Positioning Your Enterprise for Cloud, Analytics and Mobile Computing

Scoring Fast and Winning Big with Analytics on z Systems

Numerous studies show how businesses gain competitive advantage by using analytics

#1 CIOs rank analytics as the #1 factor contributing to an organization's competitiveness[1]

54% of CxOs say customers influence them to a *large* extent[5]

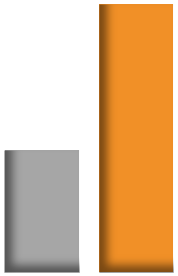

Organizations that embrace analytics are more than **2x** more likely to outperform their peers[2]

Financial outperformers are **64%** more likely to use analytics to evaluate talent supply and demand on an ongoing basis[3]

Enterprises that apply advanced analytics have **33%** more revenue growth and **12x** more profit growth[4]

[1] IBM CIO Study 2009
[2] IBM IBV/MIT Sloan Management Review Study 2011
[3] IBM CHRO Study 2010
[4] IBM CFO Study 2010
[5] IBM Institute of Business Value, "The Customer-Activated Enterprise"

Many leading businesses use IBM analytics systems and software to gain that edge

SICOOB

A Brazilian credit union realizes **200%** internet growth and **600%** overall growth, sustaining it over 2 million members

PETROL

Slovenian automotive goods and services company implements smarter commerce – suggest-selling at point-of-sale – to significantly increase sales

US-base cancer research center realizes **100%** payback in 3 months through proactive identification of fraudulent activities, and optimizes financial compliance processes

The more a business uses analytics, the better it performs

Running analytics off-platform doesn't pay for a mainframe-centric business…

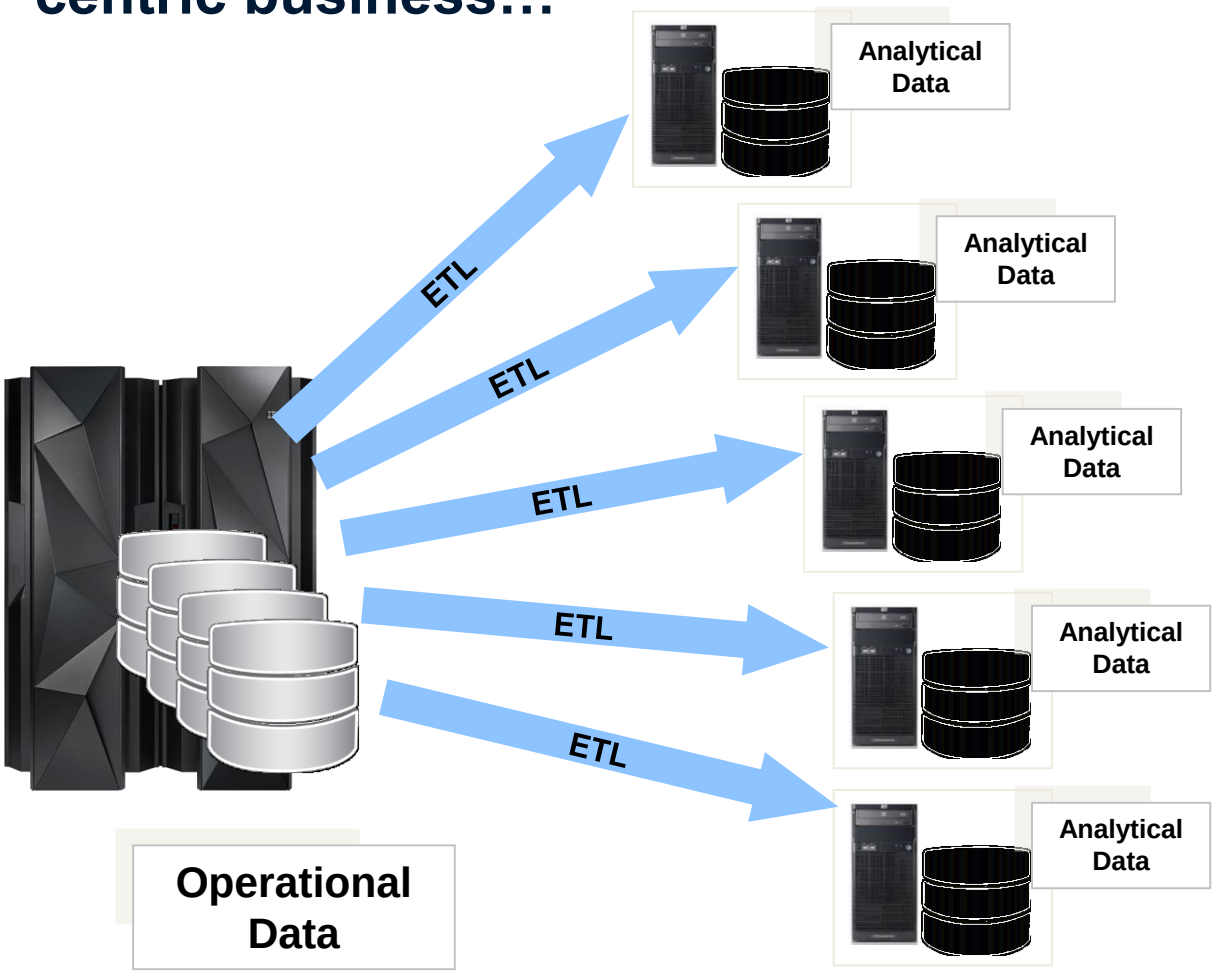

A large European bank:

- 120 database images created from bulk data transfers
- 1,000 applications on 750 cores with 14,000 software titles
- ETL consuming 28% of total distributed cores and 16% of total MIPS

A large Asian bank:

- One mainframe devoted exclusively to bulk data transfers
- ETL consuming 8% of total distributed core and 18% of total MIPS

With this strategy, IT costs grow faster than business growth

Source: IBM Eagle Studies

… Rather it leads to significant data transfer costs

Example:

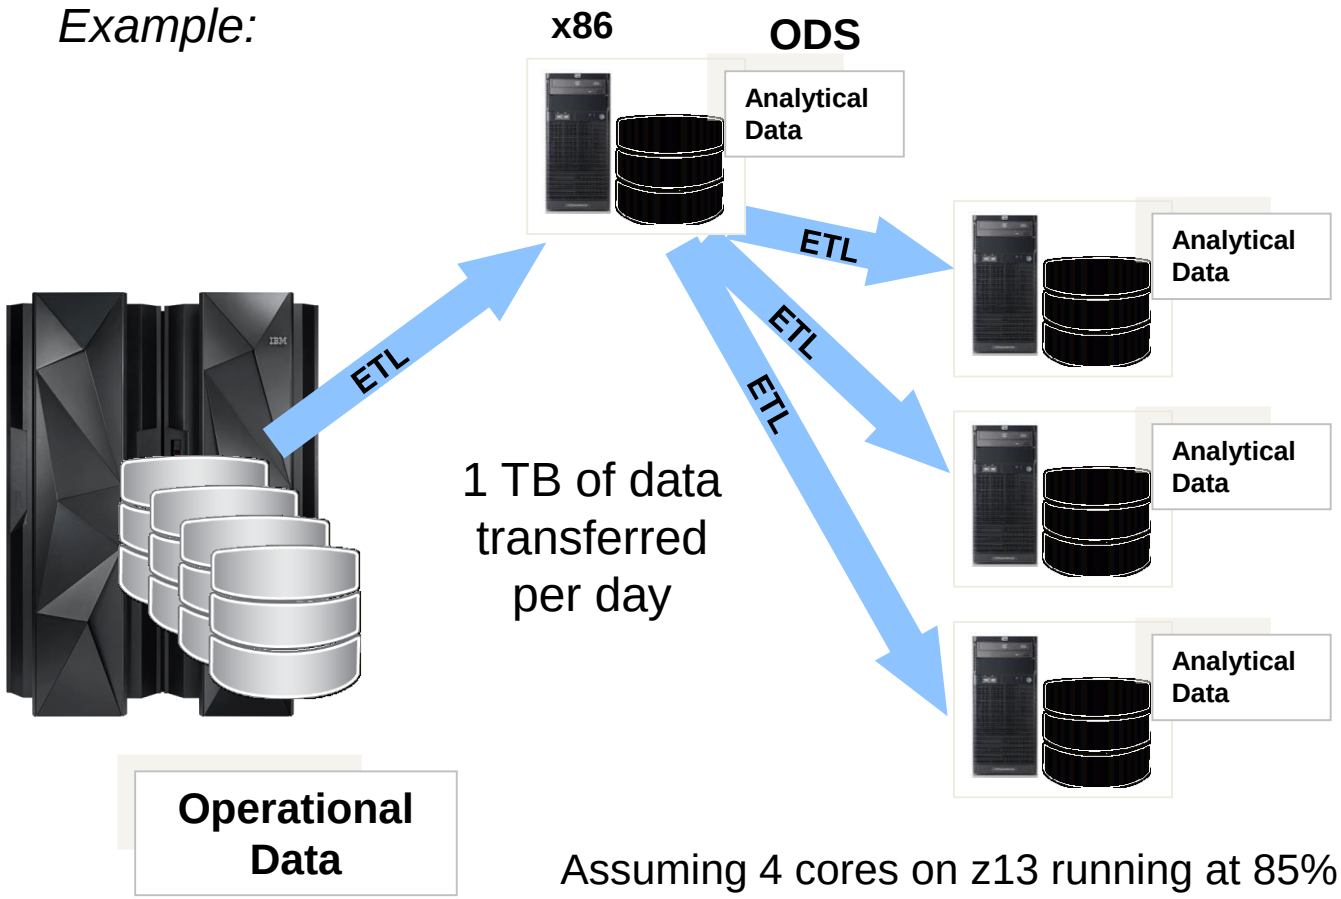

Estimated 4 yr. cost summary

System costs = $9,864,412
Labor costs = $393,927
Total = $10,258,339

Assuming 4 cores on z13 running at 85% utilization and 12 cores on x86 servers run at 45% utilization, transfer will burn **519 MIPS** and use **10 x86 cores** *per day*

Source: IBM CPO internal study

Today, z Systems are designed to run all analytics on-platform, creating a first-class System of Insight

Gain a competitive edge by co-locating analytics software with data and accelerators in the System of Record

- Accelerate operational analytics with a hybrid database management system

System of Record

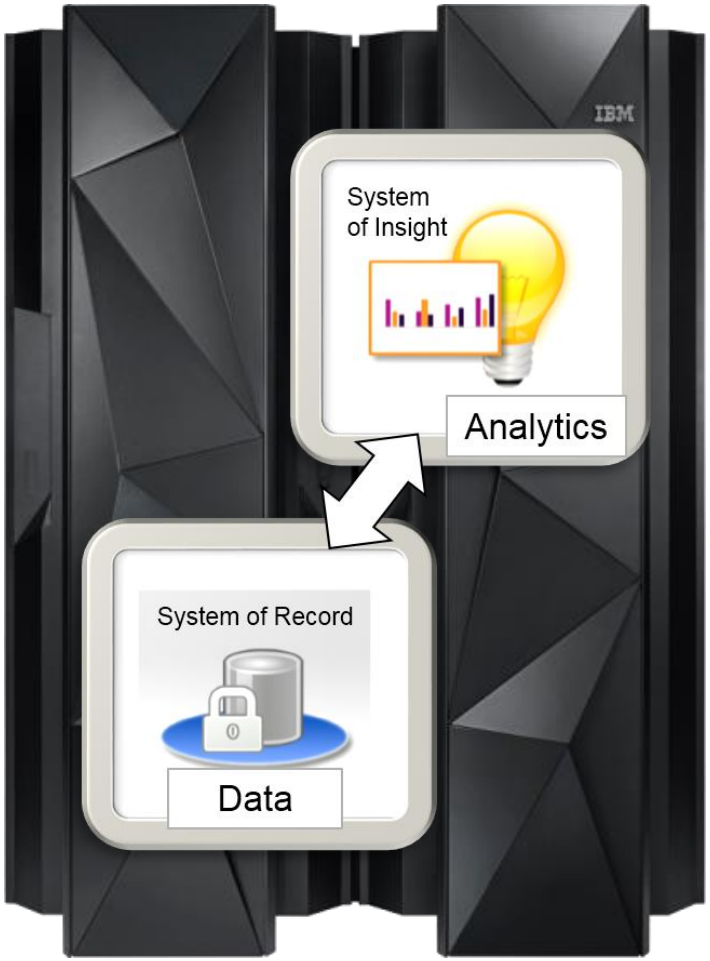

System of Insight

- Create 360º view of customers using Hadoop and descriptive analytics
- Use predictive analytics and real-time in-transaction scoring
- Leverage columnar analysis option

Run a complete solution – including query acceleration, Big Data, BI, and predictive analytics – on z Systems

Data Store
DB2 for z/OS

Big Data (Hadoop)
InfoSphere BigInsights

Business Intelligence and Reporting
IMB Cognos Enterprise

Predictive Analytics, Modeling, Scoring
IBM SPSS

BLU Acceleration
DB2 LUW

IBM z Systems

DB2 Analytics Accelerator

Green boxes denote Linux on z software.
Blue denotes z/OS software. Cognos runs on both.

DB2 for z/OS and the DB2 Analytics Accelerator create a hybrid database management system…

… to accelerate operational analytics

Data Store
DB2 for z/OS

- Uniform and transparent access for transactional and analytical applications

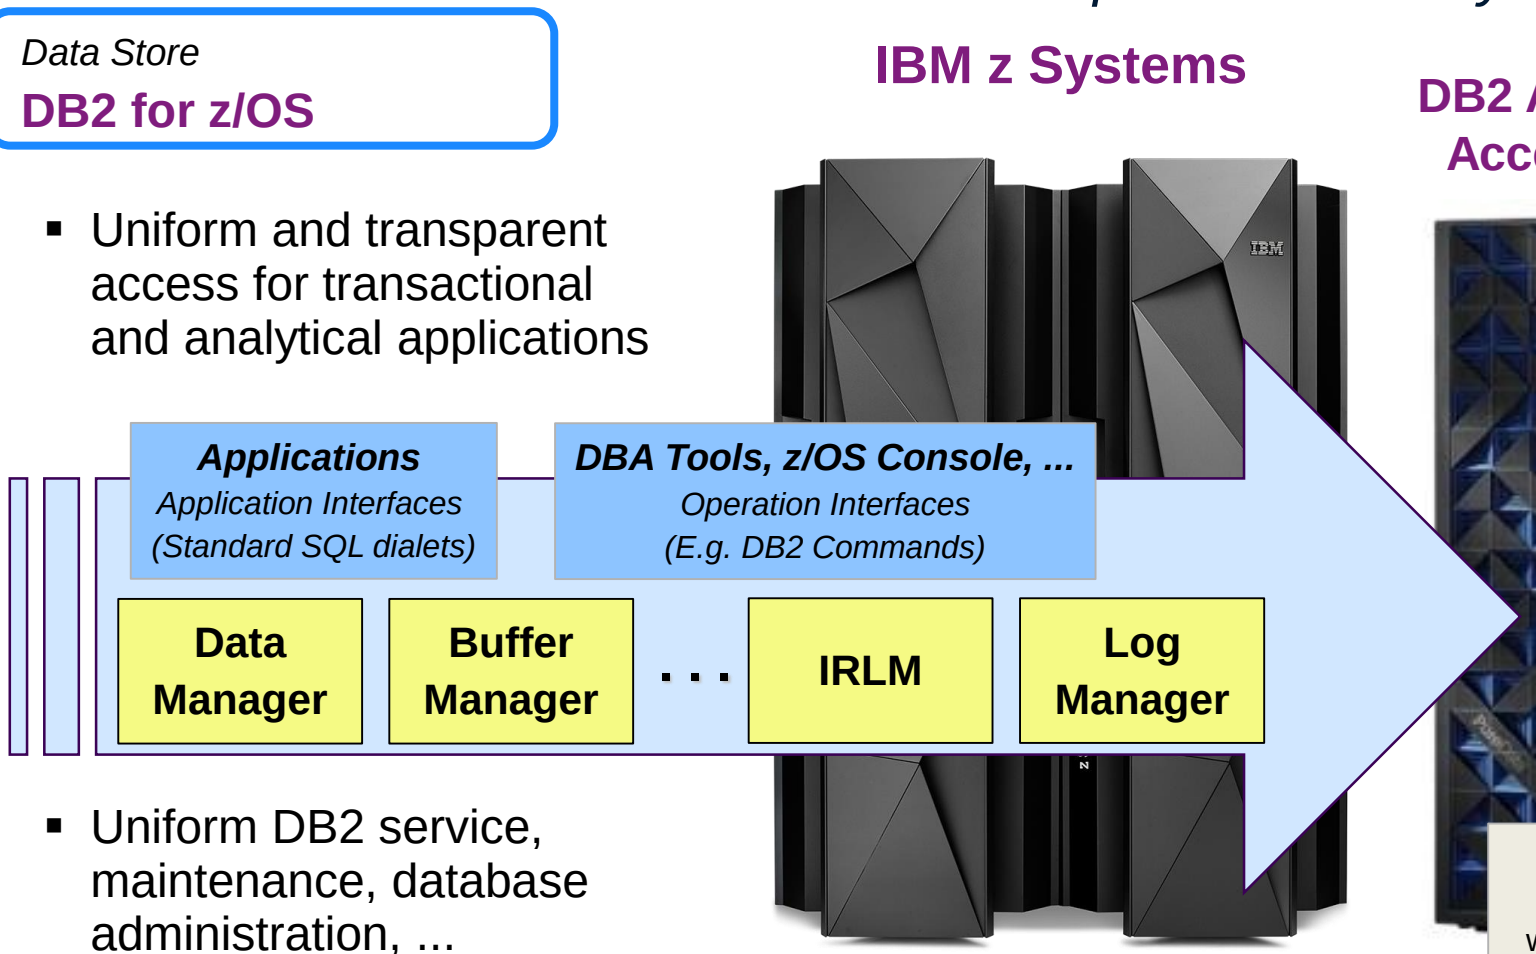

- Uniform DB2 service, maintenance, database administration, ...

Powered by PDA – A true appliance, with industry leading ease of performance

DB2 Analytics Accelerator uses S-Blade technology for industry unique data stream processing…

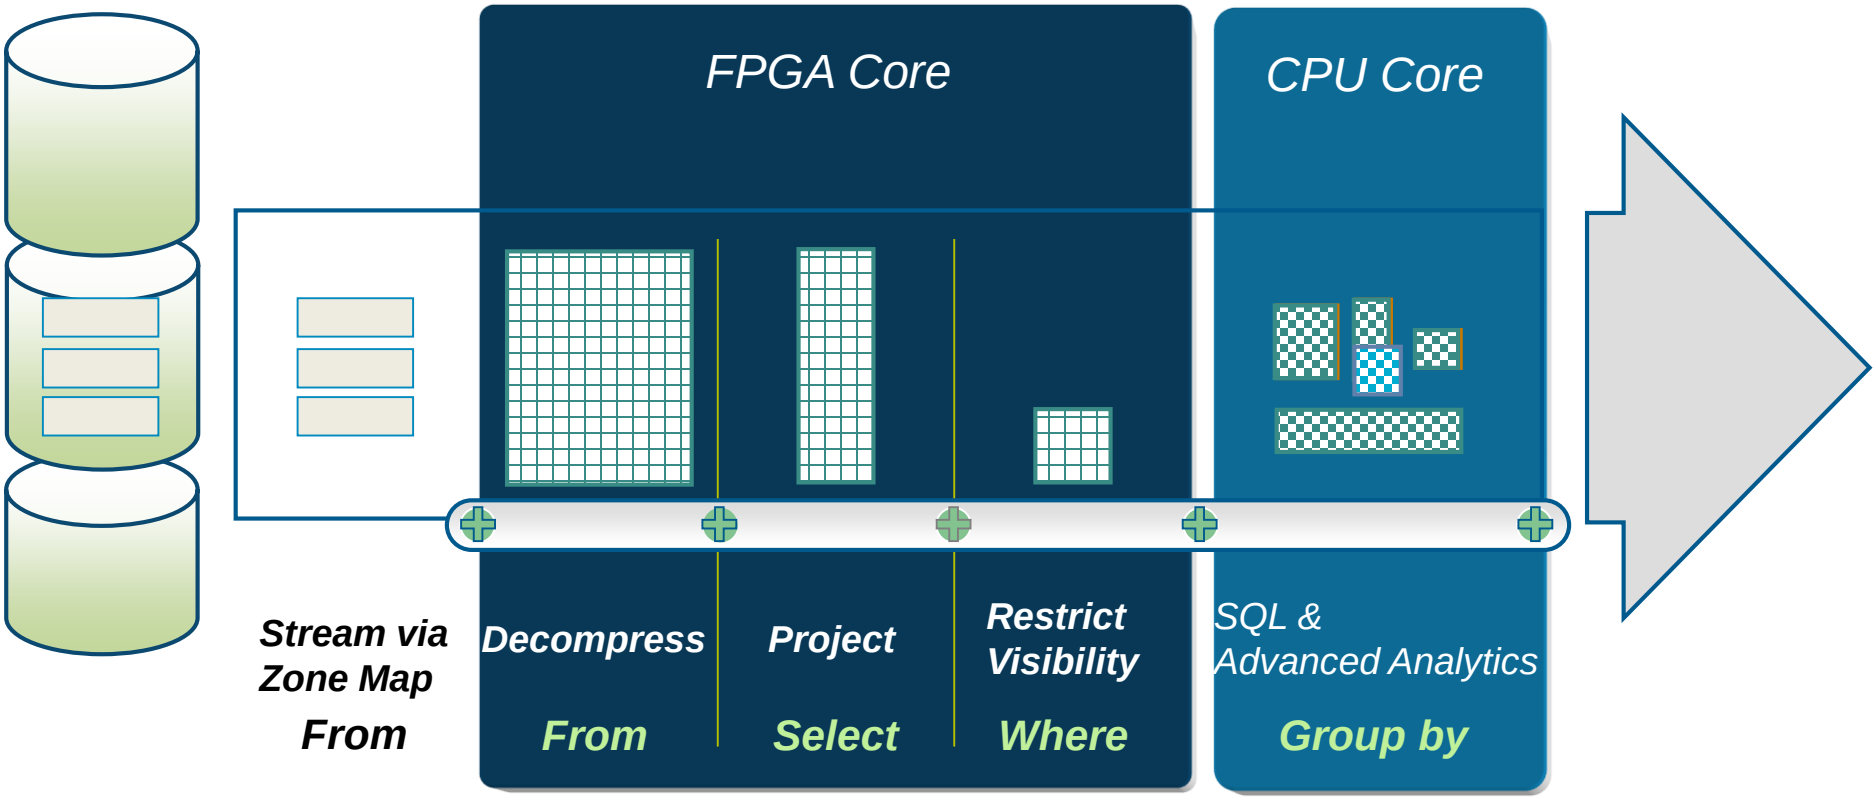

Select State, Age, Gender, count(*) From MyBigTable Where Birth Date > ‘01/01/1960’ And State in (‘CA’, ‘NY’, ‘GA’) Group by State, Age, Gender Order by State, Age, Gender

…which drives blazing speed through balanced design

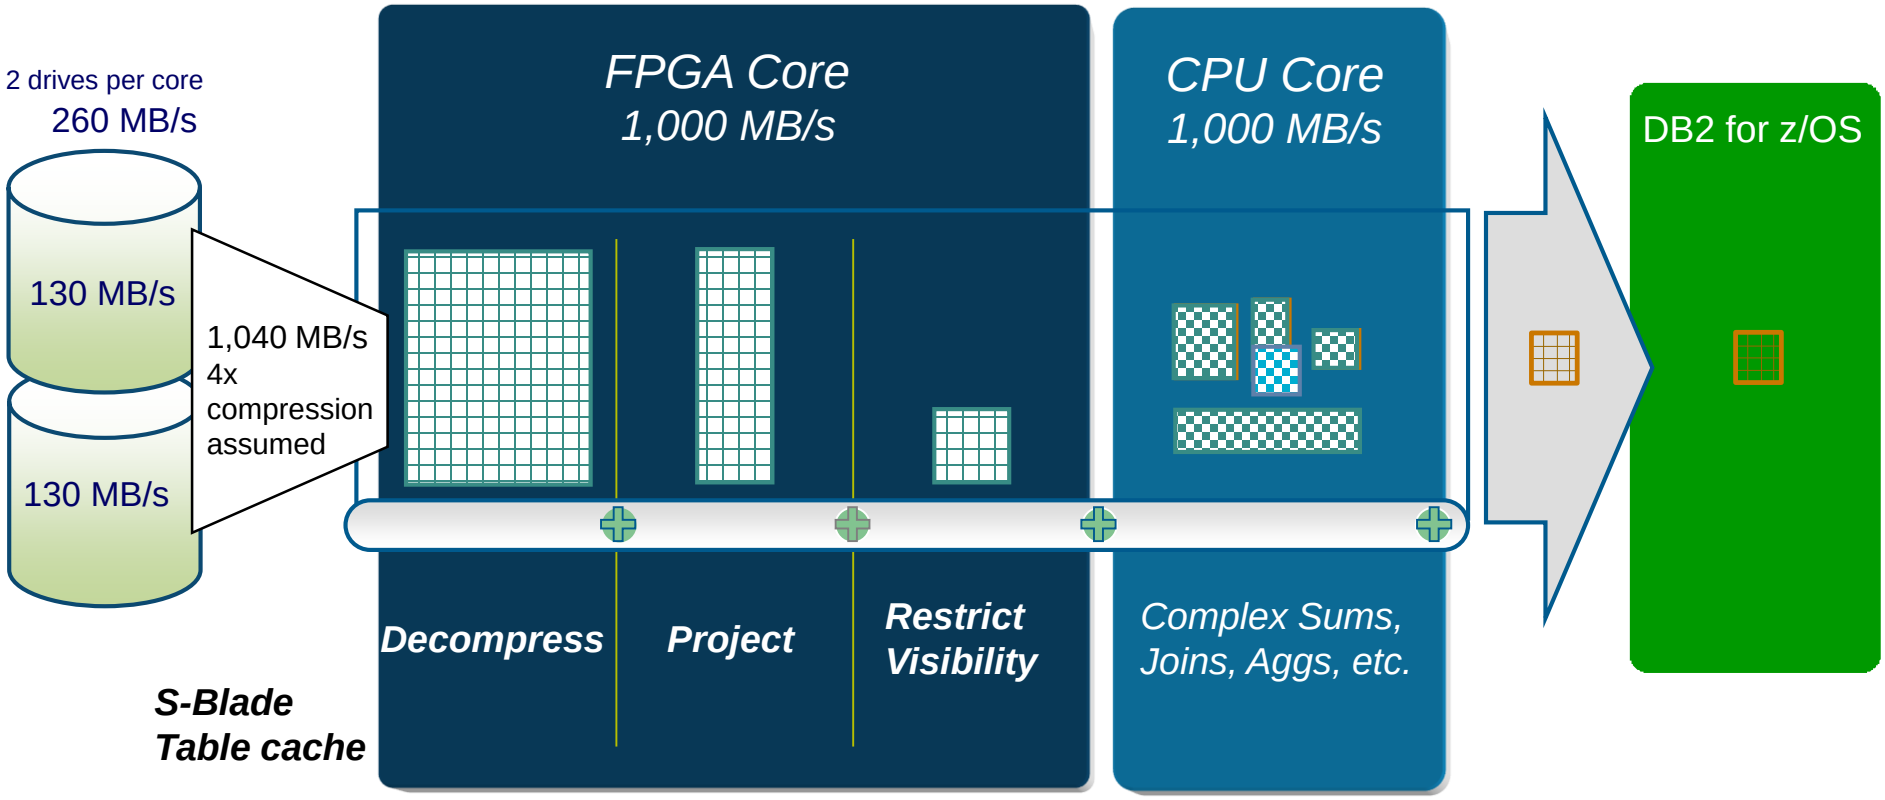

Using the DB2 Analytics Accelerator as the analytics data store saves over 88% in ETL and transfer costs

Example:

z13

IBM DB2 Analytics Accelerator (N3001-010)

Estimated 4 yr. cost summary

System costs = $1,052,901
Labor costs = $137,613
Total = $1,190,513

88% Lower cost

Estimated for systems compared

Operational Data

Assuming 4 cores on z13 running at 85% utilization and 140 x86 cores on N3001-010 running at 45% utilization, transfer will burn **260 MIPS** and use **0.44 x86 core** *per day*

This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. Test involved measuring in a controlled laboratory environment elapsed time for system and administrator to extract, send and receive 1,118GB file from z13 to DB2 Analytics Accelerator N3001-010 (Mako Full Rack. Prices, where applicable, are based on US prices as of 12/31/2014 for both IBM and competitor. Estimated amortized cost from 4 Year Total Cost of Acquisition (TCA) that includes all HW, SW (OS, DB and tools) and 4 years of service & support. For Labor costs, used annual burdened rate of $159,600 for IT Administrator for z Systems and x86. Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment. Users of this document should verify the applicable data for their specific environment.

Maintaining hardware and software currency of z Systems and DB2 will improve performance

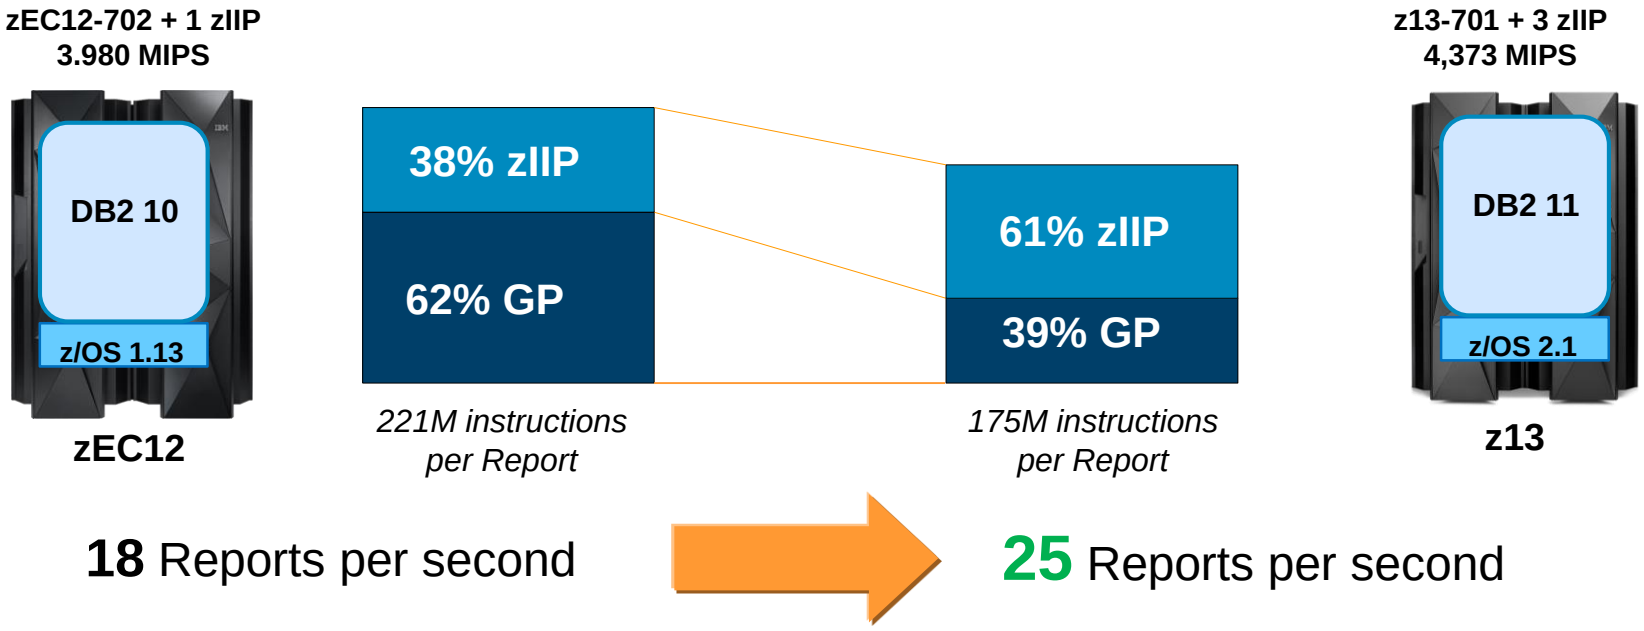

- **Over 60% zIIP offload** – from newest generation of specialty processors with SMT – yields **better price performance**
- **21% shorter path length** – resulting from DB2 for z/OS upgrade – **reduces CPU usage**
- **39% higher throughput** – from combined effects of software and hardware upgrade - **reduces elapsed time** to execute operational reports

Continuous optimizations of the platform and the Analytics Accelerator improve throughput and price performance

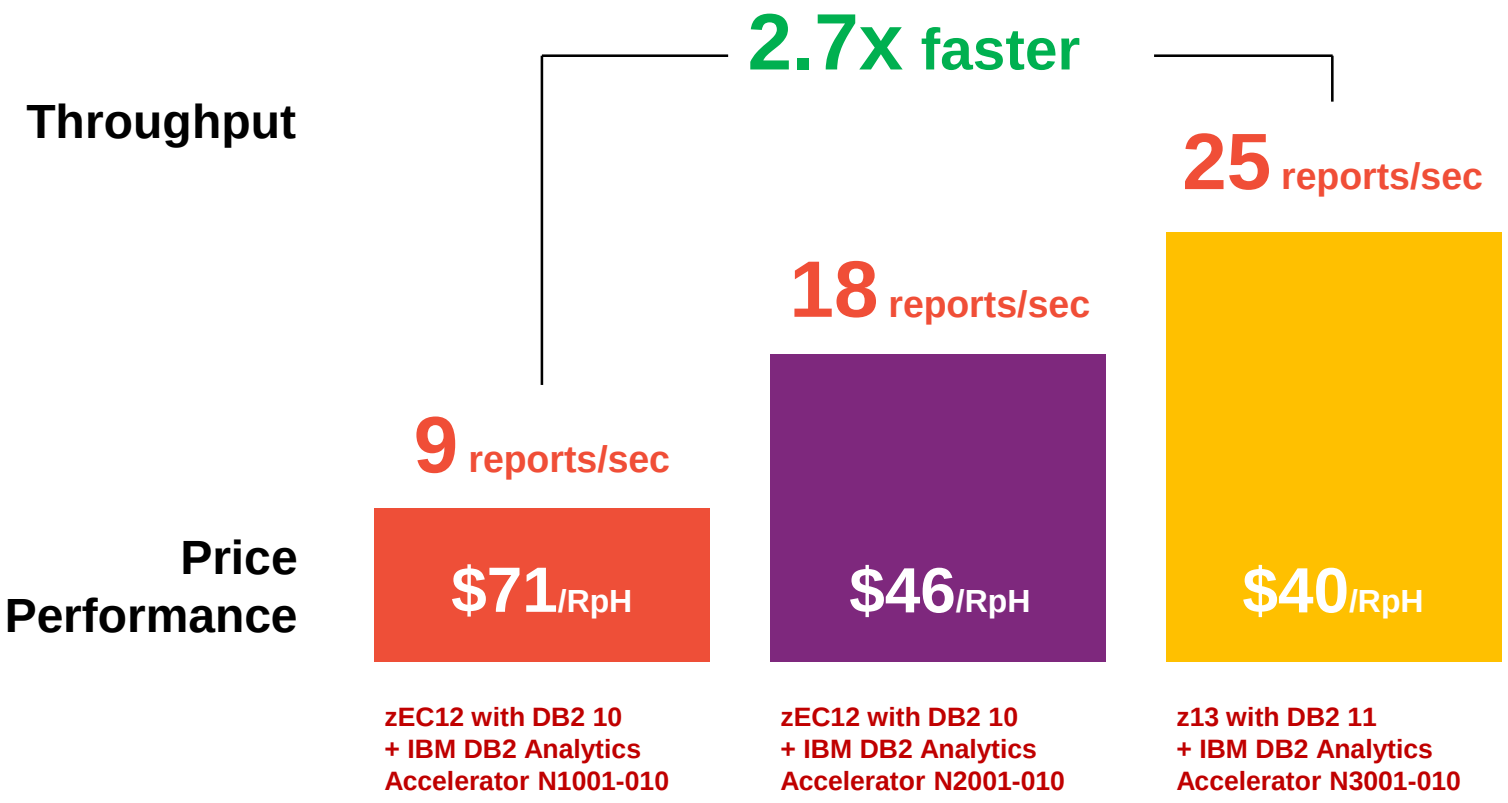

10TB BI Day Analytics Daily set of 161,166 reports, 80 concurrent users

DB2 and the Analytics Accelerator score a big win over the competition

Standalone Pre-integrated Competitor V4

$151

per Report per Hour (3yr TCA at discount)

(75% on software, 50% on hardware)

Full Unit

Estimated Workload Time*	226 mins
Reports per Hour	42,787
Competitor Full Unit (HW+SW+Storage) using discounted pricing	$6,451,161

IBM z Systems

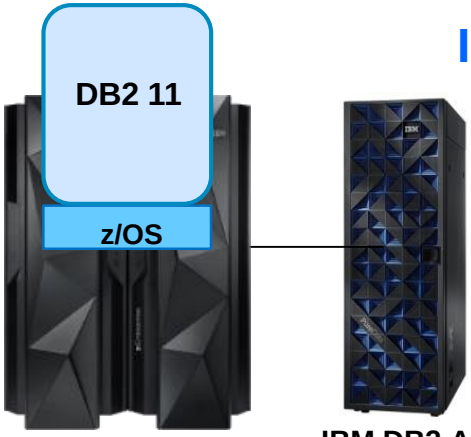

z13 with 1GP + 3 zIIP

IBM DB2 Analytics Accelerator (N3001-010)

$40

per Report per Hour (3yr TCA at no discount)

Workload Time	105 mins
Reports per Hour	92,095
z13 (1 GP + 3 zIIP, HW+SW+ Storage) + Accelerator V4.1 with PDA N3001-010 hardware	$3,652,131

2x Better performance

3.8x Better price performance

Estimated for systems compared

* Competitor Full Unit workload time estimated from Eighth Unit measurements assuming perfect linearity. Actual results will vary.
Comparing test results of an IBM zEnterprise Analytics System 9700 with an estimated performance on competitor full unit configuration (version available as of 12/31/2014), for a materially identical 10 TB BIDAY "Fixed Execution" workload in a controlled laboratory environment. BIDAY "Fixed Execution" workload measures elapsed time for executing 161,166 concurrent reports using 80 concurrent users. Intermediate and complex reports are automatically redirected to IBM DB2 Analytics Accelerator for z/OS (powered by N3001-010 hardware or Mako).. Price comparison of 3YR Total Cost of Acquisition (TCA) based on U.S. prices current as of December 31, 2014, including hardware, software, and maintenance. Used discounted pricing for competitor with 50% hardware discount and 75% software discount. Compared prices exclude applicable taxes, and are subject to change without notice. Competitor configuration: Full Unit including competitor recommended software options and features. IBM configuration: z13 platform with 1CP and 3 zIIPs with 128GB memory and DB2 Analytics Accelerator Full Rack (N3001-10) with 7 S-blades (140 Intel E5-2680v2 2.8GHz cores and 128 GB RAM), 2 Hosts (1 active – 1 passive) with 20 Intel E5-4650v2 2.4GHz cores each and 12 disk enclosures, each with 24 600GB SAS drives . Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment. Users of this document should verify the applicable data for their specific environment.

Run a complete solution – including query acceleration, Big Data, BI, and predictive analytics – on z Systems

Data Store
DB2 for z/OS

Big Data (Hadoop)
InfoSphere BigInsights

Business Intelligence and Reporting
IMB Cognos Enterprise

Predictive Analytics, Modeling, Scoring
IBM SPSS

BLU Acceleration
DB2 LUW

IBM z Systems

DB2 Analytics Accelerator

Green boxes denote Linux on z software.
Blue denotes z/OS software. Cognos runs on both.

The power of Hadoop, combined with descriptive analytics, enables businesses to have a 360º view of their customers

Hadoop:

- A framework for “distributed” storage and processing of very large data sets across clusters of Linux on z guests
- Takes advantage of massively parallel processing
- Uses simple programming models based on MapReduce

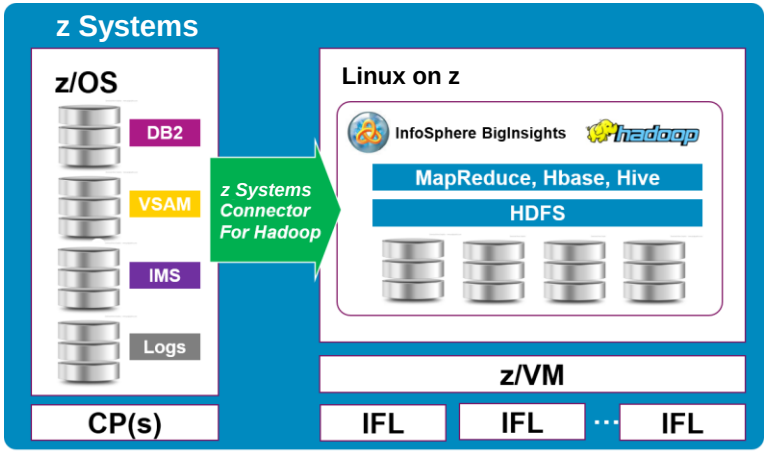

IBM BigInsights

Descriptive Analytics:

- Insight into what has happened
- Provides reports/dashboards
  - Aggregate and drill-down on data using different dimensional attributes such as by date, geography, demographics, etc.
- Visualize data using interactive charts, graphs, maps and other objects
- Runs on Linux on z and z/OS

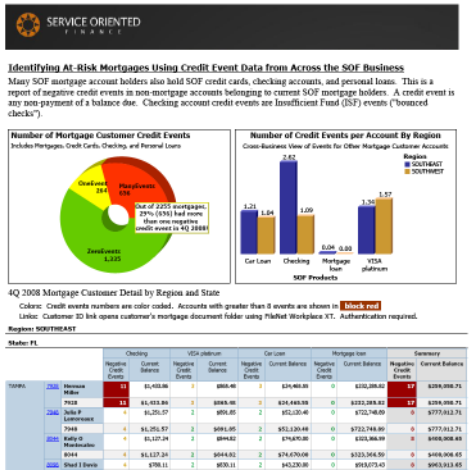

IBM Cognos Enterprise

DEMO: 360º view, from sentiment analysis combined with traditional customer data, is a critically important first step

Many businesses view this as important functionality, before getting deeper into analytics

- Use IBM BigInsights to identify good customers who have made complaints on Twitter
- Combined data from Twitter with mortgage data in the data warehouse
- Build report with IBM Cognos Report Studio to show complete customer profile

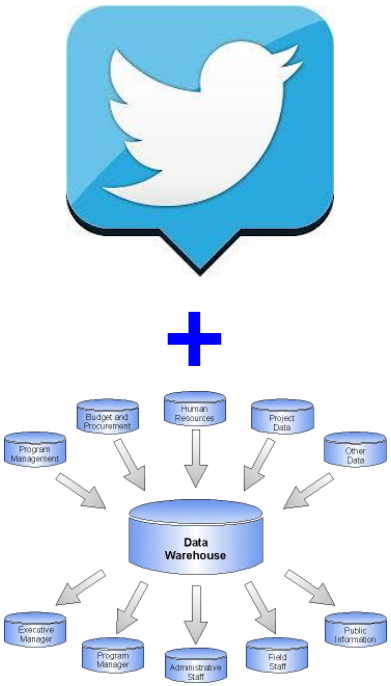

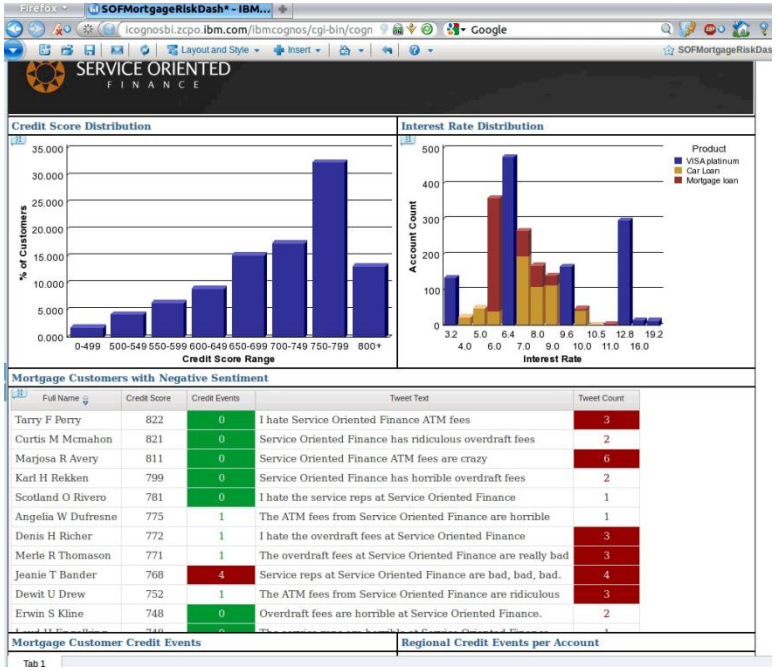

Predictive analytics truly opens up avenues for fast business insight

Predictive Analytics:

- Predicts what might happen
- Provides scores that helps in optimized decision support
  - Build models using historical data and mathematical algorithms such as clustering or classification
- Some models provide rules that can be integrated into business processes
- Runs on Linux on z

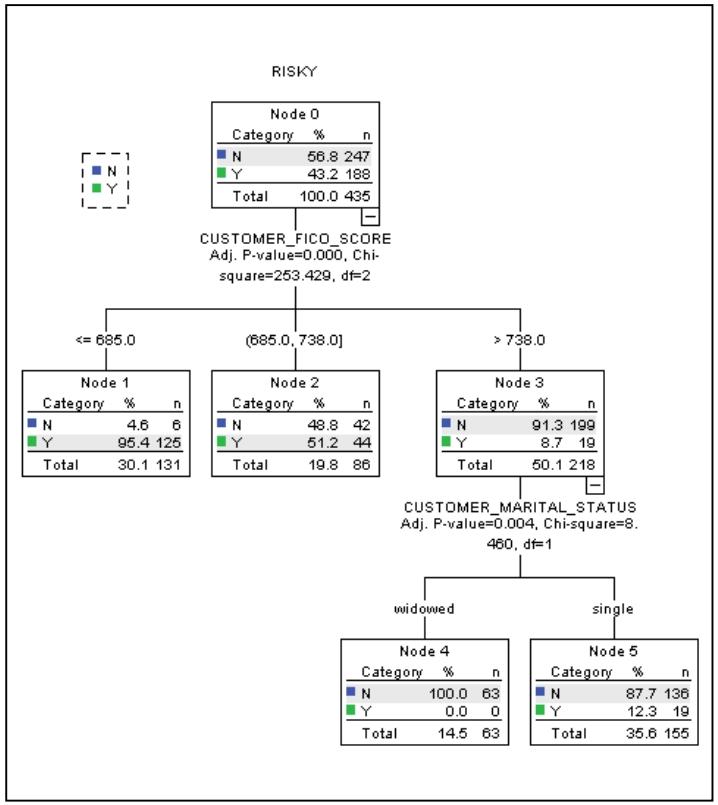

IBM SPSS Statistics and Modeler

Predictive analytics feeds into in-transaction scoring to improve business outcomes

- Instantaneous and accurate decision based on real-time information or events
- Reduce risk by putting high risk customers on “watch”
- Increase satisfaction of valued customers by providing offers using “next-best action”

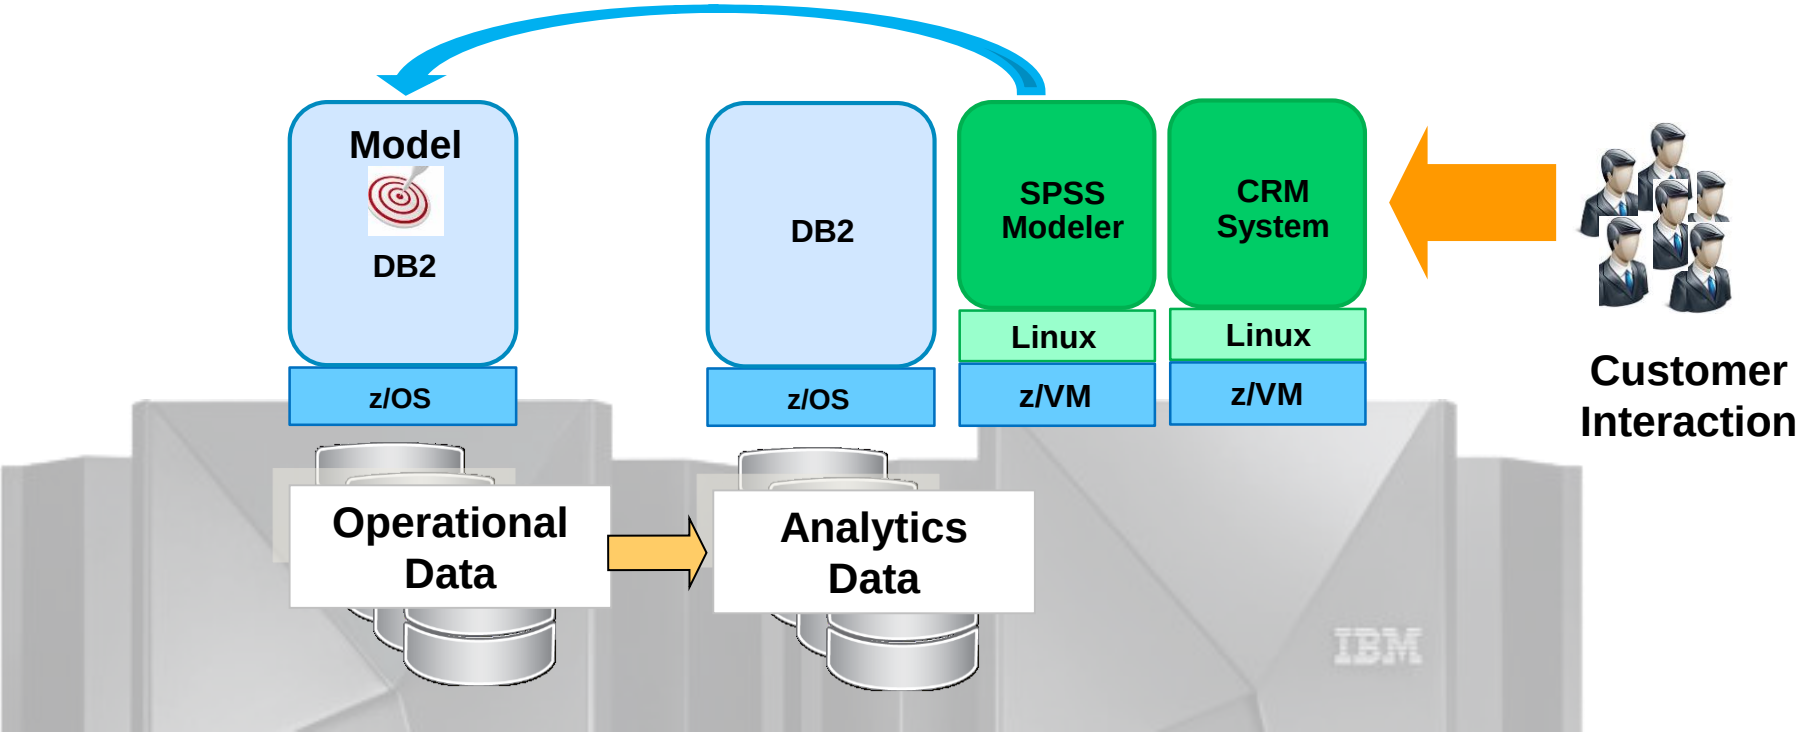

Scoring is used to determine how closely a new pattern matches a previously known pattern

Banking

Card: Use scoring to determine transaction risk based on spending history

Money laundering risk: Based on money wiring to multiple accounts keeping amount below threshold

Retail

Sales opportunity: Real-time scoring for target marketing

Government

Compliance: Score to detect non-compliant behavior and tax evasion

Social Services: Assess likelihood that individual will need multiple agency support to proactively engage various agencies to create best outcome and manage costs

DEMO: Score online banking transactions for Next Best Action and Fraud Detection

In-transaction scoring using SPSS Modeler and CICS/DB2 core banking workload

1. High value deposits with net balance between $100-$500K initiate wealth management service recommendation on welcome page
2. Multiple withdrawals within short period of time trigger fraud alert and lock the account

On-platform scoring achieves 94% cost per throughput savings

Which platform provides the lowest TCA over 3 years?

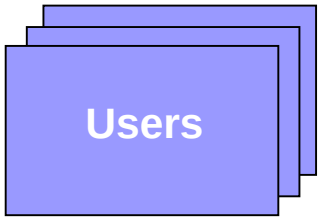

OLTP workload

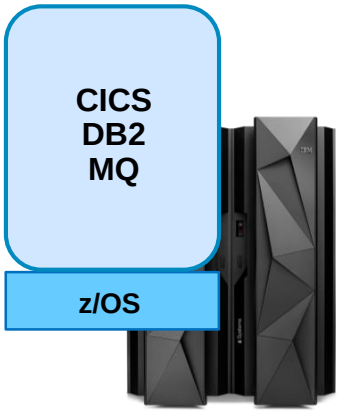

z13-705

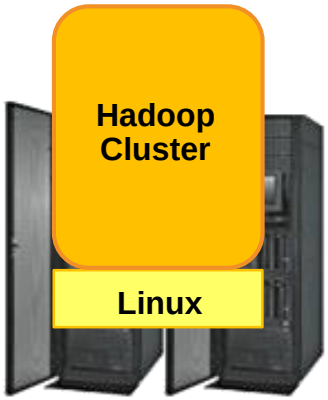

2 x Competitor x86 System
Intel E5-2650 v2 2.6GHz 16co

$6,196
per Score per second
(3 yr. TCA
no discount)

z13-704

$391
(per Score per second
(3 yr. TCA no discount)

94% Lower cost per throughput*
For systems compared

[1] Modeler scoring function only running as a UDF in DB2 for z/OS

* 3-Year TCA includes 3-year US prices for Hardware, Software, Maintenance and Support as of 12/31/2014. This is based on an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace.

Use BLU Acceleration on z Systems for analyzing data not in DB2 for z/OS

Fast Answers. Simply Delivered.

What is BLU Acceleration?

- In-memory analytic database integrated into DB2 for Linux on z Systems
- Multiple IBM innovations
  - In-memory processing of columnar data without the limitations of memory size
  - Analyze compressed data with actionable compression
  - CPU Acceleration

Analyze more data faster and more efficiently

Row-organized data can be inefficient for some analytic workloads

- Analytics queries often operate on only a small number or even a single column value across a very large number of rows
  - For example: MIN, MAX, SUM, COUNT, AVG
- Retrieving all column values is inefficient when only a small number of columns (maybe just 1) are needed

Row Organized Customer Table

	CUST_ID	FIRST	LAST	AGE	SEX
Row 1	466	Steve	Miller	49	M
Row 2	467	Pat	Smith	32	F
Row 3	478	Tina	Jones	27	F
Row...	479	Rick	Miller	42	M
Row N	481	Tom	Smith	36	M

Each colored row represents a data page

I/O

Query:

```
Select AVG(AGE) from Customer
```

466	Steve	Miller	49	M
467	Pat	Smith	32	F
478	Tina	Jones	27	F
479	Rick	Miller	42	M
481	Tom	Smith	36	M

Less Efficient!

AVG=37.2

Query only needed 1 column to compute average age, but I/Os were required to retrieve all columns for all rows

Column-organized data is better suited and more efficient for some analytic workloads

- BLU Acceleration organizes data into columns
- Column values for many records are combined into “pages” and stored on disk
- One I/O operation (to disk or RAM) can retrieve a column value for many rows
- Great for analytical workloads
  - When SPECIFIC columns are accessed for MANY records
  - No indexes required – columns are essentially “self indexing”

Column Organized Customer Table

CUST_ID	FIRST	LAST	AGE	SEX
Col A	Col B	Col C	Col …	Col N
466	Steve	Miller	49	M
467	Pat	Smith	32	F
478	Tina	Jones	27	F
479	Rick	Miller	42	M
481	Tom	Smith	36	M

Each colored column represents a data page

Query:

```
Select AVG(AGE) from Customer
```

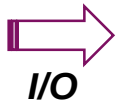

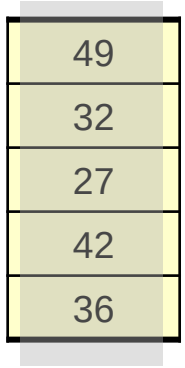

AVG=37.2

More Efficient!

Query needed 1 column to compute average age, I/O retrieved just one column

DEMO: Columnar data can speed up IT analytics

- Using BLU Acceleration IT analytics
  - Load PerfSVM System Summary data (at 30 second intervals) into BLU and show power of columnar store vs row store for query doing average of Pct. Busy (CPU) and SSCH+RSCH/s (I/O)

Interval 00:00:04-22:13:34, on 2014/12/18

	<------- CPU -------->				<--Users-->		<---I/O--->	
		<--Ratio-->					SSCH	DASD
Interval	Pct			On-	Log-		+RSCH	Resp
End Time	Busy	T/V	User	line	ged	Activ	/s	msec
>>Mean>>	.2	1.41	.7732	16.0	16	5	420.0	.2
22:01:04	8.3	1.05	.9899	16.0	16	6	858.4	.7
22:01:34	5.9	1.07	.9848	16.0	16	5	883.1	.7
22:02:04	5.0	1.10	.9594	16.0	16	4	944.9	.8
22:02:34	6.9	1.02	.9884	16.0	16	5	486.4	.2
22:03:04	7.8	1.01	.9898	16.0	16	5	480.9	.2
22:03:34	7.9	1.22	.9096	16.0	16	5	7473	.2
22:04:04	8.0	1.27	.8971	16.0	16	5	8785	.2
22:04:34	7.9	1.28	.8908	16.0	16	5	9259	.2
22:05:04	7.7	1.30	.8781	16.0	16	4	12305	.2
22:05:34	8.4	1.36	.8612	16.0	16	6	16055	.2
22:06:04	8.3	1.31	.8740	16.0	16	5	12910	.2
22:06:34	7.3	1.24	.8945	16.0	16	5	9543	.2
22:07:04	6.6	1.03	.9814	16.0	16	4	767.1	.3
22:07:34	11.6	1.12	.9515	16.0	16	5	4559	.4
22:08:04	9.9	1.13	.9440	16.0	16	4	4135	.4
22:08:34	7.6	1.27	.8953	16.0	16	5	8569	.2
22:09:04	7.4	1.22	.9091	16.0	16	4	6944	.3
22:09:34	6.6	1.19	.9234	16.0	16	5	4558	.5
22:10:04	6.2	1.15	.9385	16.0	16	4	3522	.4
22:10:34	4.3	1.08	.9790	16.0	16	5	1022	.8

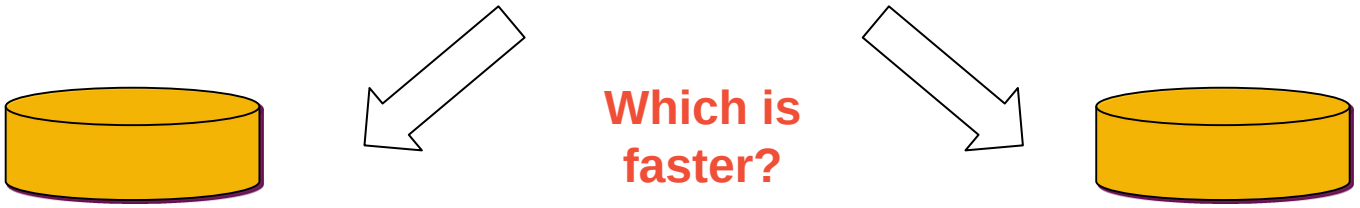

Analytics processing on z13 is simpler and faster, laying the foundation for digital business growth

SIMD technology
Speeds up processing for compute-intensive analytics workloads

SMT technology
Improves response time and throughput of data-driven workloads

10 TB Memory
Improves data buffering and in-memory analytics

Faster I/O
Reduces data transactional latency

2x Compression
Reduces CPU usage, reduces storage requirements, increases memory efficiency

z Systems – an exceptional System of Record and a first-class System of Insight

60+% zIIP offload for z13+DB2 11

39% Higher throughput for z13+DB2 11 than previous version

System of Record

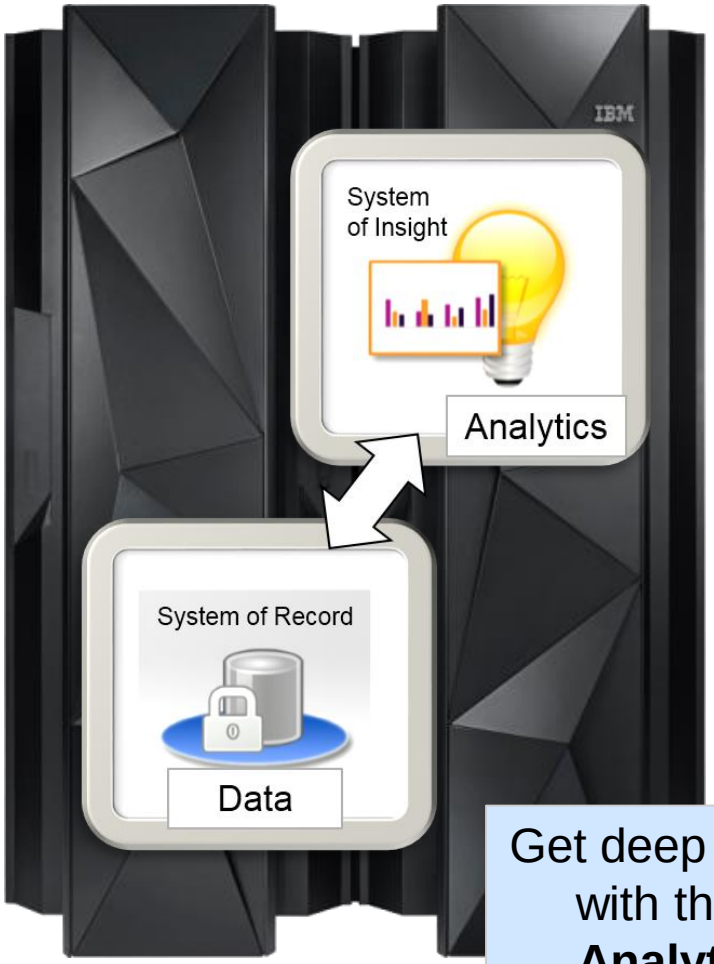

System of Insight

3.8x Better cost per workload for z13+ Analytics Accel. than competition

94% Lower cost per throughput with BigInsights on z

Get deep discounts on software with the **IBM zEnterprise Analytics Solution 9700**