



***The Modern Mainframe...
At the Heart of Your Business***

Consolidating Data on System z



© 2006 IBM Corporation

Why Put Data on DB2 for z/OS?

Corporate data is crucial to our next generation solutions.
Oracle RAC claims to provide a lower cost solution.



**Service Oriented Finance
CIO**

Lets see why the world's largest corporations rely on DB2 for z/OS.




IBM

DB2 Proven Success in the Finance Industry

- DB2 for z/OS is in 25 of the top 25 banks worldwide
- Why?
 - ▶ Highest Scalability – Near-linear scalability and workload management
 - ▶ Highest Availability – DB2 provides nearly continuous 99.999 availability.
 - ▶ Proven Security and Compliance – RACF, Encryption, DB2 Audit Management Expert
 - ▶ Better Support of Current Technology Trends – Native XML support and SAP optimization
 - ▶ Lowest overall TCO


Comparison of Data Sharing Architectures

DB2 for z/OS Centralized Sysplex Design



High speed centralized
lock manager in
coupling facility

Oracle RAC Distributed Lock and Data Design



Distributed lock
management with
high messaging **overhead**

The DB2 Sysplex Design Scales Linearly


- Beyond two members, DB2 clustering overhead is low, usually between ½% and 1% for each additional image
 - ▶ Coupling facility processor handles the workload of lock and cache management
 - ▶ Hardware invalidates local copies without local processor interrupt
 - ▶ Result is near linear scale out

- Oracle RAC does not scale well beyond 4 to 6 nodes
 - ▶ The local processor overhead grows as nodes are added
 - ▶ More overhead means less transaction throughput per local processor
 - ▶ Result is limited scale out

06 - Consolidating Data on System z v2.0.ppt

6

Centralized Coupling Facility Permits Efficient Lock and Cache Management in DB2



A, B, and C have read locks with local copies

1. B Obtains write lock
2. B Updates local copy
3. B Caches update in buffer pool
4. CF invalidates all cached copies without interrupting processors

Cache and locks are maintained with no inter-node disturbance!

06 - Consolidating Data on System z v2.0.ppt

7

Why is Oracle RAC Scalability Limited?

RAC Inefficiencies increase as a cluster grows

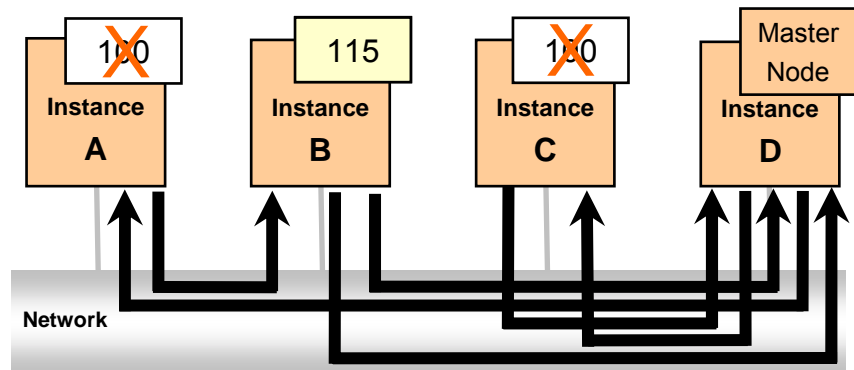
- RAC Nodes must constantly communicate to process requests to maintain distributed cache and lock data
- Adding additional nodes to the cluster results in increased inter-node communication which requires additional local processor and network time
- RAC distributed lock management overhead increases faster than the added capacity of more nodes

Let's look at some examples...

06 - Consolidating Data on System z v2.0.ppt

8

Oracle RAC: Lock Management Overhead



Lock Assume

7. B Updates local copy

Inter-node connections: 6

In a cluster with 4 nodes, an update operation may need 6 network connections and two in-memory calls (not shown).

Example based on Oracle's US Patent 7,107,319 B2.

06 - Consolidating Data on System z v2.0.ppt


9

Result: DB2 Scales Out, Oracle RAC is Limited

- DB2 for z/OS provides near-linear scalability with relatively little overhead as nodes are added
- With Oracle RAC, overhead increases rapidly as additional nodes are added and performance degrades after only 4 to 6 nodes

Sources: "Scale-up versus scale-out using Oracle 10g with HP StorageWorks", Hewlett-Packard, 2005

"Enterprise Data Base Clustering Solutions" ITG, October 2003




06 - Consolidating Data on System z v2.0.ppt

10

HP Benchmark Describes Oracle RAC Scalability

Oracle RAC has efficiency issues scaling beyond 4 to 6 nodes.


Source: Hewlett-Packard, 2005
Scale-up versus scale-out using Oracle 10g with HP StorageWorks



06 - Consolidating Data on System z v2.0.ppt

11

Oracle RAC Overhead Wastes Processing Power in Each Node



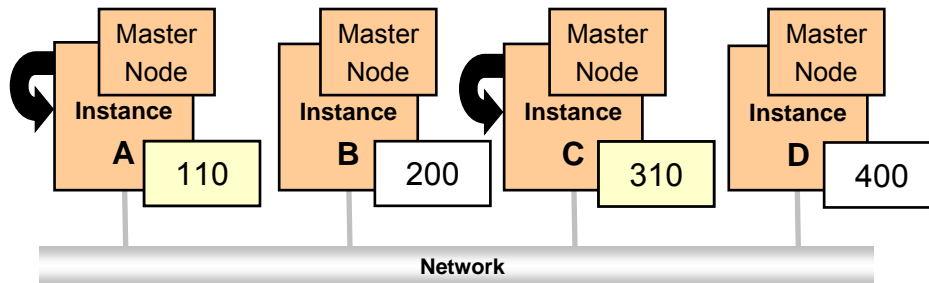
Oracle RAC source:
"Scale-up versus
scale-out using Oracle
10g with HP
StorageWorks",
Hewlett-Packard, 2005

DB2 for z/OS source:
"Enterprise Data Base
Clustering Solutions"
ITG, October 2003

06 - Consolidating Data on System z v2.0.ppt

12

Partitioning is Necessary to Effectively Use Oracle RAC Above 4 to 6 Nodes



- Partitioning associates the data, buffer, and lock manager with a particular server node
- All work requests affecting the associated data partition must be executed by the owning node
- Partitioning requires a redesign and reimplement of the database and the application (big job on Oracle!)
- Partitioning may not always be possible depending on the workload

06 - Consolidating Data on System z v2.0.ppt

13

Compare Scalability Achievements: Winter Corporation's "2005 Top Ten" Awards


- "The highest performing transaction processing system, a [DB2] z/OS implementation, executed **over one billion** SQL statements in an hour (United Parcel Service)."

- The study lists the largest known peak workload on **Oracle RAC** to be **8.6 million** SQL statements per hour

- "The largest transaction processing database, **23 TB**, was hosted on DB2 for z/OS" (UK Government Land Registry)

- The study lists the largest transaction processing database on **Oracle RAC** to be **9.6 TB**

Leading OLTP System,
DB2 for z/OS, Processed
Over One Billion SQL
Statements per Hour



http://www.wintercorp.com/VLDB/2005_TopTen_Survey/TopTenWinners_2005.asp
http://www.wintercorp.com/WhitePapers/WC_TopTenWP.pdf

06 - Consolidating Data on System z v2.0.ppt

15


Fractional Availability Improvements are Important

- Fractional Improvements Result in Millions in Savings

- Financial Impact of Downtime Per Hour for financial industry is \$1,495K

- Financial Services Company Example:

- ▶ \$300B assets, 2500+ branches, 15M customers
- ▶ Retail banking, loans, mortgages, wealth management, credit cards
- ▶ CRM System – branches, financial advisors, call centers, internet
- ▶ Number of users – 20,000+



	Unix/Oracle	zSeries/DB2
Availability %	99.825%	99.975%
Annual outage	15h 20m	2h 11m
Cost of Downtime	\$23M	\$3M

\$20 Million dollar difference!

Sources: ITG Value Proposition for Siebel Enterprise Applications, Business case for IBM eServer zSeries, 2004 & Robert Frances Group, 2005

06 - Consolidating Data on System z v2.0.ppt

16

Data Security and Compliance: DB2 for z/OS Has a Proven Track Record

DB2 for z/OS Security


- Less than 10 security related patches in the last 10 years
- Proven RACF and Multi Level Security
 - ▶ End to end security including applications disks, printers and network
- DB2 Test Database Generator
 - ▶ Ensures anonymous access to data necessary for testing
- DB2 Archive Expert
 - ▶ Allows customers to easily archive and access data
- DB2 Audit Management Expert
 - ▶ Supports compliance requirements
- End-to-end encryption via hardware assist

Oracle's Security Exposures

- **C/NET - 01/17/2007**
"Oracle plugs 51 security flaws" including 26 for database
- **ComputerWorld - 10/17/2006**
"Oracle releases 101 patches in quarterly update" including 63 for database
- **NGS Research - 11/21/2006**
"The conclusion is clear – if security robustness and a high degree of assurance are concerns when looking to purchase database server software – given these results one should not be looking at Oracle as a serious contender."

XML Usage is Increasing

- SOA
 - ▶ Web Services messages are XML
 - Web Services are the foundation of service oriented architecture*
- Business-to-Business Integration
 - ▶ Platform-independent transport mechanism.
 - Business transactions may be defined in XML*
- Forms and Document Processing
 - ▶ Government and legal industry require digital signature
 - ▶ Documents often contain sub-documents
 - Documents are increasing created and exchanged in XML*



XML – The Change Is Fundamental

- Relational is a data model
 - Relations (tables)
 - Attributes (columns)
 - Set based w/ some sequences
 - Strict schema

POID	CustomerID	ItemID
12	1	2
102	3	4
102	3	5

id	LastName	FirstName	Street	City	State	Zip
1	Pirahesh	Hamid	1 Harry Rd	San Jose	CA	95141
3	Seinger	Pat	555 Bailey Ave	San Jose	CA	95141


ItemID	Name
2	#6 wire nut
5	Small Walrus
4	Apollo moon rocket

- XML is a data model
 - Hierarchical tree structure
 - Nodes (elements, attributes, comments, etc.)
 - Relationships between nodes
 - Sequence based w/ some sets
 - Flexible schema

```
<? xml version="1.0" ?>
<paymentRequest id="12345" >
  <payment>
    <payment_name>Loan</payment_name>
  </payment>
  <payee_info>
    <payee_amt>$97.85</payee_amt>
  </payee_info>
  <remitter>
    <name>John Smith Co</name>
    <address>
      <street>1234 W. Main</street>
      <city>Yonkers</city>
      <state>NY</state>
      <zip>11111</zip>
    </address>
  </remitter>
</paymentRequest>
```

DB2 9 Supports Native Storage and Retrieval of XML Data as Well as Relational Data

Use SQL/XML to access both types of data

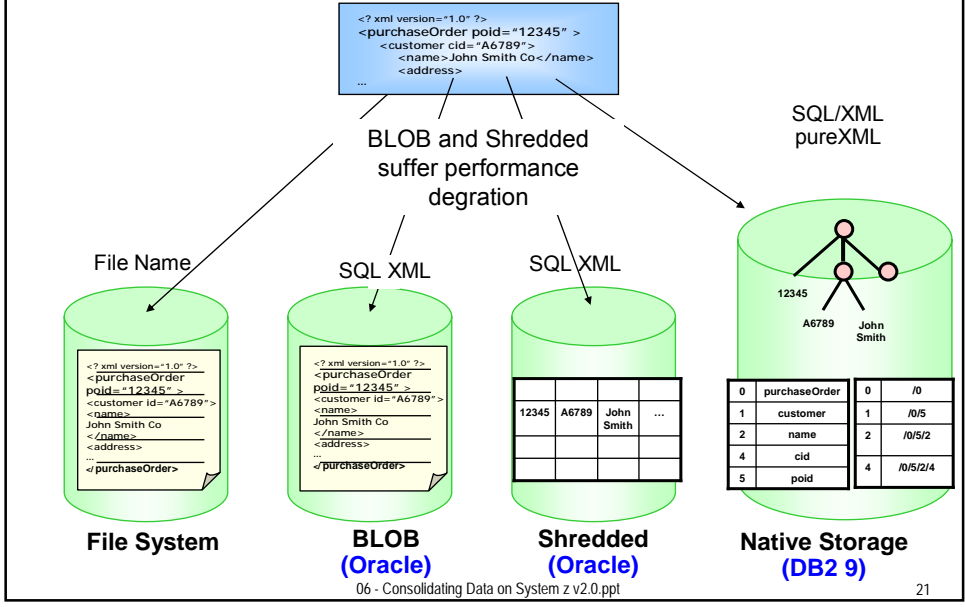


XML integrated in all facets of DB2!




New XML applications benefit from:

- Ability to seamlessly leverage relational investment
- Proven Infrastructure that provides enterprise-class capabilities

DB2 9 Has Better Support for XML



DB2 XML Performance Beats Oracle




DB2 9 performs 1.5 to 3.6 times faster than Oracle on XML operations

Note: Performance results are for DB2 UDB, DB2 9 for z/OS is expected to have similar results

DB2 for z/OS Is Optimized to be *the* Data Server for SAP

- Partnership with SAP
 - ▶ 13 years of DB2 partnership with SAP
 - ▶ Joint development team with SAP to integrate SAP and DB2 solutions
 - ▶ DB2 for z/OS V8: more than 50 features by SAP
 - V8: more than 50 features requested by SAP
 - ▶ DB2 for z/OS V9: approximately 40 features requested by SAP
 - ▶ No unique features in SAP exploit Oracle



06 - Consolidating Data on System z v2.0.ppt

23

DB2 for z/OS Is Optimized to be *the* Data Server for SAP

Examples:

- Ease-of-Use
 - ▶ Easy to clone DB2 instances, such as test environment
 - ▶ Customized SAP 'Tuner'
- Less DBA skills and activities required
 - ▶ Large Object Management, SAP uses large objects a lot
 - ▶ DB2 Recovery Expert for automatic recovery and backup
 - ▶ Real-time Statistics Utility provides automatic scheduling information
 - ▶ BACKUP and RESTORE system enhancements
- SAP-specific enhancements to DB2 Query Optimizer
 - ▶ Enhancements for SAP Business Intelligence query performance
 - ▶ Enhancements for SAP OLTP products
- High Performance
 - ▶ SAP Business Warehouse performance gains through Dynamic Index ANDing

06 - Consolidating Data on System z v2.0.ppt

24

Get More Business Results Out of Your Data

Our branch offices have separate databases.
Each branch is analyzing customers and sales on their own.



Service Oriented Finance
Marketing

Looking at data in isolation can miss larger trends and opportunities



IBM

06 - Consolidating Data on System z v2.0.ppt

25

Service Oriented Finance Needs a Data Warehouse to Make Optimal Business Decisions


- Each branch is responsible for its own marketing campaign
- Corporate marketing gets reports from each of the branches based on local results
- Corporate marketing needs to spot trends to know what campaigns are most effective region-wide
- A corporate data warehouse would give marketing the data to easily do comparisons between the branches and promote best practices

06 - Consolidating Data on System z v2.0.ppt

26

DEMO: Branch Analysis With Excel

- Excel being used against local SQL Server database
 - ▶ Local branch has good growth
 - ▶ Local campaign had positive impact on customer growth and number of products per customer



06 - Consolidating Data on System z v2.0.ppt

27

Use DB2 for z/OS to Build Your Corporate Data Warehouse

Performance features for Data Warehouse solutions

- Parallel Queries
 - ▶ Exploit multiple processors if available
- Materialized query table
 - ▶ Save and reuse previous partial query results
- Star Schema Join Enhancement
 - ▶ Performance enhancements for typical data warehouse accesses


06 - Consolidating Data on System z v2.0.ppt

28

DB2 for z/OS Automatically Parallelizes Queries

- Three techniques are supported
 - ▶ Utilize multiple processors on a single system
 - ▶ Utilize multiple processors across a sysplex
 - ▶ Execute parallel I/O requests
- These techniques are combined as configurations permit


Example: Parallelization of Queries



Use WebSphere DataStage to Load Your Data Warehouse From the Branches

■ Data Transformation and Movement

- ▶ Extract data from source
- ▶ Transform data
- ▶ Load data into data warehouse
- ▶ DataStage Designer tool creates DataStage ETL jobs



06 - Consolidating Data on System z v2.0.ppt

31

Data Stage Transforms Data on the Fly

Different field names
 Different field order
 Add Branch Identifier
 Different currency format



PROD ID	CUST ID	BRANCH ID	QTY	AMT	SALEDATE
000 101	100	01	01	10,000.00	2007-02-28
000 121	100	01	03	500.50	2007-02-28
000 101	101	01	01	20,000.00	2007-03-01

Data Warehouse



Transform


PRODUCT	QTY	CUSTNO	AMOUNT	DATE

Branch Data

06 - Consolidating Data on System z v2.0.ppt

33

DEMO: Run DataStage ETL Job



06 - Consolidating Data on System z v2.0.ppt

34

DEMO: Use Same Excel Tool to Do Corporate Level Analysis

- Excel accesses data from DB2 for z/OS seamlessly
- Campaign to offer mortgage point deduction most effective... Park City branch shows this
- The Inlet branch has the largest number of new accounts for all of the branches, however that does not necessarily equate to the most profit



06 - Consolidating Data on System z v2.0.ppt

35

DB2 for z/OS Data Warehouse Capabilities

- DB2 functional and performance enhancements
 - ▶ Parallel queries, Materialized Query Table, Star Join Enhancement
- Extract, Transform and Load
 - ▶ WebSphere QualityStage, WebSphere DataStage, WebSphere Information Integrator Editions
- Analyze/Report – connect to most popular data analysis tools
 - ▶ Excel, Alphablox, QMF, Brio, Hyperian, Business Objects, SAS, IBI
- Performance Monitoring
 - ▶ IBM Tivoli Omegamon XE for DB2 Performance Expert on z/OS
- Comprehensive Management Tools for Security and Compliance Including
 - ▶ DB2 Data Archive Expert, DB2 Test Database Generator, DB2 Audit Manager Expert, IBM Encryption for DB2 and IMS Databases

06 - Consolidating Data on System z v2.0.ppt

36

DB2 for z/OS Has the Lowest TCO

Capability is important, but cost is a big concern for us



On Demand Bank
CIO

DB2 for z/OS costs less than Oracle RAC



IBM

06 - Consolidating Data on System z v2.0.ppt

37

DB2 – Better Compression Ratio Than Oracle

- TPC-H is a well known data warehouse benchmark
 - ▶ Each vendor uses the same tables and same data
 - ▶ Oracle published their compression rates for TPC-H tables at the VLDB conference in 2003
 - ▶ IBM ran the same tests on the same tables
- Test results

Table	Compression Ratio	
	Oracle	DB2
LINEITEM	38%	58% (1.5x better)
ORDERS	18%	60% (3x better)
Entire Database	29%	59% (2x better)

06 - Consolidating Data on System z v2.0.ppt

38

Storage Costs: DB2 Provides More Storage Savings than Oracle

- DB2 for z/OS lowers TCO by reducing storage use
 - ▶ Typical tables compress in the range of 40% - 80%
 - ▶ Overall database storage savings average around 30% - 45%
 - ▶ For certain large database installations 50–80% savings were reported
- Example storage savings with DB2 vs. Oracle

	Oracle	DB2 for z/OS*
Reference price for 3TB of disk	HP XP12000 Storage	IBM System Storage DS8100
Overall database compression ratio (using TPC-H benchmark results.)	29%	59%
For 100 TB uncompressed data storage needed	75 TB of HP Storage	42 TB of IBM Storage
Cost of storage	\$3.1M + \$0.225M** = \$3.34M	\$1.45M
With compression, storage for DB2 costs <u>56% less</u> than for Oracle		

* DB2 for z/OS achieves similar compression ratios to those of DB2 for LUW
 ** HP Storage Software charge

06 - Consolidating Data on System z v2.0.ppt

39

New zIIP Processor Dramatically Lowers Cost

What Workloads Can Be Run on a zIIP?

- How much DB2 workload can typically be run on a zIIP?
 - ▶ Queries received via DRDA Remote Access Protocol (**Database Server scenarios**)
 - Up to 40%
 - ▶ Parallel queries (**Data Warehouse scenario**)
 - Up to 80%
 - ▶ Some of **index maintenance** utilities
- Offloads to zIIP specialty processor reduce DB2 load and charges on general purpose processors
 - ▶ For sub capacity pricing, the offload must occur at a time that will reduce billable rolling average
- IBM has tools to help customers estimate their off load potential

06 - Consolidating Data on System z v2.0.ppt

40

Example: Consolidate SAP Database Server on Mainframe

Existing Mainframe



Existing processors:
9 general purpose
4000 MIPS
running existing DB2
workloads

Add 1 LPAR for New SAP Data Server w 42 TB Storage



966 MIPS
additional
workload

3 year
cost of
acquisition
\$3.40M

Add two processors:
1 zIIP's
386 MIPS (40%)
1 General purpose
580 MIPS (60%)

Or add HP Integrity rx8640 Server w 75 TB storage



84,042 *
RPE's

3 year
cost of
acquisition
\$5.45M

* Production RPE's required = 966 x 87 = 84,042

06 - Consolidating Data on System z v2.0.ppt

41

SAP Data Server Incremental Cost Breakdown


Mainframe Hardware				Mainframe Software			
OTC		ANNUAL		OTC		ANNUAL	
1 General Processor	\$1,452,500	Processor Maintenance	\$80,868	Utilities	\$0	DB2 MLC Incremental	\$33,840
1 zIIP Processor	\$125,000					zOS MLC (zNALC)	\$34,944
IBM Storage (42TB)	\$1,449,801	Storage Maintenance	0				
TOTAL	\$3,027,301	TOTAL	\$80,868	TOTAL	\$0	TOTAL	\$68,784

Distributed Hardware				Distributed Software			
OTC		ANNUAL		OTC		ANNUAL	
HP Processors	\$603,939	Processor Maintenance	\$123,139	Oracle EE	\$640,000	Oracle S&S	\$297,440
HP storage (75TB)	\$3,107,469	Storage Maintenance	\$30,951	HP-UX	\$126,048	HP-UX S&S	\$35,819
						Prepaid \$107,456 in year 1 for 3 years	
TOTAL	\$3,711,408	TOTAL	\$154,090	TOTAL	\$991,148	TOTAL	\$333,259


06 - Consolidating Data on System z v2.0.ppt

42

Oracle on Distributed Cost Comparison With MF (Incremental DB2 Workload For SAP)



Total cost = \$3,395,389



Total cost = \$5,448,404


1.6 times more expensive

06 - Consolidating Data on System z v2.0.ppt

43


Example: Consolidate New Data Warehouse Application on Mainframe

Existing Mainframe



Existing processors:
2 general purpose

Existing Disaster Recovery Site



Existing processors:
Pay for one general purpose processor for hot disaster switch over and one "dark" DR processor at \$30K

Add 1 LPAR for New Data Warehouse w 42 TB Storage

Prod

1954 MIPS additional workload

Add four processors:
3 zIIP's
1464 MIPS (75%)
1 General purpose
489 MIPS (25%)

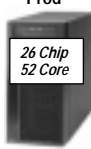
And Add Disaster Recovery

Prod

Pay for four additional DR processors at \$30K each

3 year cost of acquisition \$4.90M

Or add Superdome 9000 Server w 75 TB storage



26 Chip
52 Core

169,998 *
RPE's

Disaster Recovery typically not considered

3 year cost of acquisition \$8.24M

* Production RPE's required = 1954 x 87 = 169,998

06 - Consolidating Data on System z v2.0.ppt 44

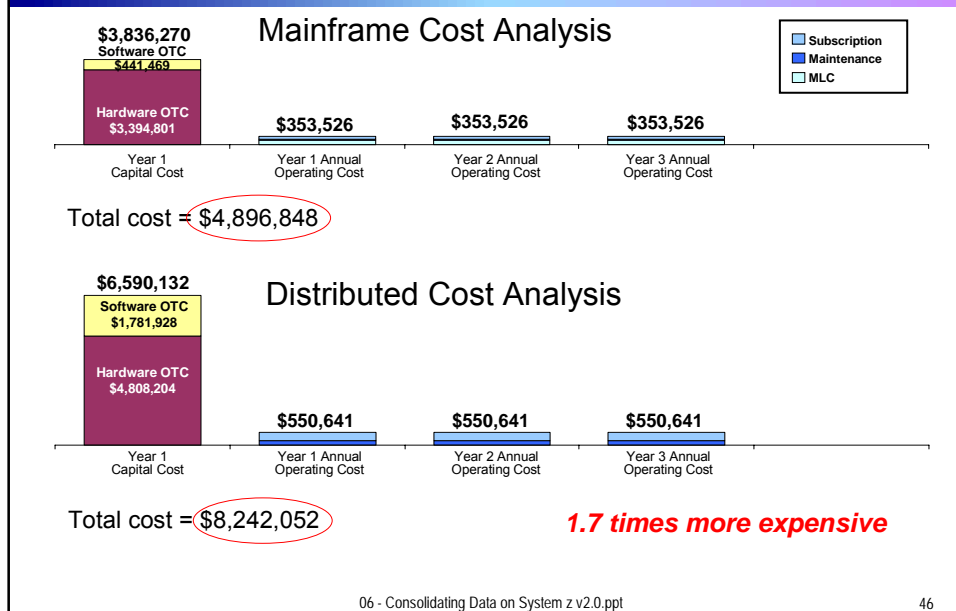
Data Warehouse Incremental Cost Breakdown

Mainframe Hardware				Mainframe Software			
OTC		ANNUAL		OTC		ANNUAL	
Z Processors	\$1,825,000	Processor Maintenance	\$123,540	Utilities	\$441,469	Utilities S&S	\$44,454
4 DR Processors	\$120,000					DB2 MLC	\$72,240
IBM Storage (42TB)	\$1,449,801	Storage Maintenance	0			QMF MLC	\$34,716
						zOS MLC	\$78,576
TOTAL	\$3,394,801	TOTAL	\$123,540	TOTAL	\$441,469	SubTotal MLC	\$185,532
						TOTAL	\$229,986

Distributed Hardware				Distributed Software			
OTC		ANNUAL		OTC		ANNUAL	
HP Processors	\$1,700,735	Processor Maintenance	\$164,044	Oracle EE & Utilities	\$1,352,000	Oracle S&S	\$297,440
HP storage (75TB)	\$3,107,469	Storage Maintenance	\$30,951	Unix	\$204,828	Unix S&S	\$58,205
				HP Storage SW	\$225,100		
TOTAL	\$4,808,204	TOTAL	\$194,995	TOTAL	\$1,781,928	TOTAL	\$355,645

06 - Consolidating Data on System z v2.0.ppt 45

zIIP Processors Lower the Cost of Acquisition



The World Relies on DB2 for System z – You Can Too

- Over 10,000 Licenses World Wide
- Over 8 Million Clients
- Over 3,000 TBs of Production Data
- Over 700 ISV Applications and Packages
- Owns 96% of Relational System z Market Place

