2012 13:56:00

0

SYS1*,PROCESSOR -- # CP processors online [8D0D20]

Time Range: 04/13/2012 13:55:45 - 04/13/2012 13:56:00

3

SYS1*,PROCESSOR -- % IIP on CP [8D3550]

Time Range: 04/13/2012 13:55:45 - 04/13/2012 13:56:00

0

SYS1*,I/O_SUBSYSTEM -- i/o activity rate [8D0E90]

Time Range: 04/13/2012 13:55:45 - 04/13/2012 13:56:00

NaN

DB2 Analytics Accelerator Status: disabled

ACCELERATOR	MEMB	STATUS	REQUESTS	ACTV	QUED	MAXQ
DEM0IDAA	DSN9	STOPPED	17	0	0	0

DISPLAY ACCEL REPORT COMPLETE
 DSN9022I -DSN9 DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION

Enable Accelerator Disable Accelerator Display Status

38 10.05.2012

© 2012 IBM Corporation



Thank You



Patrick Hempeler

*Client Technical Professional
Pan Europe DWH/BI on System z
IBM Software Group*

*IBM-Allee 1
D-71139 Ehningen*

*Mobil: 0151 / 1513 8920
Email: patrick.hempeler@de.ibm.com*

