

Bringing Big Data to the Enterprise - Gaining new insight with Big Data capabilities

Dan Wardman

***Vice President, Information Management,
Mainframe Software, IBM Software Group***



An abstract graphic composed of various colored 3D cubes and rectangular blocks in shades of red, pink, orange, and black, arranged in a cluster.

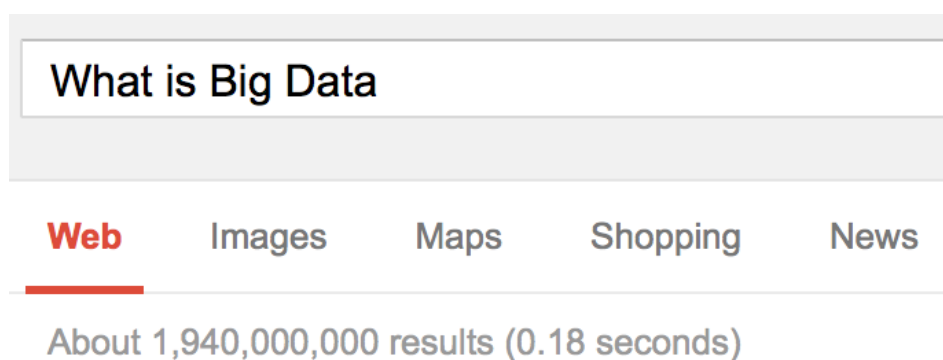
Agenda

- **Big Data**
- **zEnterprise & Big Data**
- **zEnterprise & Big Data Analytics**
- **Client Case Studies**
- **Analytic Solutions on zEnterprise**



What is Big data?

- Google can give you nearly 2 Billion options
- Vendors have even more definitions



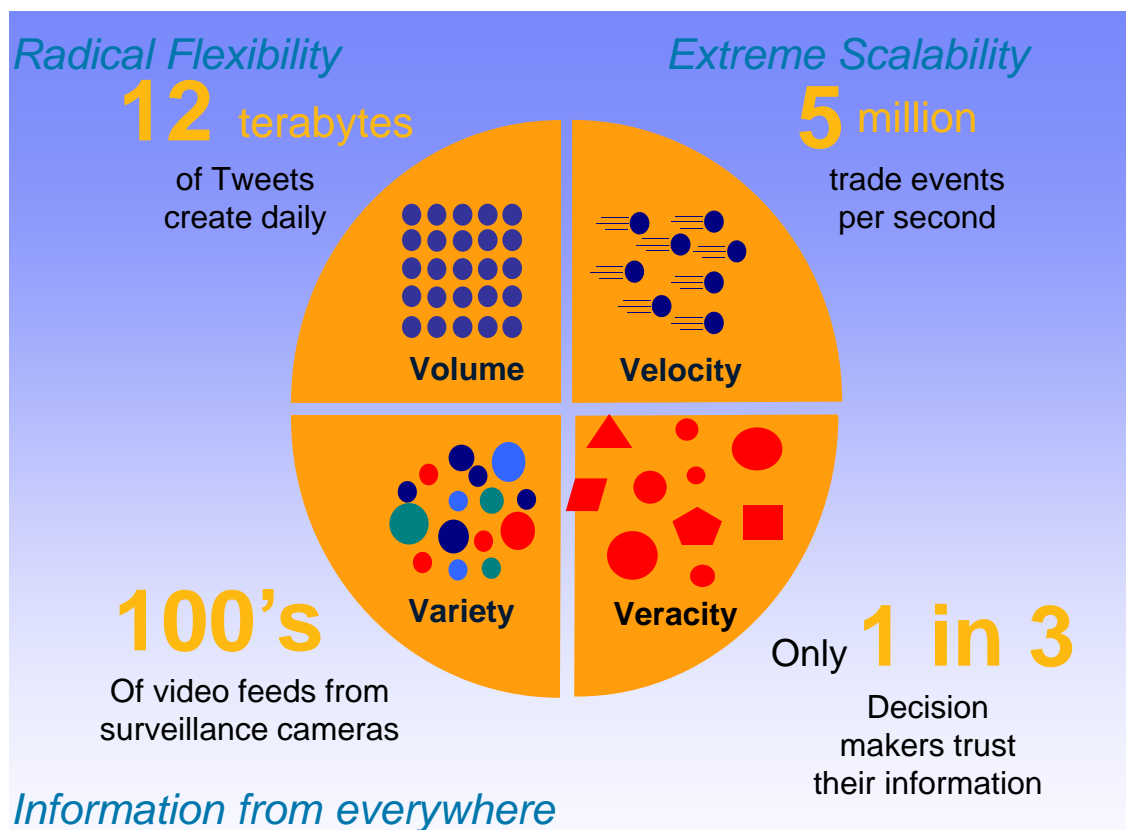
Here is how Gartner defines Big Data

- Big data is high-volume, high-velocity and high-variety information assets that demand **cost-effective, innovative information processing for enhanced insight and decision making.**



Gartner®

We've Moved into a New Era of Computing - V⁴



“We have for the first time an economy based on a key resource [Information] that is not only renewable, but self-generating.

Running out of it is not a problem, *but drowning in it is.*”

– John Naisbitt

Majority of today's analytics based on relational / "Structured" Data

- Analytics and decision engines reside where the DWH / transaction data is
- "Noise" (veracity) surrounds the core business data
 - Social Media, emails, docs, telemetry, voice, video, content
- What data are you prepared to **TRUST?**
- Where do you put your trusted Data?



"Circle of trust"

Demand for differently structured data to be seamlessly integrated, to augment analytics / decisions

- Analytics and decision engines reside where the DWH / transaction data is
- “Noise” (veracity) surrounds the core business data
 - Social Media, emails, docs, telemetry, voice, video, content
- Expanding our insights – getting closer to the “truth”
 - Lower risk and cost
 - Increased profitability



**“Circle of trust”
widens**


Enterprise Integration and Governance... the key to success of incorporating Big Data

■ Information Integration

– Insights from big data must be incorporated into the warehouse and analytics/ decision engines

■ Information Governance

– Companies need to govern what comes in, and the insights that come out






The Big Data Starting Point

Where are organizations getting the most return on Big Data projects?

“What types of data/records are you planning to analyze using big data technologies?”



Most big data use cases hype its application for analysis of new, raw data from social media, sensors, and web traffic, but we found that firms are being very practical, with early adopters using it to operate on enterprise data they already have.

Source: 2012 IBM Global Big Data Online Survey



...and our webcast survey said...

▪ Have you already implemented or are you planning to implement any Big Data based initiatives within the next 6 months?	Yes	31%
	No	69%
▪ How would you rate the value of being able to integrate insights from social media, telemetry, unstructured data into your analytics and decision making processes?	High	50%
	Medium	36%
	Low	14%
▪ Do you see the IBM System z platform as pivotal to the success of Big Data initiatives?	Yes	90%
	No	10%

Five key findings from the study about big data:

1. Customer analytics are driving big data initiatives
2. Big data is dependent upon a scalable and extensible information foundation
3. Initial big data efforts are focused on gaining insights from internal data
4. Big data requires strong analytics capabilities
5. Adoption of big data is focused upon delivering measureable business value, which happens in four stages:
 - Educate: focusing on business as usual with casual understanding of big data;
 - Explore: developing strategy and roadmap based on business needs and challenges;
 - Engage: creating pilots to validate value and requirements; and
 - Execute: deploying two or more big data technologies and continuing to innovate



Big Data Use Cases



Big Data Exploration

Find, visualize, understand all big data to improve decision making



Enhanced 360° View of the Customer

Extend existing customer views (MDM, CRM, etc) by incorporating additional internal and external information sources



Security/Intelligence Extension

Lower risk, detect fraud and monitor cyber security in real-time



Operations Analysis

Analyze a variety of machine data for improved business results



Data Warehouse Augmentation


Integrate big data and data warehouse capabilities to increase operational efficiency



The role of zEnterprise in Big Data Analytics



- A large percent of the data that is accessed for analytics originates/resides on IBM zEnterprise
 - 2/3 of business transactions for U.S. retail banks
 - 80% of world's corporate data
- Businesses that run on zEnterprise
 - 66 of the top 66 worldwide banks
 - 24 of the top 25 U.S. retailers
 - 10 of the top 10 global life/health insurance providers
- 1,300+ ISVs run zEnterprise today, more than 275 of these selling over 800 applications on Linux
- The downtime of an application running on System z equates to approximately 5 minutes per year
- The System z mainframe can run over a thousand virtual Linux images on a single frame the size of a refrigerator



Traditional Approach to workload optimized Systems

Operational Applications

Transaction Processing



Shared Everything DB

High volume business transactions and batch reporting running concurrently

Analytic Applications

Data warehousing



Shared Nothing DB

Low volume complex queries



The zEnterprise Hybrid Solution

Mixed Workloads for Next Generation Business Analytics

**Mixed Workload
Applications**


Enhancing Big Data Analytics with IMS and DB2 for z/OS

<ul style="list-style-type: none"> ▪ Much of the world's operational data resides on z/OS 	<ul style="list-style-type: none"> ▪ Unstructured data sources are growing fast
---	---

▪ Two significant needs:

1. Merge this data with trusted OLTP data from zEnterprise data sources
2. Integrate this data so that insights from Big Data sources can drive business actions

- IMS & DB2 are providing the connectors & the DB capability to allow BigInsights to easily & efficiently access each data source
- DB2 is providing the connectors & the DB capability to allow DB2 apps to easily and efficiently access hadoop data sources





IBM PureData System for Hadoop

Accelerate Hadoop analytics with appliance simplicity

Accelerate Big Data projects with built-in expertise

- Explore new ways to use all data
- Unlock new insights from unstructured data
- Establish a cost efficient on-line data archive

Simplify with integrated system management

- InfoSphere BigInsights software
- Compute and Storage hardware

Ensure production grade security and governance

Easily integrate with other systems
in the IBM big data platform



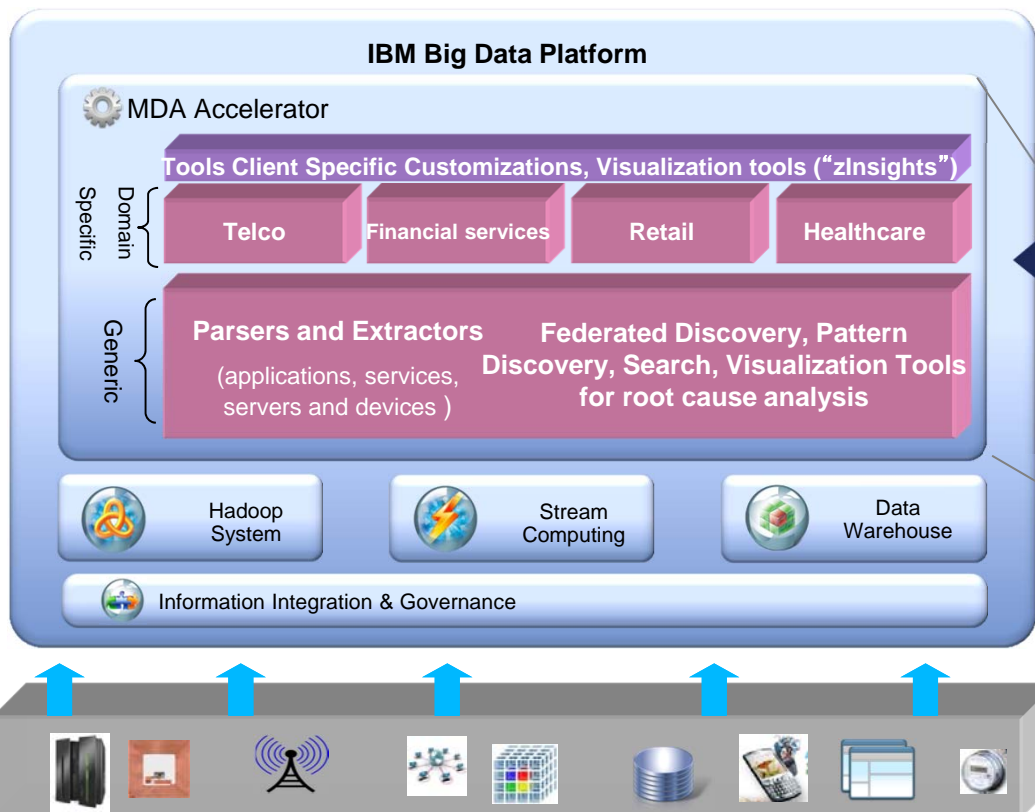



DB2 11 Support for Big Data

- **Goal: integrate DB2 for z/OS with IBM Hadoop based BigInsights Bigdata platform**
 - Enabling traditional applications on DB2 z/OS to access Big Data analytics.
- **Analytic jobs can be specified using JSON Query Language (Jaql)**
 - Submitted to BigInsights
 - Results stored in Hadoop Distributed File System (HDFS).
- **A table UDF (HDFS_READ) reads the Bigdata analytic result from HDFS, for subsequent use in an SQL query.**
- **Must have a variable shape of HDFS_READ output table**
 - DB2 11 supports generic table UDF, enabling this function



Machine Data Accelerator




IT use cases:

- Server, performance, troubleshooting

Business use cases:

- Click stream and transaction analysis
- Optimiz- ... ce planning





Big Data Innovation with IBM zEnterprise



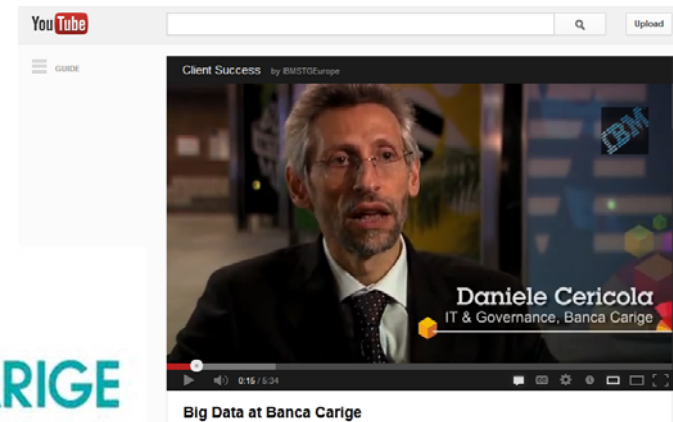
Operations Analysis

Analyze a variety of machine data for improved business results



Banco do Brasil purchases the **largest ever** DB2 Analytics Accelerator solution to drive customer insight from operational data. The 120-way system can hold 1.28 Petabytes of data. Queries that previously took **11 hours to run now complete in 26 seconds**, over 1500 times faster!

Banca Carige chooses System z to provide real time analytics as part of their Big Data client solution





Big Data Innovation with IBM zEnterprise

Data Warehouse Augmentation

Integrate big data and data warehouse capabilities to increase operational efficiency

With healthcare reform posed to add 30 million new members,



Aetna looks to expand membership by as much as 75% using System z which can now provide insight 1700 times faster without impacting existing applications & infrastructure

Implemented a clinical dimensional data warehouse with billions of patient diagnostic records with superior scalability and 24x7 availability, surpassing industry privacy requirements





Implementing a Mission Critical Big Data Application

Big Data Exploration

Find, visualize, understand all big data to improve decision making



JOHN DEERE

GPS & sensor information volumes exceeded the capabilities of the existing system. It was redesigned as an enterprise mission critical application using DB2 for z/OS and System z data sharing to now provide the availability and scalability to meet the current and future requirements for this solution.



zEnterprise Analytics solutions deliver...



Run complex queries up to **2000x Faster**

Meet SLAs & score **3000-5000+ transactions in Real-time**

95% savings in host disk space for historical data

30%- 45% performance improvement with zEC12 over z196

80% less capacity for Data Warehousing

Ave. **87% savings in CPU, 96% savings in servers** over 5 yrs. for BI deployments

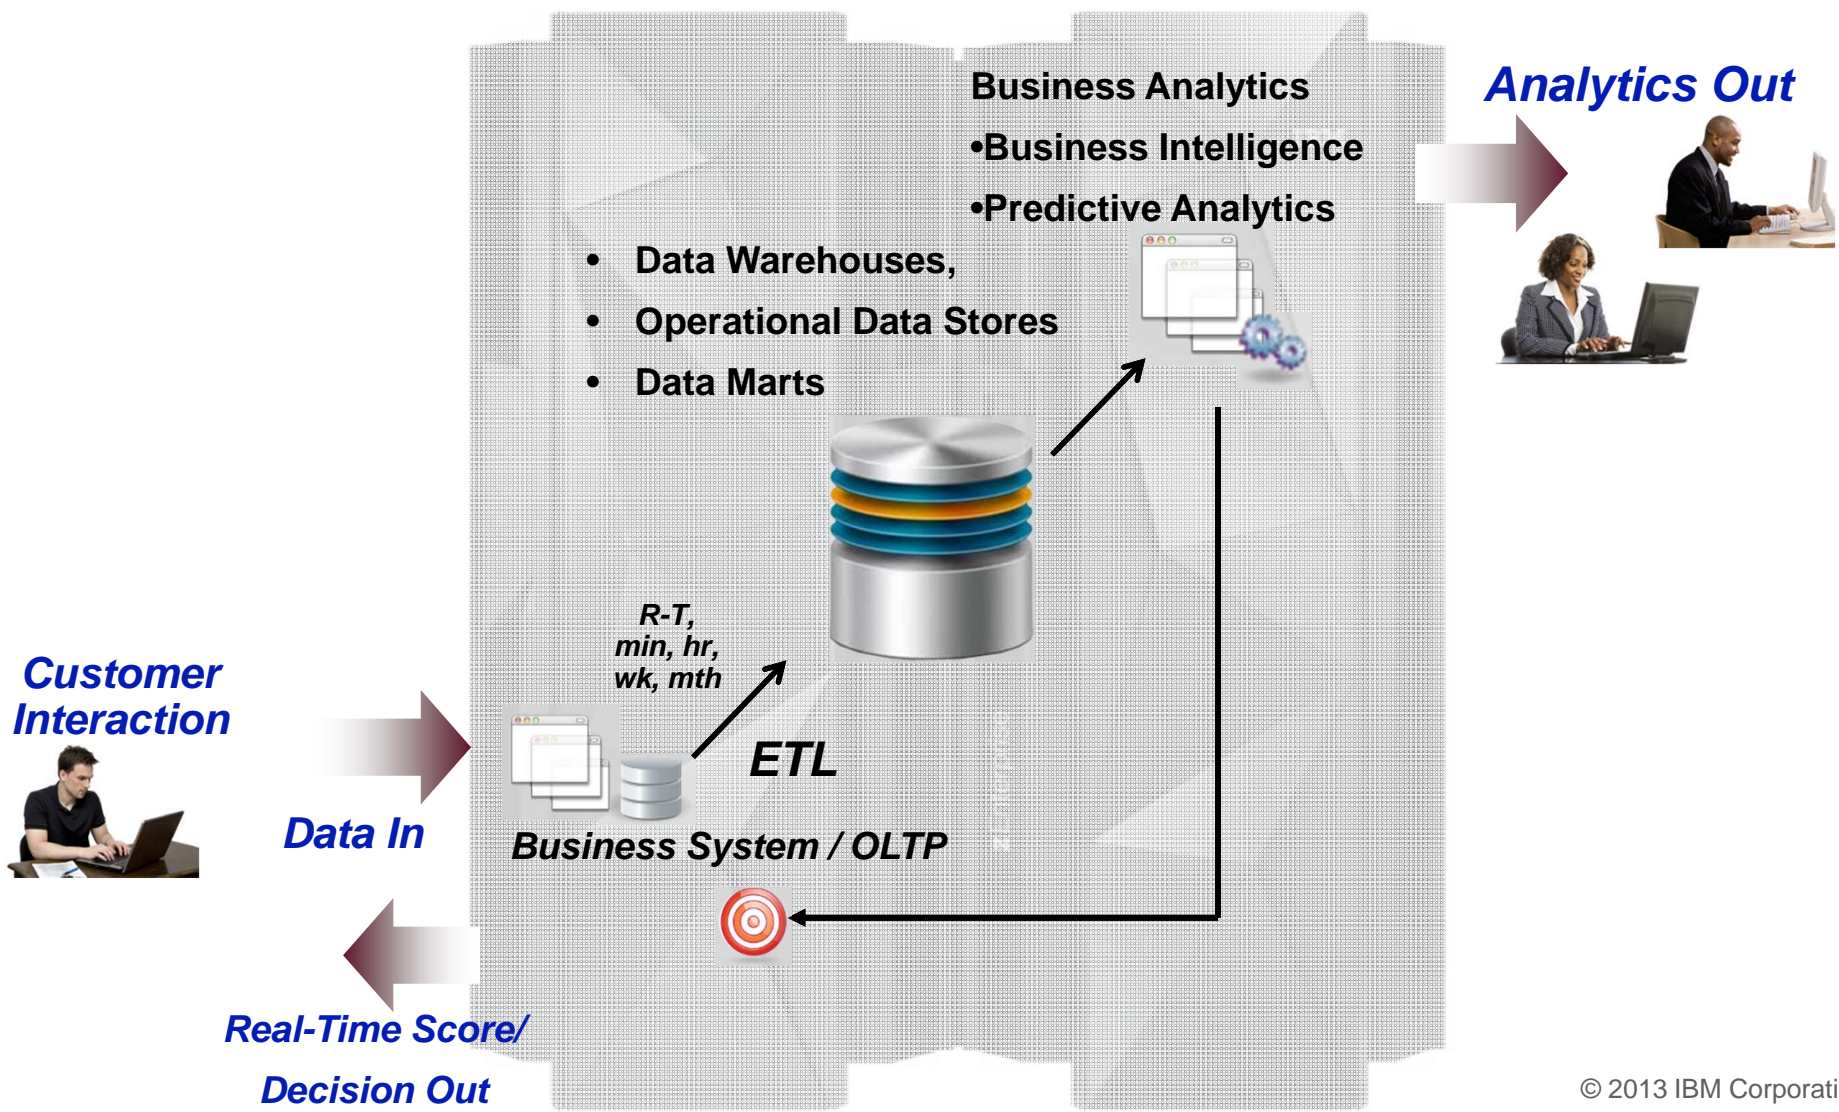
BI system **admin savings alone, pays for the HW investment** in 5 years

BI servers that run at **90%+ capacity** without impacting SLA



IBM zEnterprise

An end to end rock solid foundation for analytics





DB2 11 Major Themes

Building on the Success of DB2 10

▪ Performance Improvements

- Improving efficiency, reducing costs, no application changes
- 0-5% for OLTP, 5-15% for update intensive batch
- 5-20% for query workloads
- Exploitation of new zEC12 hardware features

▪ Continuous Availability Features

- Improved autonomies which reduces costs and improves availability
- Making online changes without affecting applications
- Online REORG improvements, less disruption

▪ Enhanced business analytics

- Faster, more efficient performance for query workloads
- SQL improvements and IDAA enhancements
- Temporal and SQLPL enhancements
- Transparent archiving

▪ Simpler, faster DB2 version upgrades

- No application changes required for DB2 upgrade
- Access path stability improvements
- Product stability: support pre GA customer production





DB2 Analytics Accelerator for z/OS

Blending zEnterprise and Netezza technologies

A high performance analytics accelerator appliance for IBM zEnterprise, delivering dramatically faster complex business analysis transparently to all users.



Fast

Complex queries run up to 2000x faster while retaining single record lookup speed

Cost Saving

Eliminate costly query tuning while offloading complex query processing

Appliance

No applications to change, just plug it in, load the data, and gain the value

What's new in DB2 Analytics Accelerator V3

- ✓ *Lowering the cost of historical data*
- ✓ *Better decisions through lower latency of data*
- ✓ *Dramatic improvement in scale and growth opportunities*
- ✓ *Lowering the cost of analytic computing*



- **High Performance Storage Saver**
 - Significantly reduces the cost for storage resources
 - Option to store data only once: in the accelerator
- **Incremental Update**
 - Data changes are propagated to the accelerator as they happen
 - Uses change data capture technology
 - Extends the accelerator use to reporting on operational data
- **New optimization**
 - Tables or partitions refresh much faster and less resources intensive
 - Optimized unloading data from DB2
- **High Capacity**
 - Capacity has been extended to 1.28 PB for a single Accelerator
- **New functions**
 - More queries eligible for acceleration



Introducing the new N2001

The fastest performance of Netezza technology to date!

*Accelerate Performance
of Analytic Queries*

- **3X faster performance¹**
for Big Data analytics
- **128 GB/sec effective scan rate per rack²**
to tackle Big Data faster

*Increase Efficiency
of your Data Center*

- **50% greater data capacity per rack³**
helps optimize data center efficiency
- **More capacity and less power per rack**
than both Oracle and Teradata

*Simplicity and
Ease of Administration*

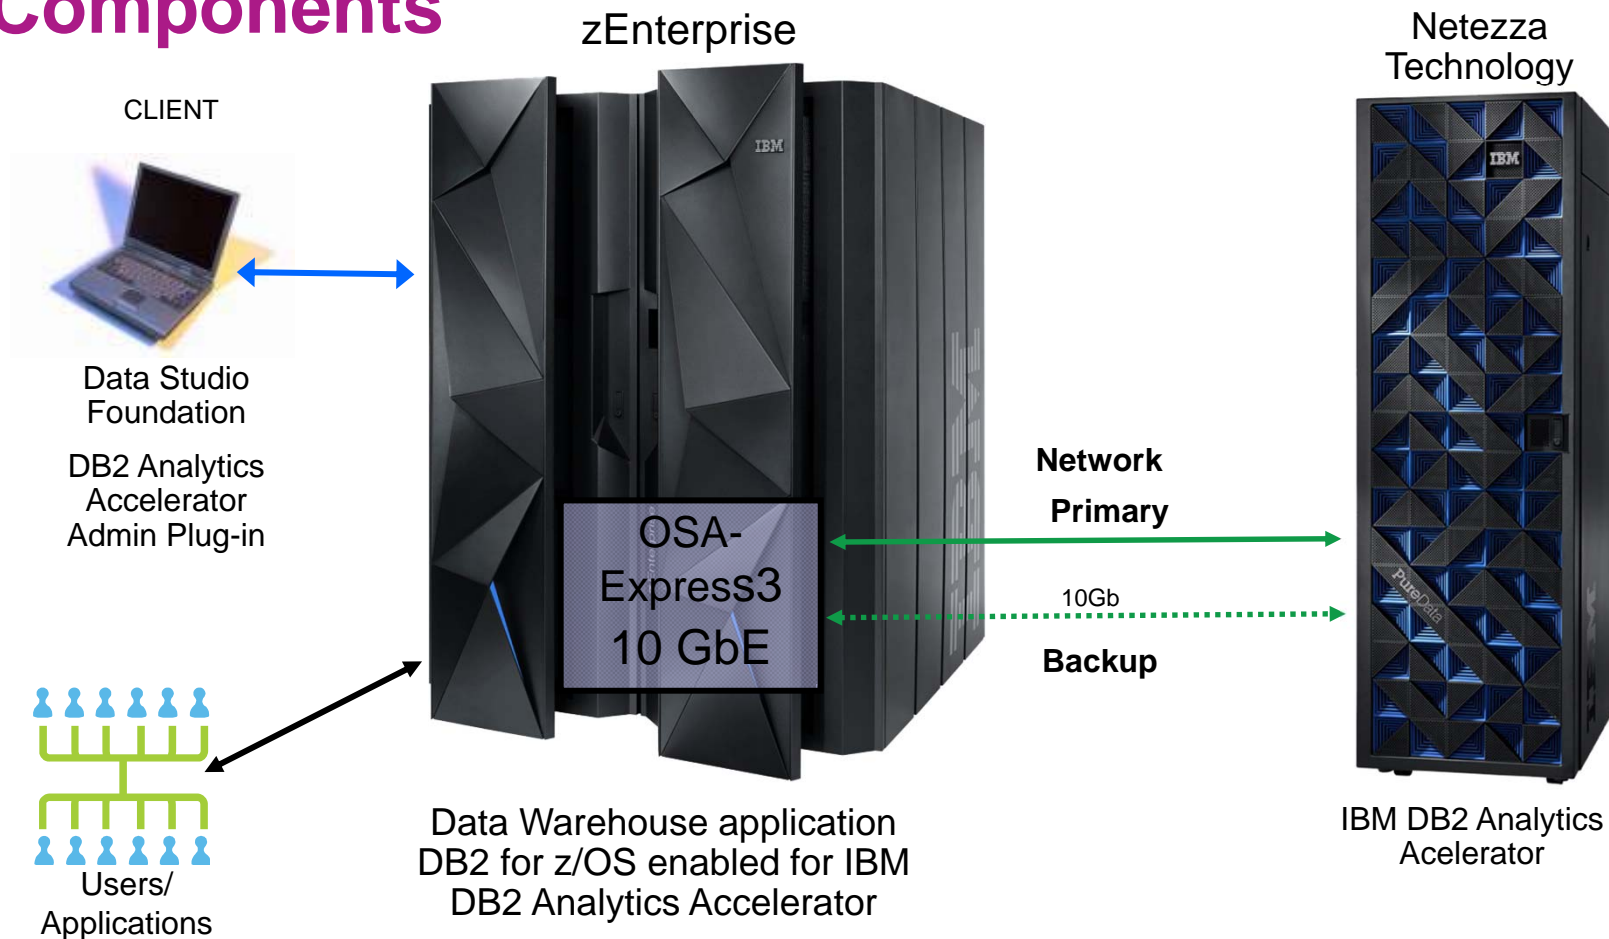
- **Improved system management and resilience**
*to spend less time managing and more time
delivering value*

¹ Based on a comparison of the IBM PureData System for Analytics N2001 to the IBM PureData System for Analytics N1001. The performance speed refers to the query times on both macro-analytic and mixed workload tests as conducted in IBM engineering lab benchmarks. The N2001 query times were an average of 3x faster than those of the N1001. Individual results may vary.

² 128 GB/sec scan rate assuming an average of 4x compression across the system. Individual results may vary.

³ Capacity of IBM PureData System for Analytics N2001 compared to previous generation IBM PureData System for Analytics N1001

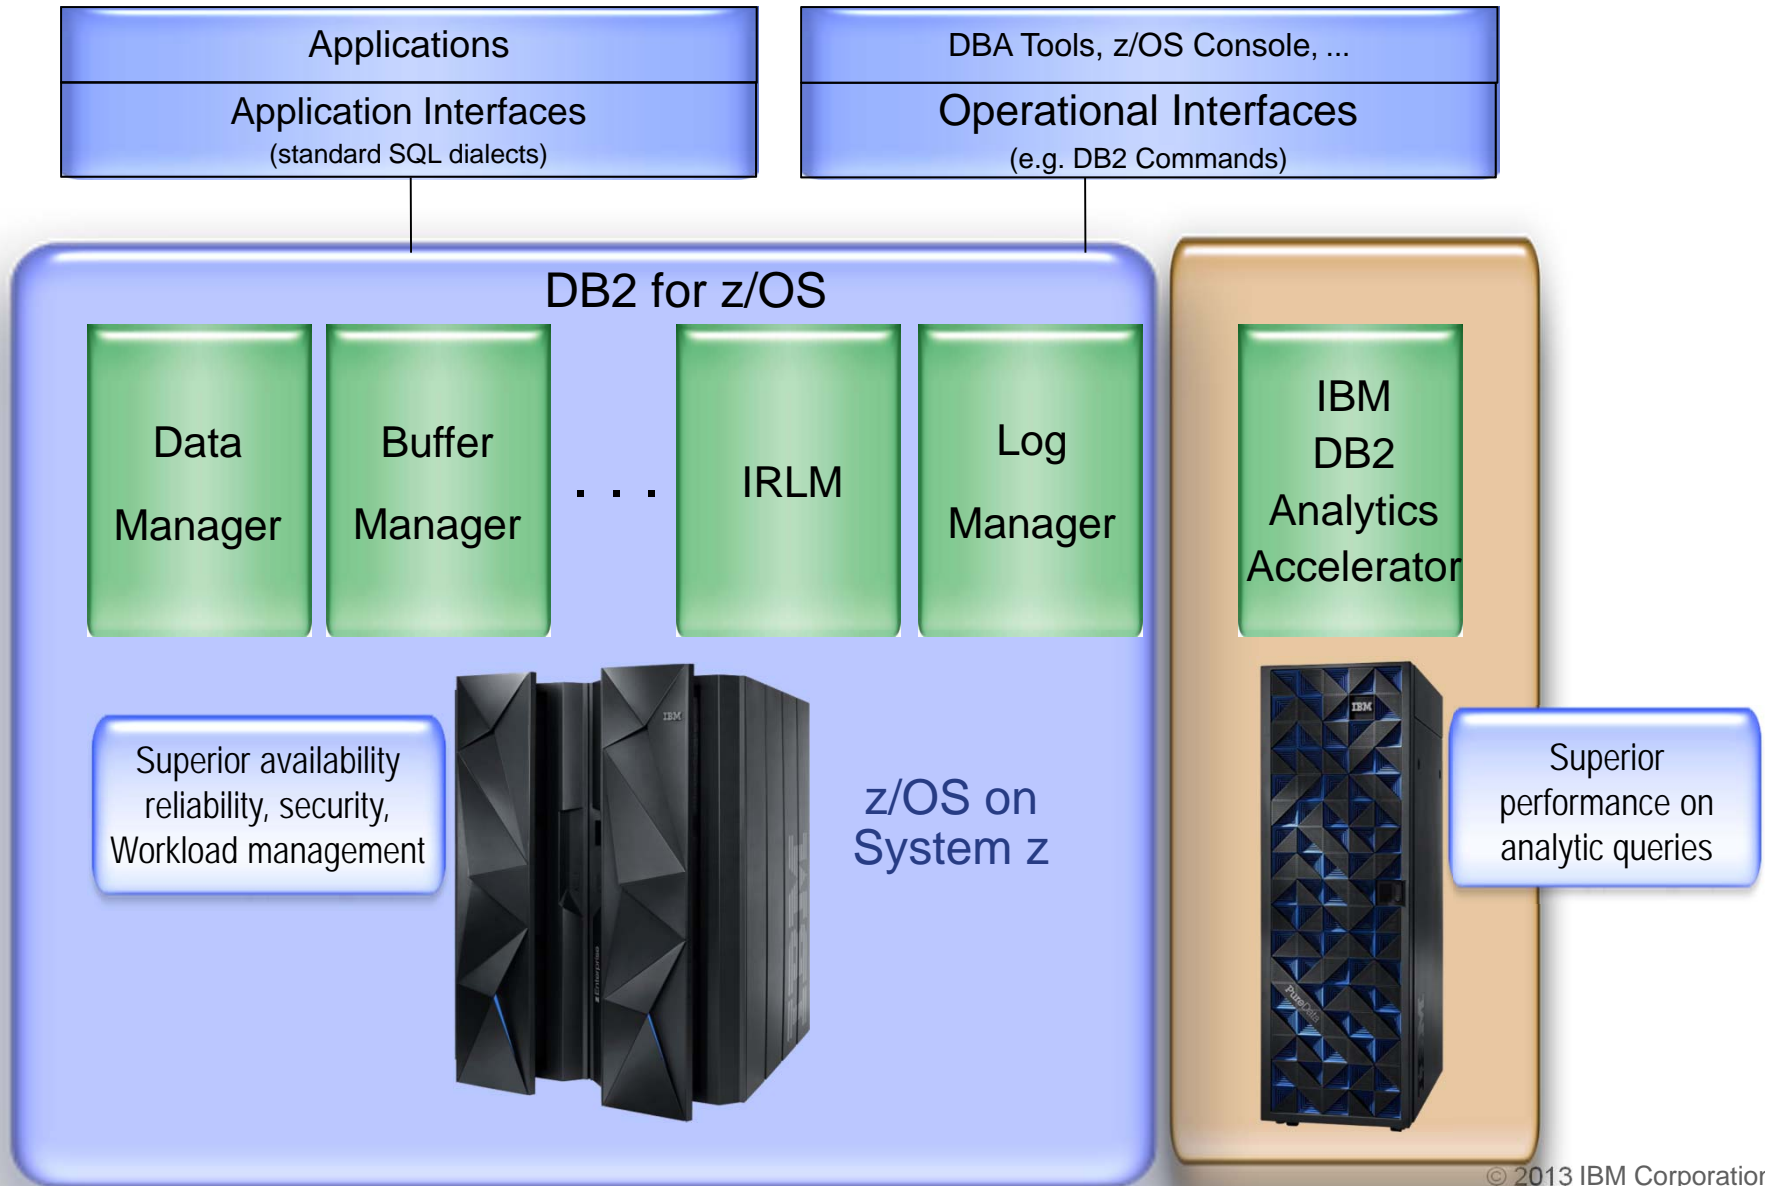
IBM DB2 Analytics Accelerator Product Components



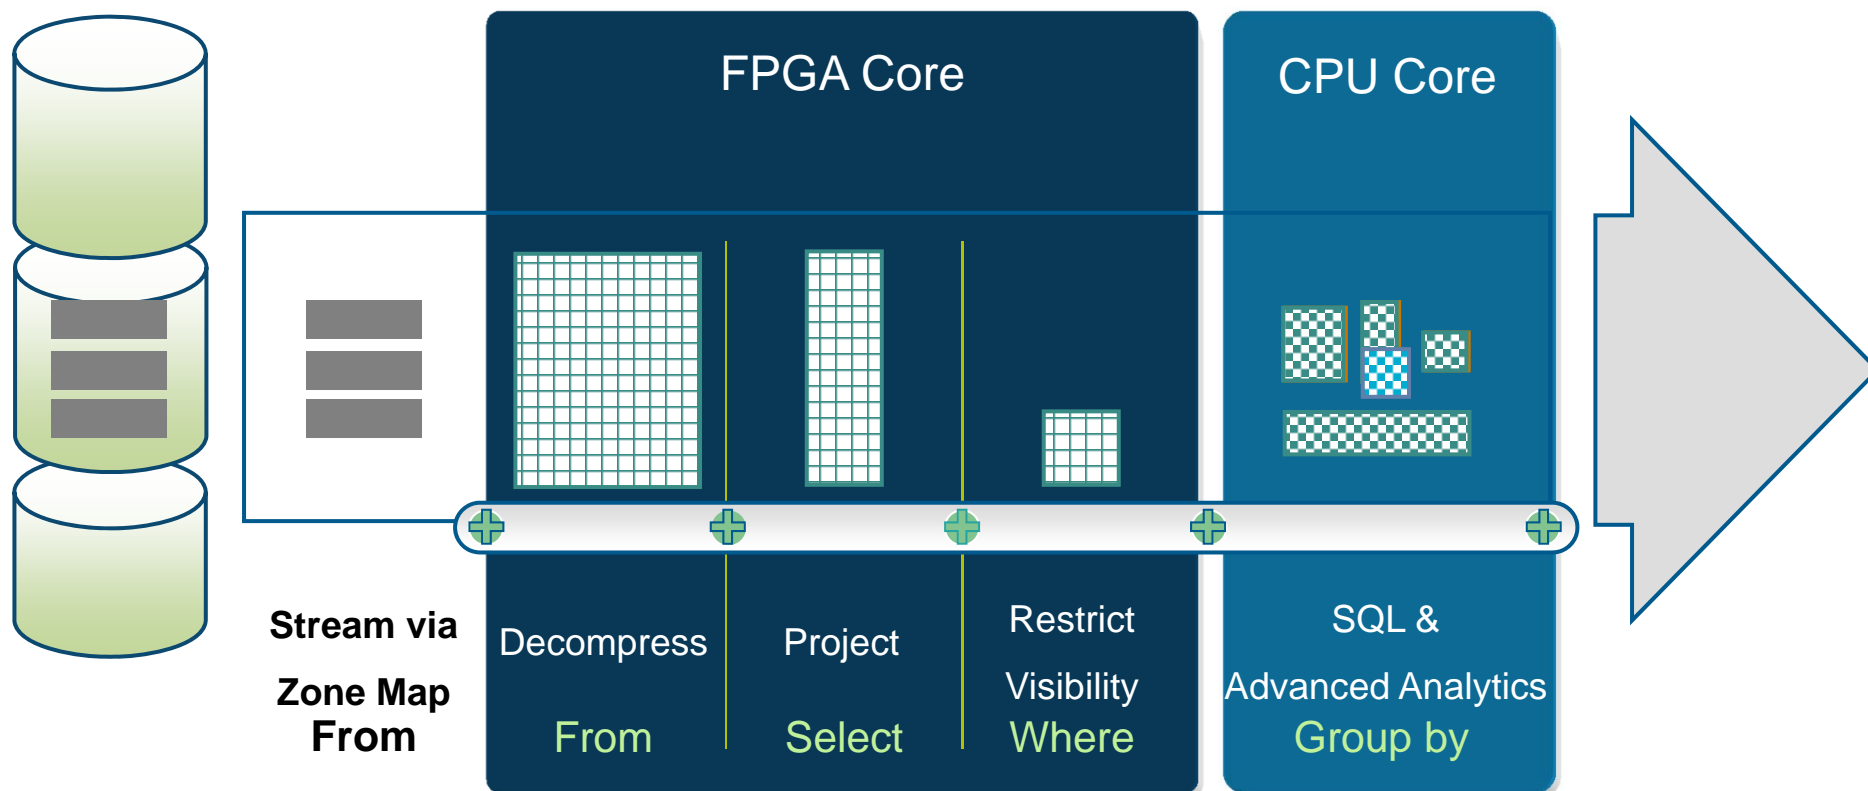
Note: There are several connection options using switches to increase redundancy



Deep DB2 Integration within zEnterprise



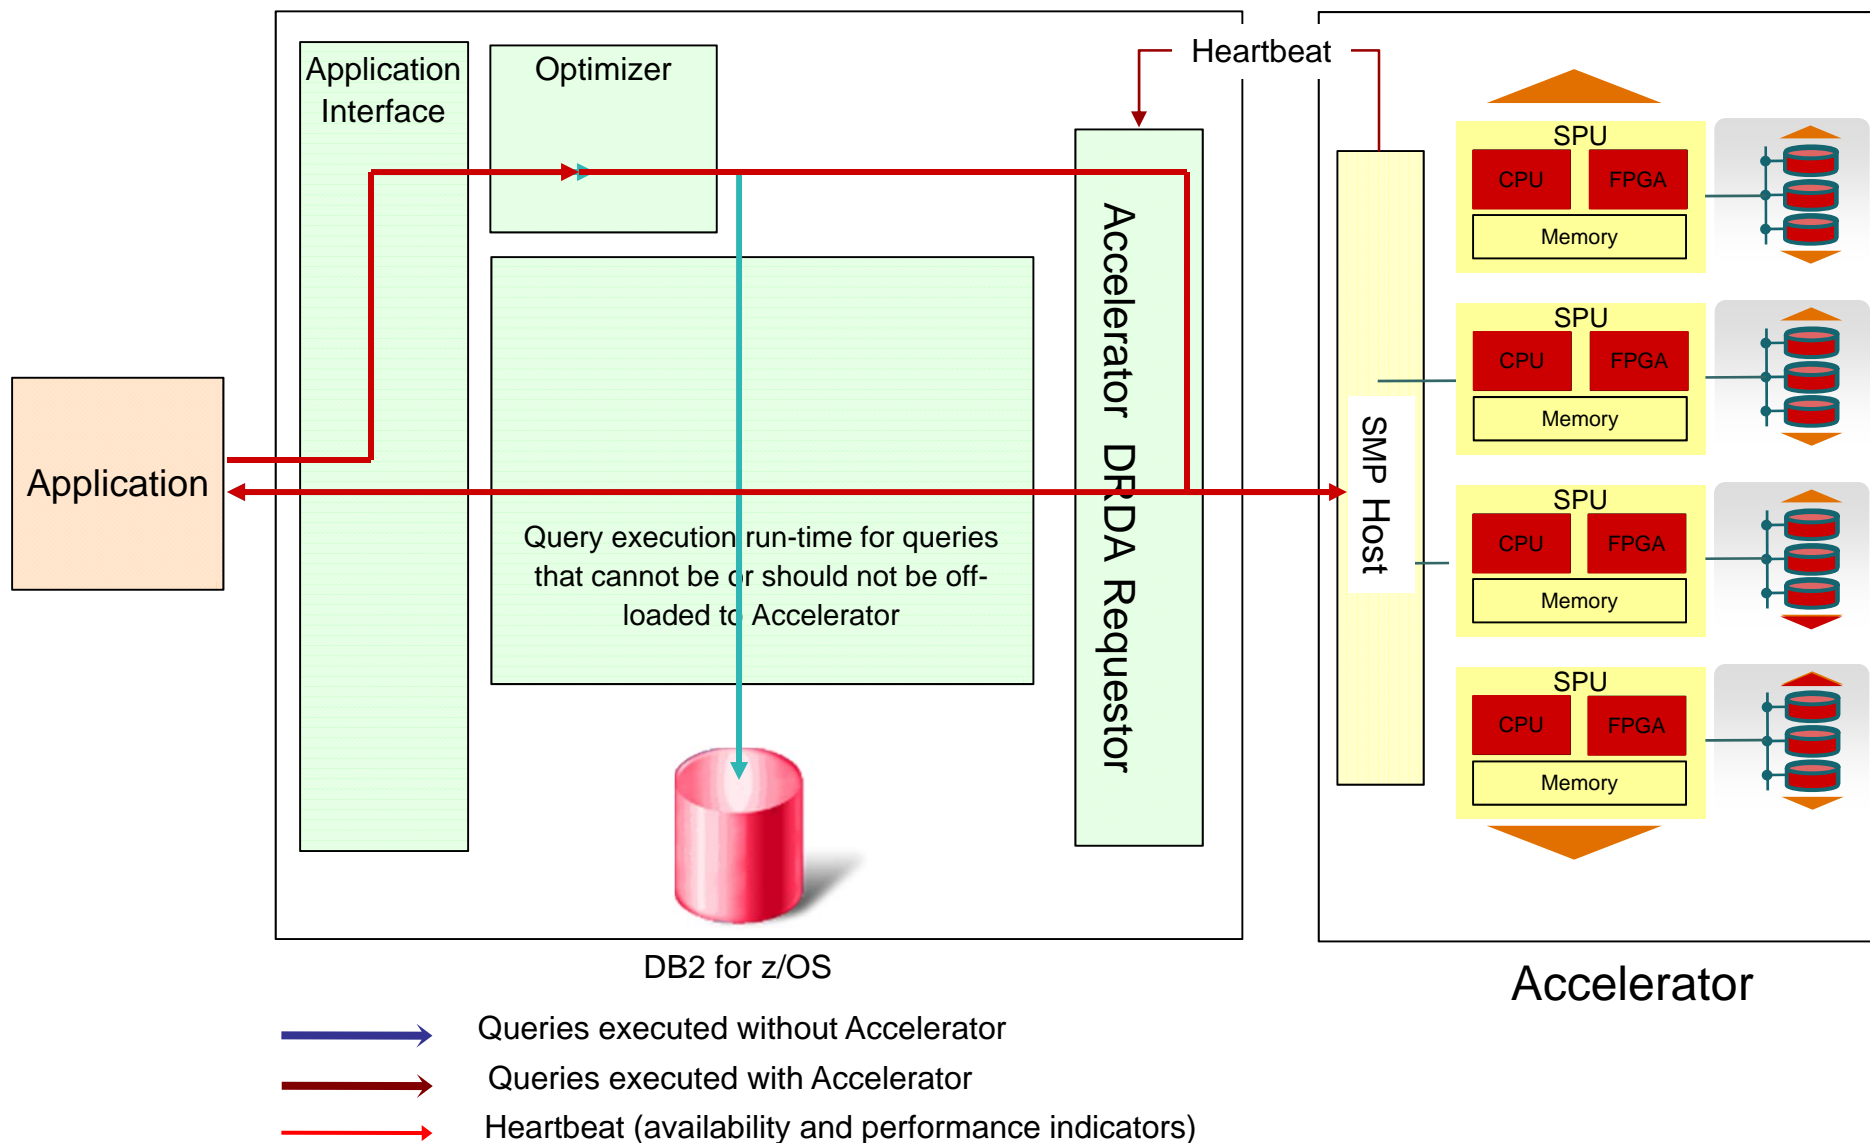
The Key to the Speed



```

Select State, Age, Gender, count(*) From MultiBillionRowCustomTable BirthDate BirthDate1960
And BirthDate < '1/1/1960' ( And, State in ('FL', 'CA') Group by State, Age, Gender, State, Age, Gender
Order by
State, Age, Gender
    
```

Query Execution Process Flow




High Performance Storage Saver

Reduces the cost of high speed storage

Ability to store historic data on the Accelerator only

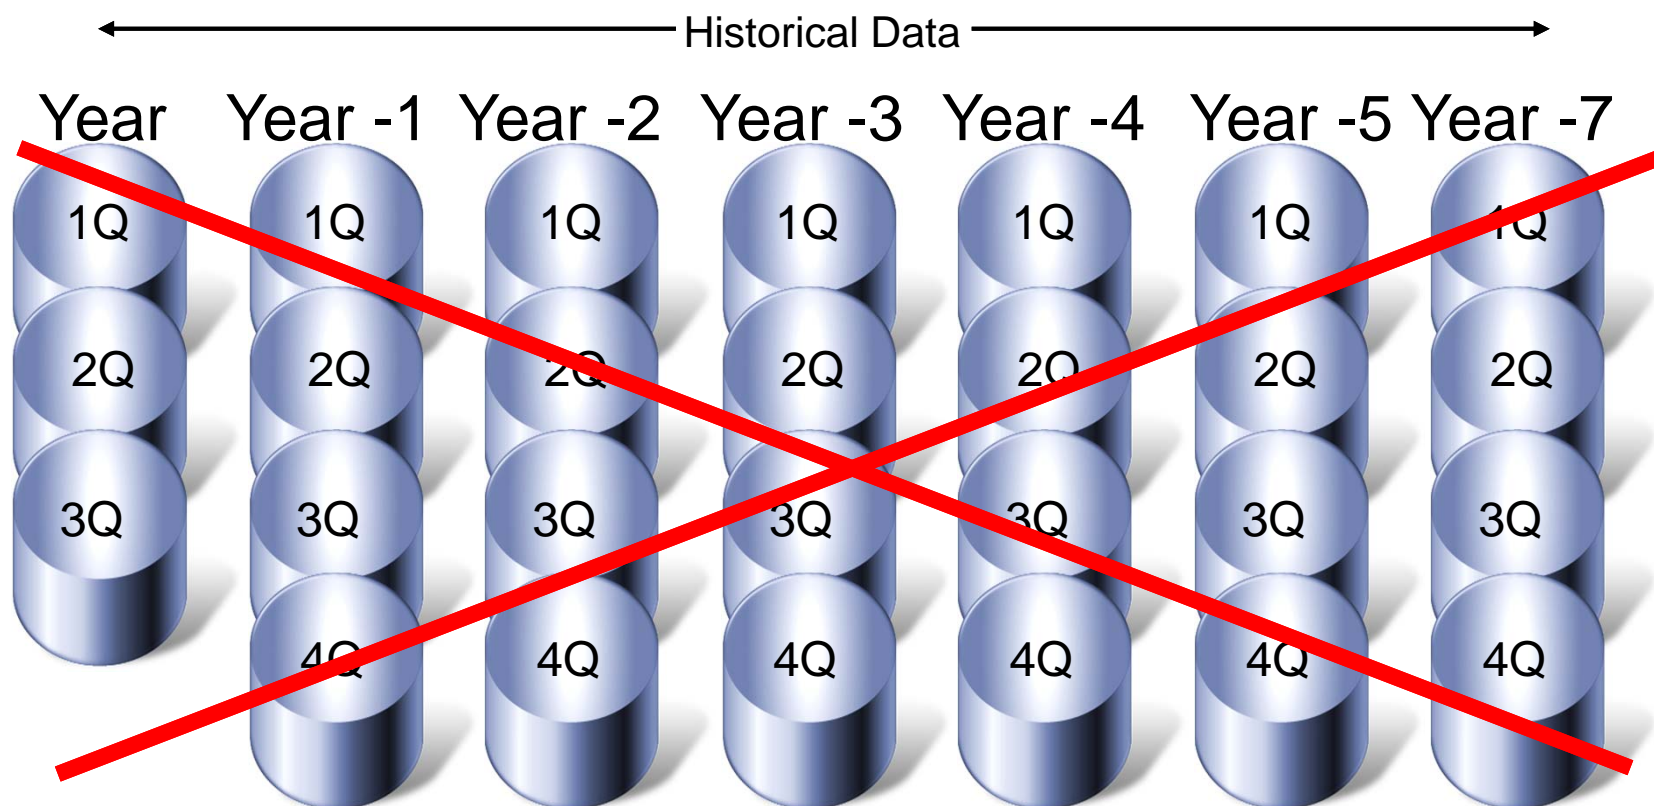
➤ You can choose the disk location for historical data to reduce host data warehouse storage usage by over 95%

- When data no longer requires updating, reclaim the DB2 storage
- Tables can be resident on:
 1. DB2 Only
 2. DB2 and Accelerator
 3. Accelerator Only
- Special Registers control behavior
 - CURRENT QUERY ACCELERATION
 - CURRENT GET_ACCEL_ARCHIVE
- Managed by zParms
- Enhanced Heuristics



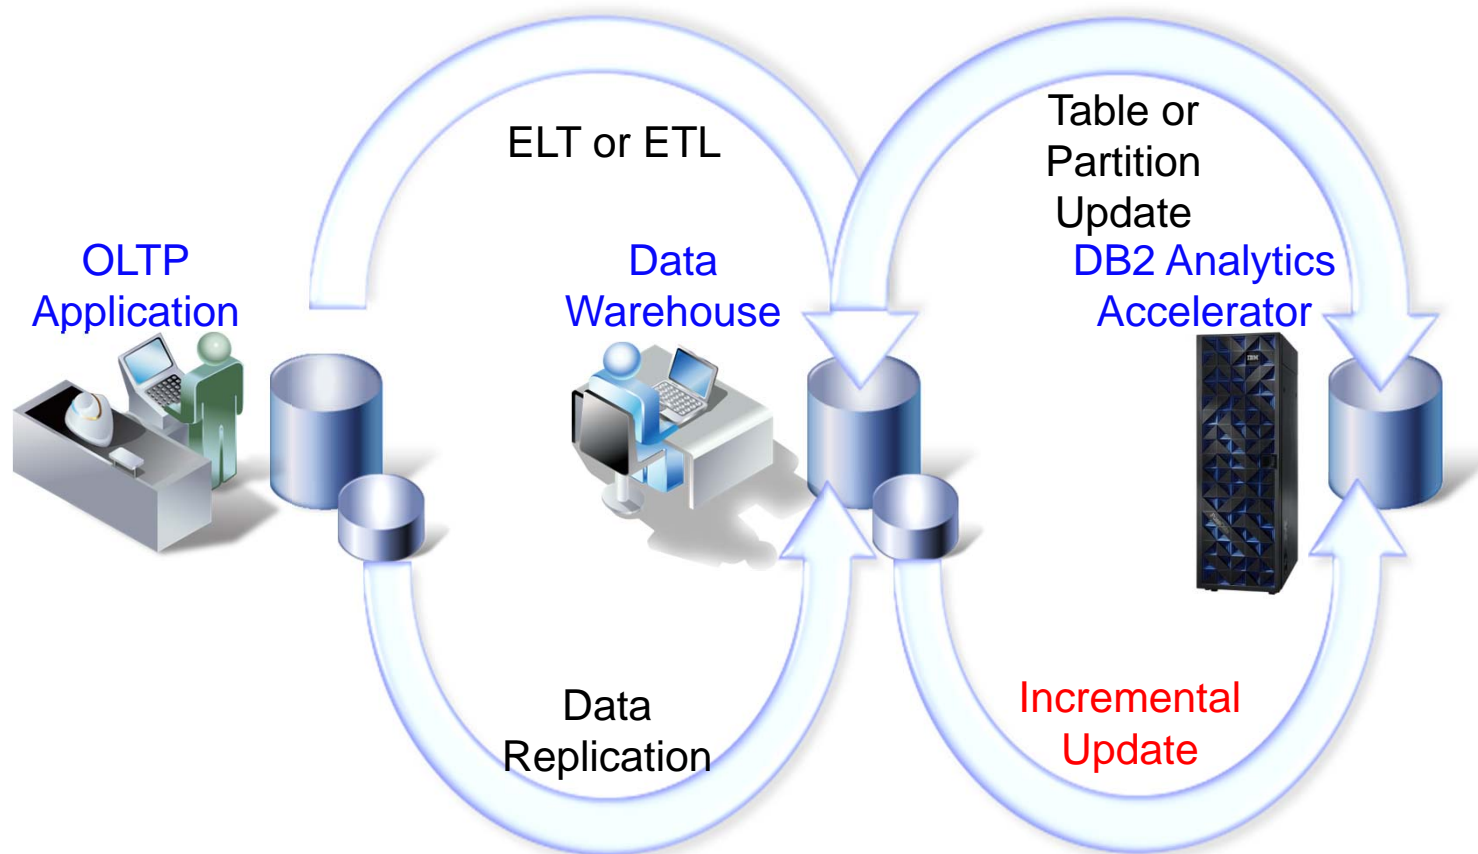


Save Over 95% of Host Disk Space for Historical Data



One Quarter = 3.57% of 7 years of data
One Month = 1.12% of 7 years of data
One month = 2.78% of 3 years of data

Incremental Update



Synchronizing data to lower data latency from days to minutes/seconds



DB2 Analytics Accelerator Future Directions

- **Acceleration of More Queries**
 - Greater workload acceleration with Static SQL support
 - Access path stability improvements
- **Increased Control and Monitoring**
 - High Perf Storage Saver partition read only
 - Incremental Update Trace Improvements
 - Improved Incremental Update Utility execution
- **Performance Improvements**
 - Exploitation of new N2001 hardware features
 - Extract Load Transform
 - Multi-row fetch for increased throughput
 - Workload balancing between accelerators
 - Improved Workload Management for high priority work
- **Scoring & Modeling**



IBM zEnterprise Analytics System 9700

Mixed Workloads for Next Generation Business Analytics



The next generation of System z analytics; an integrated solution of hardware, software and services that enables customers to rapidly deploy cost effective game changing analytics across their business.

Preselected

All the necessary components are identified and integrated into an end-to-end solution

Pretested

Over 20 different customer typical configurations are presized and tested

Solution Priced

Aggressively priced for a cost-effective add-on or new deployment for customers with critical data operations



IBM zEnterprise

An end-to-end, integrated solution for big data analytics

Bring your analytics to your data:

- 70% of the data used for analytics originates on zEnterprise

Easily delivers on modern analytics requirements for:

- Timely, accurate and secure
- Superior availability, scalability and performance
- Rapid deployment and expansion
- Reduced cost and complexity

Evolves with your business:

- Start where you want and grow without re-architecting





Thank You


Dan Wardman
Vice President
wardman@us.ibm.com
IM Mainframe Software
Site Exec, Silicon Valley Lab



Back up proof points DB2 Analytics Accelerator for z/OS



Typical Customer Example:



270 of the Mixed Workload Queries



Executes in DB2 returning results in seconds or sub-seconds

30 of the Mixed Workload Queries took minutes to hours

Query	Total Rows Reviewed	Total Rows Returned	DB2 Only		DB2 with IDAA		Times Faster
			Hours	Sec(s)	Hours	Sec(s)	
Query 1	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	3,425,29	724	0:44	2,640	0.0	2	1,320
Query 9	4,130,107	137	0:42	2,520	0.1	193	13

Successfully accelerated the problem queries without affecting the rest

Typical Experience - Fast Time to Value

- **IBM DB2 Analytics Accelerator**
 - ➔ 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 ...**
 - ➔ 400 GB “Loaded” in 29 Min
 - 570 million rows (Loads of 800GB to 1.3TB/Hr)
- **Actual Query Acceleration ... 1908x faster**
 - ➔ 2 Hours 39 Minutes to 5 Seconds
- **CPU Utilization Reduction**
 - ➔ 35% to ~0%



Actual customer results, October 2011



Moving the Data Warehouse back to zEnterprise *Performance Based Quotable Quotes...*

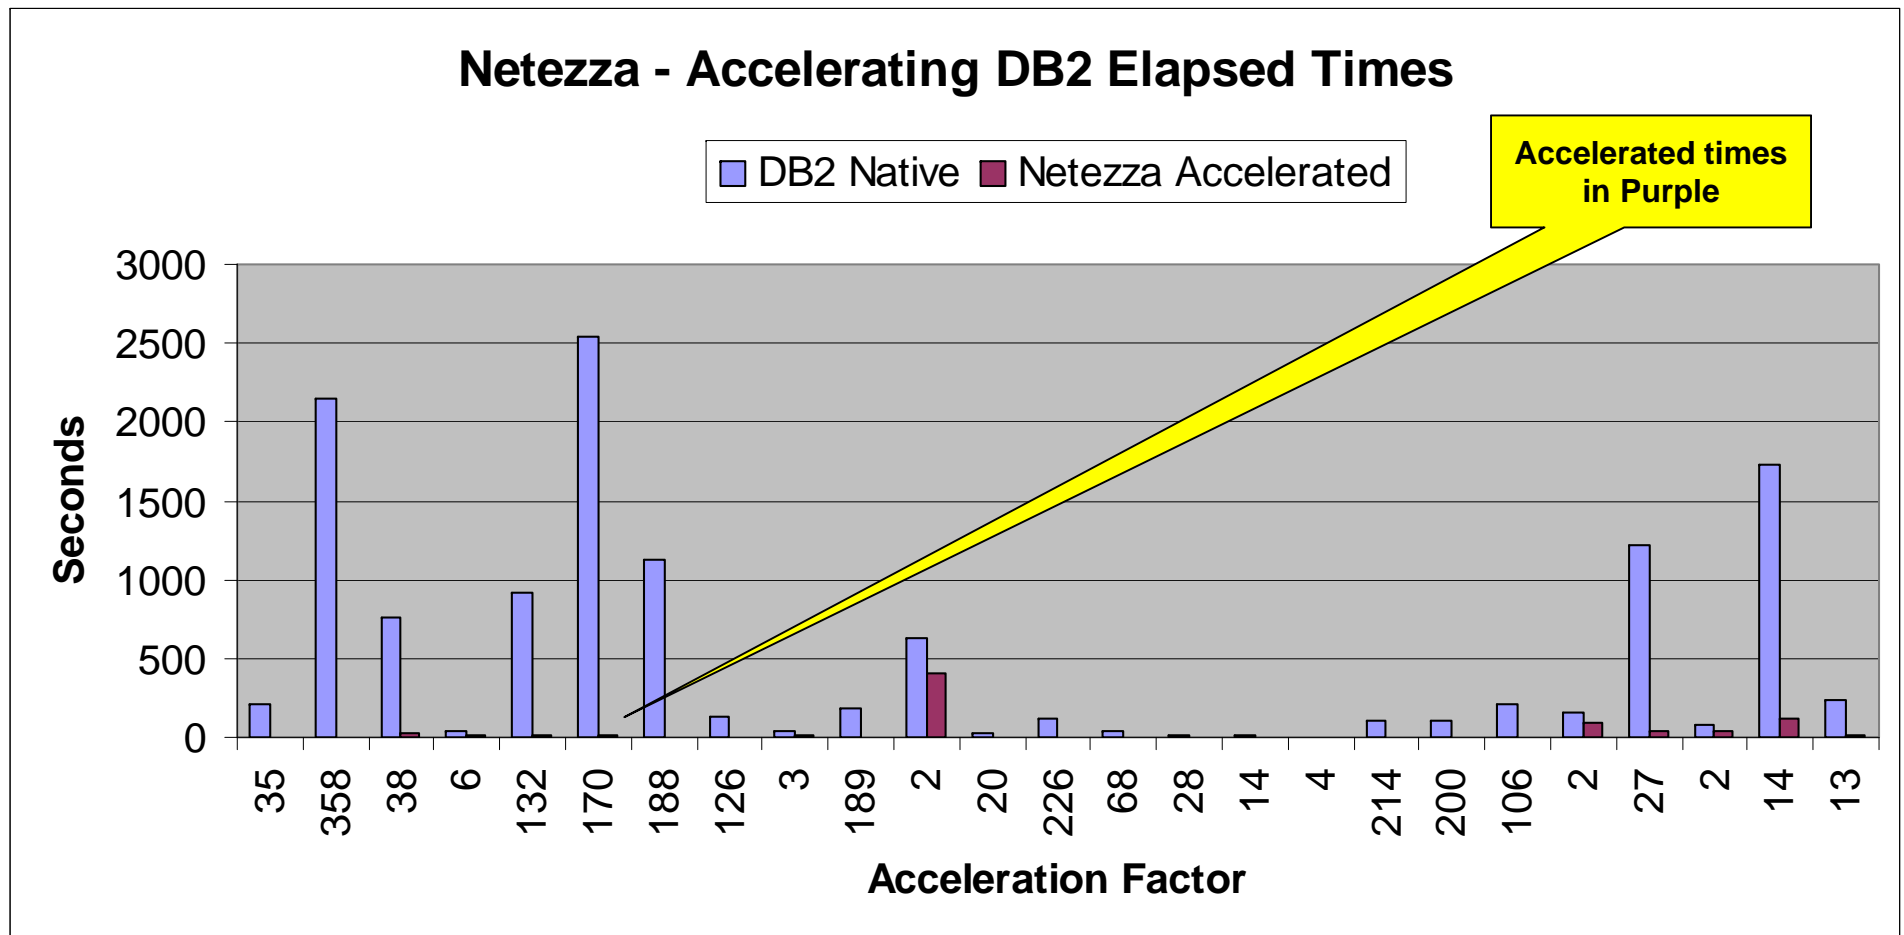
- **During first query submissions... “The accelerated version of that query just finished in 8 seconds, the DB2 version is still running” ... and it ran for 27 more minutes!**

- **“That one used to run in 50 minutes, now it runs in 47 seconds.... wait a minute, last time it ran in 34 seconds... oh, who cares when you are used to 50 minutes!”**
 - The difference ended up being a heavy LOAD in progress.
 - This “tens of minutes down to seconds” theme carried forward.

- **Asked the user to remove the filters in Business Objects, the response was "Are you sure you want me to do that?" as normally they would wait forever. When the screen came back at the "speed of light", the user was “Ecstatic”.**



'Wicked Fast' Performance for Complex DB2 Queries



Query Details

Several Ran Sub-Second

Up to 358x Faster!

Accelerated EXEC_DATE	Without Acceleration (seconds rounded)	Without Acceleration CPU (seconds rounded)	Accelerated (seconds rounded)	X Factor	ROWS_RETURNED	RESULT SIZE	STMT_TEXT
5/9/12 12:27 PM	210	7	6	35	26787	628K	T3.CNTR_NMBR, T3.MBR_ID,
5/9/12 12:43 PM	2146	261	6	358	3868	60.4K	T1.PRFX_CD,
5/9/12 12:45 PM	755	33	20	38	624	30.9K	T1.PRFX_CD,
5/9/12 12:46 PM	40	14	7	6	72627	2.98M	T1.CNTR_NMBR, T1.MBR_ID,
5/9/12 12:47 PM	921	186	7	132	4799	136K	PRFX_CD, CNTR_NMBR,
5/9/12 12:48 PM	2543	1695	15	170	3826	299K	T1.MBR_TRNS_EFCTV_DT,
5/9/12 1:02 PM	1128	536	6	188	122948	2.81MB	CNTR_NMBR,
5/9/12 12:56 PM	126	43	1	126	0	0 B	S.CNTR_NMBR, S.MBR_ID,
5/16/12 12:00 AM	2	0	not accel (not expected to)				
5/10/12 8:15 AM	34	13	11	3	528701	9.38 MB	TBL1.PRFX_CD,
5/10/12 8:37 AM	0	0	not accel (not expected to)				A.FR_MKT_SBSEG_CD AS
5/10/12 8:48 AM	189	24	1	189	131508	1.63 MB	AS CL_MKT_SEG_CD,
5/10/12 9:31 AM	623	418	409	2	18904956	631 MB	K_AFLT_ID, AFLT_ID AS
5/10/12 9:47 AM	20	8	1	20	131508	1.38 MB	K_CNTR_ID, CNTR_NMBR
5/10/12 9:51 AM	0	0	not accel (not expected to)		299		A.K_BS_RPT_UNT_ID AS
5/10/12 12:54 PM	113		1	226	1	33B	'FS AAE NONCNTR',

Avoided CPU Consumption

Avoided Redirecting




Existing Data Warehouse Performance Based Quotable Quotes...

- **Sample query run... native DB2 800 seconds elapsed time, DB2 Analytics Accelerator ran it in 6s 610ms.**
- **You are like a dealer... Now I expect everything to be this fast**
- **It is really impressive... queries that did not finish in DB2 can now be run**
- **178x faster is the current leader (among those that finished in DB2)**



Service Provider Query Results





Extending Netezza Technology to zEnterprise Data and Processes

- **Bring Massively Parallel Processing (MPP) technology to System z processes to accelerate queries and processes**
 - If it already runs on z, or all the data is sourced by z
 - Leverage core infrastructure, processes and people

- **Reduce cost of running existing processes**

- **Run queries that hit the Resource Limit Facility (RLF) limit before**

Operational Workload Assessment

Query Summary

	Total	DB2 natively	With potential	Uncertain	W/o potential
Queries	11804	10896 (92%)	908 (8%)	0 (0%)	0 (0%)
Elapsed Time (s) [1*]	35813.03	11127.10 (31%)	24685.92 (69%)	0.00 (0%)	0.00 (0%)
Elapsed Time (s)	42416.47	15134.51 (36%)	27281.96 (64%)	0.00 (0%)	0.00 (0%)
CPU Time (s) [1*]	4494.65	955.70 (21%)	3538.95 (79%)	0.00 (0%)	0.00 (0%)
CPU Time (s)	5631.97	1260.64 (22%)	4371.32 (78%)	0.00 (0%)	0.00 (0%)

[1*] - Considers each query only once


Ignore non-select statements

Query classification	Unique Queries	Executed Queries	CPU %	CPU eligible %
Total	11804	941528	100%	78%
> 60 min elapsed time (info)	0	0	0%	0%
10-60 min elapsed time (info)	2	2	2%	0%
1-10 min elapsed time (info)	183	196	66%	51%
< 1 min elapsed time (info)	11619	941330	32%	27%


Reason breakdown for 0% queries with no potential and 0% of elapsed time with no potential

Reason	# Queries	% Queries	% Elapsed Time

Start trace time	End explain time	Min time stmt cached	Max time stmt cached
Apr 22, 2012 5:27 AM	May 8, 2012 12:36 PM	Apr 22, 2012 5:48 AM	May 8, 2012 12:36 PM




Elapsed time for best DB2 native processing
 Elapsed time with acceleration potential
 Elapsed time with uncertain potential
 Elapsed time without acceleration potential



Query CPU time for best DB2 native processing
 Query CPU time with acceleration potential
 Query CPU time with uncertain potential
 Query CPU time without acceleration potential

N2001 Hardware Overview



- 12 Disk Enclosures
 - 288 600 GB SAS2 Drives
 - 240 User Data, 14 S-Blade
 - 34 Spare
 - RAID 1 Mirroring
- 2 Hosts (Active-Passive)
 - 2 6-Core Intel 3.46 GHz CPUs
 - 7x300 GB SAS Drives
 - Red Hat Linux 6 64-bit
- 7 PureData for Analytics S-Blades™
 - 2 Intel 8 Core 2+ GHz CPUs
 - 2 8-Engine Xilinx Virtex-6 FPGAs
 - 128 GB RAM + 8 GB slice buffer
 - Linux 64-bit Kernel

Scales from
1/2 Rack to 4 Racks

- User Data Capacity: 192 TB*
- Data Scan Speed: 478 TB/hr*
- Load Speed (per system): 5+ TB/hr

- Power Requirements: 7.5 kW
- Cooling Requirements: 27,000 BTU/hr

* 4X compression assumed
* 4X compression assumed



The experience so far...

- **Up and running in 1 day, quick start training completed**

- **1300+ tables loaded**

- **First query... in DB2, it ran for 11 minutes, 31 seconds 41 milliseconds. In the DB2 Analytics Accelerator, it ran for 1 second and 576 milliseconds.**
 - Customer really wants to shave off that 576 milliseconds.... we are working on it.

- **Typical query...**
 - DB2 Native Run - Elapsed Time 17 minutes, Normalized CPU Time 1 hour 33 minutes, **Cost \$166.00**
 - DB2 Analytics Accelerator Run - Elapsed Time 59 seconds, Normalized CPU Time 10 minutes, **Cost \$12.50**
 - Quote of the day... **"Just think, these run nightly and there are hundreds of them."**



DB2 10 for z/OS

- **CPU reductions for transactions, queries, and batch**
 - Out-of-the-box CPU reductions of 5-10% for traditional workloads
 - Up to additional 10% CPU savings using new functions or avoiding constraints
 - Out-of-the box CPU reductions of up to 20% for new workloads
- **Scales with less complexity and cost**
 - 5-10x more concurrent users – up to 20,000 per subsystem
 - Significant scale-up capabilities in addition to existing scale-out support
 - Consolidate to fewer LPARs and subsystems
- **Improved operational efficiencies and lower administration cost**
 - Automatic diagnostics, tuning, and compression
- **Even better performance**
 - Elapsed time improvement for small LOBS and Complex Queries

➤ ***64 bit Evolution
Virtual Storage
Relief***

➤ ***Temporal Data***

➤ ***Integrated XML
Support***

➤ ***Query Processing
Enhancements***

➤ ***Business Security
& Compliance***

➤ ***Better
Productivity***